Remarks at Dedication of Kraken Machine at the National Institute for Computational Science

April 3, 2008

Daniel E. Atkins
Director, National Science Foundation,
Office of Cyberinfrastructure

The machine we are dedicating today has been named Kraken. According to a German-English dictionary Kraken mean octopus -- and so I initially thought about 8 core chips or perhaps 8 equipment racks radiating from a central hub. But then I realized that this machine's namesake is the monstrous Kraken of sea sagas that engulf and crush large ships -- just as, this Kraken will engulf and crush large problems that confront science/engineering in service of human well being. -- a very good naming choice. Ironically contemporary HPC has vividly brought Kraken back to life in John Knoll's Academy Award winning special effects in the Pirates of the Caribbean series.

Events like this help me realize how fortunate I have been to have a professional career spanning most of the modern computer age. It also made me feel old. As a graduate student at the University of Illinois I played a small part in one of the first transistorized computer projects - Illiac 2. (I played a much larger part in Illiac 3 and 4) Illiac 2 was built from discrete transistors costing $80 a piece in (1960 dollars), occupied a room larger than the Kraken facility, and provided much less compute power than *one* of the AMD CPU cores used in Kraken. In those days the word "core" referred to a ferrite donut providing one-bit of storage in the computer's main memory. I could have never imagined that core would become the name for a complete central processor in today's massively parallel multi-cores machines. I expect that in my life time we will have single chips containing more CPU cores than the early machines had memory ferrite cores in their main memory.

Kraken is the second leadership computational resource provided under the NSF Office of Cyberinfrastructure's "Track 2 initiative" and likely the first machine in this series that will hit a petaflop of peak performance. This award was made under a new acquisition policy at NSF that emphasizes a holistic notion of "impact on science" as the overarching criteria for making HPC awards. This science-driven award process involves close cooperation between the experts providing high performance computation resources, and the science and engineering research communities who now use these resources to pursue transformative discovery and learning. Benchmarks and model programs performance as surrogates of the algorithm-types important to these communities are expected to be included in proposals and those winning awardees are expected to meet or exceed them. The performance criteria are multi-faceted including not only the raw performance of the machine, but also effective education and outreach mechanisms, institutional competency and commitment for service to the national research community.

These Track 2 awards are also made under what NSF calls a “cooperative agreement.”, The awardee and NSF cooperate in developing a plan of action, consistent with the submitted proposal, with mutually agreed performance criteria and milestones along the way. We then cooperate in realizing the plan and mitigating risk since these are leading edge machines. In the case of Kraken, I was very pleased to learn that this spirit of cooperation has extended generously to the relationship between University of Tennessee and Oakridge National Labs in the National Center for Computational Science and with CRAY and AMD. I have been told that you have all conceived this as a collective opportunity and have gone the extra mile to insure that Kraken is a dramatic step forward in offering computational resources for the open science research community... and there is a large amount of science bang for the buck. The dedication of Kraken today and all that has lead up to it has set a very high standard in both process and performance for the entire future of HPC provisioning by NSF. You have set a very high bar.

I am also delighted about the opportunity that this award provides for tangible cooperation between the NSF and the DoE Office of Science. Without question, the combined investments of the Office of Science and the NSF are providing our Nation’s researchers the world’s most pervasive and powerful HPC facilities by orders of magnitude.

Computational and data processing resources such as provided by Kraken are now absolutely essential to frontier research in almost every field of science and engineering.   Computationally based-modeling, simulation, and prediction across huge scales of time and space are the norm in exploring emergent complex systems such as embedded in social and physical global climate change models, understanding and predicting turbulent weather, understanding of the universe, the fundamentals of life, and global patterns of social interaction. Computational science is now well established as a third mode of discovery complementing theory and physical experimentation. Science and engineering research is also increasingly data-intensive due in large part to enormous breakthroughs in sensing and observing technologies. Kraken and associated visualization services are also a tool for finding “knowledge needles” in enormous “digital data haystacks” and then integrating these needles into spires of insight and breakthrough discovery.

This is a very exciting day for 21st century research and education. This day has been made possible by the cooperative efforts of many people and numerous organizations.

In the spirit of Kraken being a team effort, I also acknowledge and thank Steve Meacham and Jose Munoz, members of the NSF OCI team, for are responsible for the OCI HPC programs.

Although we focus today on this new facility, it is important to realize that it is part of a larger network of resources known collectively as the TeraGrid. The TeraGrid is a consortium of network resources increasingly integrated as a collective resource for the Nation’s research communities. Kraken brings a quantum leap in computational power and memory capacity to the TeraGrid and is a major step forward towards the next goal of computing at sustained rates at the petascale level. The TeraGrid enables researchers to link and use multiple computers with diverse architectures, data bases, and visualization facilities independent of where they are and who operates them. Increasingly TeraGrid resources are being used transparently through web-based science gateways that are tailored to both research and education in a growing number of fields. A high school student can design and analyze nano-structures using the same tools as the researchers, or even make original contributions to astronomy through data mining in a digital sky. These science gateways are entrances into virtual communities that enable participation in authentic discovery and learning independent of time and place. They open new opportunities in science for what we do, how we do it, and who participates.

Thank you all for your past and future contributions to enabling discovery and innovation essential to the well-being and vitality of our world. Congratulations and all the best. If now is the time for applause, let me ask that it be hearty applause for the entire team that worked so hard to bring us to this day of celebration.

Thank you.