Remarks by Dr. Raymond L. Orbach
Under Secretary for Science
U.S. Department of Energy
at the 2007 INCITE Awards Announcement
Council on Competitiveness
Washington, DC
January 8, 2007

It is a great privilege and joy to be back in this room to talk about this, in my view, quite remarkable announcement that we will be making today. I want to thank Debra Wince-Smith and the Council on Competitiveness for their support over the years.

When back in 2003 we announced 5.5 million CPU hours, we thought it was a big deal. We were dealing with computers whose peak speeds were at least an order of magnitude less than what we have today. And at the time when people talked about a petaflop, which is a measure of speed that we all are striving for, I thought it was just poppycock, that I wouldn't have to worry about that while I was in the Department of Energy, but that was the goal.

If the budget proposals of the Department of Energy are funded, we will have a petaflop in 2008, and it will give the United States leadership over every other country in the world. We are the leaders in high-end computation and we intend to continue.
One of the reasons that we have been successful is that these machines are not simple to work with. In order to make them work there is a match between people and machines, and I will come back to that in just a moment.

In parallel, therefore, with the actual physical construction of these large scale facilities, we have a program in the Office of Science called SciDAC, which brings together computational scientists, mathematicians, and the scientists themselves as teams to begin to look at how to take the physical or biological sciences problem that we are trying to solve and put it on these machines to get it to run efficiently at these speeds. Sometimes just by working on the codes themselves you can increase the speeds by factors of two or three.

What we in the DOE Office of Science have done is optimized the relationship between man and machine, between person and machine, and that is what makes us special. If you look around the world, there are other fast computers. They are not as fast as ours, but ours work, and they work on the class of problems now that is spanning not only science but industry.

Initially when we started, we had four winners of our INCITE competition. They were all from science. I am delighted that today we will announce five new winners from industry which, together with the four from last year which are continuing on, will mean that something like 10 percent of the winners in a peer-reviewed competition are from industry. To me, this is the beginning of what will truly be an exciting development for the United States.

We are going to turn these machines into machines of competitiveness. This may be at the heart of what we call the virtual prototype, shortening time to market, reducing the cost of development, and, as we will hear today, actually moving into new directions, things we can't do right now in the laboratory. That is one of the great promises that high-end computation offers.

It is often said that science is based on two pillars, namely experiment and theory. In fact, high-end computation, especially through simulation, is the third pillar of science. It is actually a staple support as well. It gives us the ability to simulate things which cannot be done either experimentally or for which no theory exists -- simulations of complex systems which are simply too hard or so far are not amenable to analytic approaches. So you need to think of high-end computations as a third pillar for discovery, and it gives the United States an advantage in industry, in science, and an instrument of discovery that I believe cannot be achieved elsewhere.

It is with obvious enthusiasm and excitement that I am here today to talk about 95 million CPU hours that we will be allocating and four computational platforms of very different architectures: one at Oak Ridge National Laboratory, one at Argonne National Laboratory, one at Pacific Northwest National Laboratory and one at NERSC at Lawrence Berkley National Laboratory.

You may ask, why have four large scale computational structures? The answer is, each of them is different. We are exploring a region that has never been looked at before. We don't know which architecture is best suited to which scientific problem. We are going to learn which one is best – or for that matter, what the flexibility is amongst the different architectures – and indeed the architectures themselves are changing as the demands of the computations are becoming ever more severe.

What you are seeing is a change in the nature of the way that computation is addressed. All of us have laptops. All of us work there in front of the screen. But we are now talking about teamwork. If you look at the INCITE winners this time, you'll find that, for example, principal investigator Robert Harrison, from Oak Ridge National Laboratory, has 18 collaborators. That is a team of 19 people. If you look at the National Center for Atmospheric Research (NCAR), principal investigator Warren Washington has 12 collaborators.

What we are seeing now is a transformation in the way we do computation. You need to think of high-end computation the same way you think of high energy physics or of nuclear physics, or of large scale teams coming together to deal with these issues. People that bring individual skills on the mathematics side, on the science side, on the computational side.

It's a new sociology for science, for achievement, for discovery. All this is happening in real time, and it's wonderful that the United States of America is the world leader. We intend to remain at that forefront. What you are seeing today is in some way an opening salvo to declare that this is the future of our country in innovation and scientific discovery.

With that, let me thank you very much for coming in, and give it back to Debra to begin the program.