Council on Competitiveness First Annual High Performance Computing Users Conference
Washington, D.C.
July 13, 2004

John Marburger
Director, Office of Science and Technology Policy
Executive Office of the President

Thank you for inviting me to speak today about high performance computing. Thirty years ago I was trying to simulate high intensity laser beams with the most powerful computers then available. I was consulting at Lawrence Livermore Laboratory, and had access to their Control Data Corporation machines (CDC 6600 and 7600), at that time the world's fastest computers. Everyone knows that the power of today's workstations, routinely available to professional scientists and engineers, vastly outstrips those monster machines. We could only dream in those days of having such capabilities, and even in the wildest of those dreams we could not envision the environment that has grown up to support high performance computing today. Our vision of what problems might be attacked was severely constrained by the weakness of known technology.

This issue of vision has turned out to be important in the history of computing, and I would like to reflect today on how visions have affected our choices over time, and on our current vision of what we need to do next to sustain our leadership in computing-related enterprise. On one hand, lack of vision has lost many opportunities to make faster progress. On the other, unrealistic visions of the opportunities that lie ahead have led to many failures, most spectacularly the bursting of the dot-com bubble in the last decade. But some visions have had revolutionary impacts on society, and we need to learn their lessons.

Most in this audience know about the famous remark by Thomas J. Watson, Sr. who may or may not have said "I think there is a world market for maybe five computers." If he did say that, it was in 1943, the year the University of Pennsylvania signed a contract with the Army Chief of Ordinance to build what became the ENIAC, the world's first electronic computer, badly needed for ballistic calculations. Recall that this was the time of major artillery engagements in the European theater of World War II, and the beginnings of the Manhattan project. A year later, Princeton mathematical genius John von Neumann became interested in ENIAC and proceeded to develop his groundbreaking idea for a device that could store different programs to perform a variety of computing tasks. His interest, and the early impetus for vastly improved automatic computation, was impelled by wartime needs.

Immediately after the war, in 1946, Watson's 32 year old son, Thomas J. Watson, Jr., learned about ENIAC and told his father that "We should put this thing on the market! Even if we only sell eight or 10, we'll be able to advertise the fact that we have the world's first commercial electronic calculator." Watson, Jr. energized the company, became its CEO ten years later, and in 1964 introduced the System/360 concept which launched the world's first operating system that made it possible to transport code from one machine to another within the 360 family, vastly reducing the cost associated with hardware upgrades. The strength of Watson's conviction can be judged by the $5 billion he was willing to spend to develop

System/360 (that's in 1964 dollars!) His vision made such an impact on American business that Forbes lists Watson, Jr. among the top 20 most influential businessmen of all time.

Despite this early recognition of the importance of issues *other than hardware* in the development of what we now call information technology, IBM failed to anticipate how quickly the low end of the hardware market would become "commoditized" and how crucial operating systems would be in the ensuing spread of computing to mass markets that began in the mid-1970's and accelerated sharply in the '80's. The rise of Microsoft on IBM's PC platform is legendary – the product of yet another vision vigorously pursued by Bill Gates. Another important early vision was Gordon Moore's law on the growth rate of hardware capability, stated in 1965, and eerily fulfilled ever since. It is important to understand the nature of "Moore's law" – it is not a law at all but a bold challenge to two generations (so far) of scientists and engineers who applied themselves relentlessly to make it true.

It was not only computing that Moore's "law" made increasingly cost-effective, but also the whole universe of information processing applications, from the input/output needs of the computers themselves to the increasingly complex switches at the core of telephone system networks. A host of other applications sprang up around the electronics produced initially for the computing market – all the digital toys and devices such as cell-phones, cameras, audio and video products, scanners, bar code readers, printers, fax machines, and a huge recreational industry centered on interactive games. The availability of these products to the general public, coupled with the ability to link them through the Internet, is changing the world. Many visionaries, among them the science fiction author William Gibson (a founding father of "cyber-punk" fiction), perceived what might happen in this new information-rich environment. Unfortunately it was difficult even for Wall Street to disentangle the possible from the merely visionary, and we suffered the economic phenomenon we now call the dot-com bubble.

As the information technology based economy snowballed toward its present state, we all learned the lesson that Microsoft taught IBM. Hardware is a means, not an end. What is important is gathering data and making it available for answering questions and accomplishing other tasks easily and efficiently. The generic hardware and operating systems that responded to the needs of business and personal productivity applications easily handles most of the spectrum of today's information technology activity. For the most part, the commoditized hardware has taken care of itself, advancing magically along the curve of Moore's law, dropping in price per performance, and enabling a cornucopia of new business models that in this post-dot-com era finally seem to be working.

The reason we are here today is that something is happening in the world of computing that is about to alter this picture. The hardware we take for granted is not capable of doing all that we should like to do, and we know today that much more power is potentially available to us. If that potential is realized it will once again transform ways of doing business. In today's globally competitive economy we cannot afford to leave this opportunity to others.

Thomas Watson Senior's apocryphal statement about the world market for computers was not completely wrong for his time. Certainly he was reacting to the expense of building and maintaining ENIAC, and the awkwardness of replacing burned out vacuum tubes twice a day. And probably he was reacting to the highly specialized intentions its wartime designers had for its use. Johnny von Neumann would soon identify what today we call "grand challenges" for electronic computing, including weather prediction and the realistic simulation of hydrodynamic flow. (NOAA's Geophysical Fluid Dynamics Laboratory at Princeton is a von Neumann legacy.) These and similar more contemporary "grand challenges" are applications that until now only governments could afford. They define the case for "supercomputing" and distinguish it and its small market from the commoditized machinery on which the mainstream information technology revolution is currently advancing.

What is happening today is that the gap is widening between the known potential of high performance computers and the capability that is actually being realized with commercial off the shelf hardware and associated software. The miracle of Moore's law is bringing a new order of powerful computing within reach of a much larger market than a few governments of developed countries, but that market has not yet grasped the potential. We are approaching a "tipping point" beyond which entirely new applications of computing will bring a new wave of transformations in our industrial ways of life, and further disrupt older ways of doing business. Today, however, we are on the near side of this tipping point, and more than ordinary vision is required to perceive the way.

I give the Council on Competitiveness, DARPA, and the Department of Energy much credit for seizing the initiative to cultivate this vision. The United States Government, its laboratories and contractors, have long been the primary customers for high performance computing, but its power can become available to a much wider base of industrial and commercial users. Like other transforming concepts, high performance computing extends the horizon of applications beyond the range of our imagination. We need to spread the word about the new capabilities and build confidence in the new visions to motivate public and private investment in them.

Others in this conference are giving a glimpse of what may be possible – dramatically shortened intervals between ideas and production, individually tailored products from cars and clothing to pharmaceuticals and prosthetics, large reductions in the cost of complex operations in manufacturing, chemical processes, and material transport. My role today is to assure you that the promise of high performance computing is a priority in this Administration, and that we are doing everything we can think of to mobilize the relevant Agencies to release its potential to stimulate innovation and maintain competitiveness throughout broad sectors of the nation's economy.

Networking and information technology R&D (NITRD), currently funded at over $2 billion per year, is one of the President's highest inter-agency research priorities. My Office of Science and Technology Policy (OSTP) manages a formal interagency working group on NITRD, and a National Coordination Office, whose Director, Dave Nelson, is here today. High end computing (HEC) is one of six major areas of formal interagency coordination within the NITRD program. Last month I met with NITRD agency principals to reinforce the priority of HEC investments for the administration in budget requests for FY06. We have formed a High End Computing Revitalization Task Force (HECRTF) and have worked to develop an interagency R&D roadmap for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and recommendations relating to federal procurement of high-end computing systems. Ten different agencies participated actively in preparing its final report, which was released in May.

These are highlights of a deliberate Administrative initiative to capture the power of high performance computing for a widely expanded set of applications in both public and private sectors. This conference is an important step in bringing the message of high performance computing to a broader audience, and lowering barriers to the realization of our collective vision of a new era of competitiveness through leadership in computing. Thanks again to the Council on Competitiveness for initiating the conference, and to all of you for making it successful.