
HPC@IDC USER FORUM JULY PLANNING MEETING

7/24 and 7/25/00 Meeting Notes

Participants (53 in total):

David Bailey, Lawrence Berkeley Lab
Jim Bidwell, Boeing
Terry Blanchard, NAVO
Paul Buerger, OSC
Kurt Carlson, ARSC
Tom Crimmins, ARL
Candace Culhane, NSA
Larry Davis, DOD MOD
David Filkin, Dupont
Norm Glick, NSA (for George Cotter)
Davis Kahaner, ATIP
Jim Kasdorf, PSC
Roland Kaluzniacki, NSA
Bill Kramer, NERSC
Jose Munoz, DOE/ASCI
Charles Nietubicz, ARL
Jim Pool, Caltech
Vince Scarafino, Ford Motor Co.
Barry Sharp, Boeing
Suresh Shukla, Boeing
Horst Simon, NERSC
Erich Strohmaier, University of Tennessee
Al Stutz, OSC
Vendors: Compaq, Cray, Etnus, Fujitsu, HP, IBM, SGI, SRC, Sun
IDC: Debra Goldfarb, Earl Joseph, Chris Willard, Steve Conway

STEERING COMMITTEE MEETING: 7-24-00

Introduction (Earl Joseph, IDC):

- Plan is to hold Steering Committee meeting before and after each User Forum Planning Meeting. The Steering Committee will provide input on User Forum meeting agendas; access how well the meeting went; make suggestions for future meetings; guide the research agenda; etc.
- Central issue for today's User Forum meeting is performance measurement. We'll hear how several more users approach performance measurement, a proposal from Erich Strohmaier, and move toward consensus on what to do about this issue.
- We've had multiple requests from HPC users wanting to attend individual meetings but who are not yet ready to join. We'll charge them \$2,000/person for the September meeting, \$1,500/person for each planning meeting.
- We've also been asked how we would approach large studies that are beyond the scope of the User Forum and IDC, as for example the development of a new performance metric. Our model will be to hire outside resources, where appropriate, to do this kind of thing for the User Forum on an ad hoc basis, and collect funding for the project as a separate activity.
- Update: We've postponed the first meeting of the Japan User Forum till the January-February timeframe in order to get past SC2000.
- We want to make sure this organization is really central to members' needs, and to create a "User Agenda". (Earl asked attendees to submit the Blue Form handouts to help identify high-priority issues beyond the performance metric.
 - (Boeing): I'm interested in what supercomputers will look like after 2004, also open source issues such as security.

- (NSA): We have a set of 12 benchmarks that show the types of codes we run without sharing/disclosing actual production codes. Our benchmark codes are very relevant for system requirements. We share the benchmark codes with our vendors.
- (DOD MOD): We're looking at a wide variety of benchmarks.
- Comment: We're interested in lists like the Top500, Top100, etc. These bother me because they can have tremendous impact. We need another list that's more a reality list, because the current list can force us to be driven into bad directions. We need a list we can point to that shows our management that we're a leading user organization.
- Comment: So do we. That's the reason we joined.
- Comment: We want of a list based on the overall capabilities of actual computers.
- (Ford): The Top500 list isn't inaccurate, just irrelevant for real computing needs.
- (NERSC): We lack an honest, neutral assessment of Japanese supercomputers for climate and fusion studies. Our users feel behind in these fields. Second, most of us in this room represent traditional supercomputers, but the real pressure is from PC clusters. Third, we're interested in knowing more about how to approach open source.
- (Ford): Interoperability is missing from open source and we need to fight for this.
- (NSA): My opinion, not NSA's, is that I'm worried about expanding this User Forum to include the low end and clusters because I'm concerned we'll lose focus on the high end. We high-end users are a minority and would lose out in a bigger forum.
- (DOD MOD): I agree we should avoid broadening the scope of the User Forum too much.
- (Debra, IDC): We're not trying to make this very broad. We're focusing on supporting scientists and engineers. But clusters are increasingly important and have been taking on much of the traditional high-end work. The needs of high-end users will be the central focus of the User Forum.
- (Earl, IDC): We're talking about expanding the focus to look at the whole spectrum of computers used by HPC users, not just the capability portion, or we may miss important trends. The primary focus will always be technical HPC.
- (Ford): We should concentrate on capability. Clusters do capability work on certain things, but not on things that have never been done before. Clusters can't do that.
- (Debra, IDC): We're not going to address e-commerce, etc., just technical computing. We need to hear from many people. We need to decide what's relevant and why—for example, grid computing. Let's put out definitive information on what works where—for example, where grids and clusters don't work, despite what's in the trade media. People can use these results with their managements. We look at applications profiles—where people are using a system, the kinds of efficiencies they're getting. The IDC concept is "strategic fit," a taxonomy of where you'd run various types of applications. Let's get some rigor and analysis into this.
- (Earl, IDC): We're looking to have 60-75 people at the September meeting, with demographics similar to today's meeting. Many sites have mixed capability/capacity workloads, so we don't want to limit User Forum membership to capability only.
- (Earl, IDC) On vendor specific Customer Advisory Board meetings. The idea to do NDA vendor meetings comes from our first meeting. Vendors want guidance about what to do at the high end. Three vendors beside Cray want to hold meeting too: IBM, SGI and one other. The intent is not selling, but discussing technologies and issues related to the vendor companies.
- We want to discuss how the Steering Committee should function. People have told us they want to be involved, but want IDC to take on the logistics: organizing and facilitating meetings, etc. We feel the Steering Committee should formulate an issues agenda for the User Forum.
 - (DOD MOD): Clarify who can be on the Steering Committee.
 - (Earl, IDC): To date, anyone user member except vendors.

- (DOD MOD): I don't like the idea of electing the Steering Committee. We may want to limit it to a workable size sometime in the future, or just continue letting any user participate.
- (Boeing): It is important to have representation from various segments of the community.
- (Ford): Committee size isn't an issue if we can keep focused on the issues. If we're pulled in different direction, for example if some people want to focus on PCs, then we have to look at the structure of the Steering Committee.
- (Debra, IDC): Are people comfortable with the division-of-labor between users and IDC?
 - Comment: Can it be formalized?
 - **(Earl, IDC): We will draft a division-of-labor.**
- (Debra, IDC): What about the creation and management of SIGs?
 - Comment: We did some of that last time, didn't we?
 - (Earl, IDC): Yes. We started last time with two: high-end, and clusters. Feedback was that the high-end one made sense, the cluster one would be better handled as a study group to review and guide IDC research.
 - Comment: How large is the User Forum now?
 - (Earl, IDC): 65. We expect 80-90 by September. We're on target for growth. We're not going after multiple 100s of members. We don't want to lose focus on technical HPC. We have a list of close to 1400 HPC names and are creating a database to identify all high-end users around the world. Our top goal is to make the HPC industry healthy. Two-way sharing of information promotes this goal. We recently formed a Performance Advisory Group (PAG). No vendors are in this SIG. The PAG is about ¾ users, ¼ outside industry experts, e.g. Tony Vacca has been asked to join.
 - (Debra, IDC): Do we need a SIG on open source?
 - (Ford): We may have our hands full with the current three SIGs.
 - Comment: What about a SIG on large special studies?
 - (Earl, IDC): IDC does many of these. Usually, vendors fund or help fund them. User Forum members automatically get a summary report and help set the direction for studies. The other model is where IDC doesn't have the expertise, for example on performance benchmarking. In this model the vendors and some users will help fund the study, and we will use outside resources to carry it out. IDC acts as an impartial body and also provides commentary and analysis on the findings. The Steering Committee will help set the direction for this research and review draft results. Some of you have said you want to merge your benchmark development into the IDC effort – and this would be great. It's a bit different in Japan, where the main reason for hpc@idc for many members is that customers don't know what to buy next. Please think today about how the Steering Committee should work. We want to discuss this further in tomorrow's Steering Committee meeting and at the September meeting.
 - (NERSC): Let's not rely on PITAC to do much for us. They no longer seem concerned about the high end. I'm glad IDC's taking leadership on this.

USER FORUM PLANNING MEETING: 7-24-00 (53 ATTENDEES)

- **Welcome (Debra Goldfarb, IDC) and Introduction (Earl Joseph, IDC):**
 - First Customer Advisory Board meeting is tomorrow, Cray. Three other vendors are interested in doing these vendor specific NDA meetings.
 - Over lunchtime, vendors will present new developments. No sales pitches.
 - The Green Sheet is the draft September agenda, please make suggestions for improving the meeting.
 - Please fill out the Blue Sheet to prioritize the issues the User Forum addresses.
 - Erich Strohmaier will talk about his performance white paper today.
 - Also in front of you are the minutes from the April meeting.

- (Question): Will there be a website? Earl: Yes, it's at www.idc.com/hpc. It's totally public at the moment. In the future, User Forum members will have a secure area.
- Steering Committee meetings are being held in conjunction with User Forum meetings, both before and after each meeting.
- We formed a Performance Advisory Group (PAG) to address the performance measurement issue.
- We want to note the recent passing of our industry colleague Dick Sherman.
- We have added the ability for non-members to attend meetings -- \$2,000 per person for those who want to attend the annual meeting.
- **Introductions -- Users identified issues of importance to their organizations:**
 - (NERSC): Open source; PC clusters; cost-benefit analysis of Japanese vector machines.
 - (Boeing): What should future supercomputers look like in 3-5 years? Users and vendors are both confused. Something's missing today. This group can help define it.
 - (OSC): Performance metric; also making hardware more seamless.
 - (DOE/ASCI): Next-generation architectures; open source.
 - (DOE/Office of Science): Performance; ease-of-use.
 - (Boeing): General-purpose supercomputers; future of vectors.
 - (Boeing): Imbalance between vector and scalar capabilities. Scalar lags, forcing us to have multiple machines.
 - (Caltech): We purchase machines before they exit, so benchmarks aren't important for us.
 - (NSA): PITAC isn't doing much for the high end. This community needs to do it .
 - (NSA): Shared memory programming models with low latencies. MPI isn't low latency. Lack of new HPC architectures.
 - (NSA): We need to keep stressing performance; also end-user ease-of-use.
 - (ARL): We're interested in a real Top500 list, not peak performance. We need to make the case for the mission-critical nature of HPC.
 - (NAVO): Are clusters real? What can they really do? We need to look at applications performance, not just system performance.
 - (ATIP): How to engage the community outside the U.S. to help us.
 - (DOD MOD): Effective mapping of applications and algorithms to architectures.
 - (Dupont): Performance of third-party applications on affordable systems.
 - (IDC): Market and user demand for various architectures.
 - (OSC): Ease-of-use software, not just hardware.
 - (Ford): Grand challenges -- pushing the capability envelope.
 - (PSC): Matching applications to architectures.
- **"Why Worry About HPC Performance Metrics?" (Earl Joseph, IDC)**
 - The current range of list price/peak performance (vectors excluded) is 8x for highly competitive, well selling products, \$6,300 to \$40,000 per peak GF. This shows that the price/peak performance metric has little meaning today.
 - Linpack correlates almost perfectly with peak performance, at peak = 1.48 times Linpack.
 - Goal: to both help users buy, and vendors explain, their products. Show advantages of all types of HPC computers, not just one category.
 - We created PAG (showed member list). It includes individual experts, as well as HPC users.
 - Need something that can be implemented readily and quickly—within one calendar year.
 - We are proposing the creation of the "IDC2000" metric:
 - Must have some correlation with price (market forces), at least better than 8x.
 - Immediately make an interim metric available, then a fuller metric within a year or so.

- **“Initial Investigations & Examples” (Earl Joseph)**
 - During the past 10 years, peak has lost all meaning in measuring and comparing performance.
 - Linpack correlates with peak for Top500 systems installed in both 2000 and 1999.
 - We also looked at another approach: peak times bandwidth. The correlation with price is better and it looks a little promising, but there are still some major outliers, so a better metric is required.
 - The point is that none of these simplifications work well as a metric. There's a need for a better metric.
 - (Question): Are we assuming market pricing is currently rational? (Earl): Yes on the average and with a wide degree of error. There are many exceptions, but a range of 8x is not due to market irregularities, perhaps a range of +/- 25% is due to market irregularities. We continually ask users about the value of their systems and compare prices, performance and a host of other metrics.
- **Ford (Vince Scarafino) Presentation**
 - We looked at the machines shown in the IDC report to see if we could get any correlation (with prices) on our own benchmarks. We used C90 units and looked at four types of machines:
 - PCs linked via the Internet
 - Workstations linked via high-performance connection
 - Commodity-based SMPs
 - High-performance supercomputers
 - We found about a 3x relationship between the categories:
 - Workstations 3x as expensive as PCs
 - Commodity-based SMPs 9x as expensive as PCs
 - High-performance supercomputers 27x as expensive as PCs
 - The point is that you get what you pay for. Some problems are unsolvable on PCs. It takes too long. Commodity-based SMPs are the heart of our capacity workload, and high-performance supercomputers are the heart of our capability workload. Bigger machines such as SMPs are easier to administer than PCs.
 - PCs can run CFD and parallelize. It's about one month to get the answer.
 - Workstations can run CFD and parallelize. About one week for the answer.
 - SPs can handle the low end of our safety computations.
 - Origin 2000s can handle most of our safety workload and some NASTRAN.
 - SV1s are between Origins and the high-end. They can handle more NASTRAN.
 - High-performance supercomputers like the T90 are used for NASTRAN and our toughest safety work.
 - In the past, capacity machines paid for capability. No longer.
 - (NSA): Today's procurement practices are irrational. They're geared toward cheap hardware and don't include manpower costs.
 - We want to build a vehicle that can go 100,000-200,000 miles without any service needs. We can't do this without high-end systems much more powerful than anything available in the global market today. Micro-based systems can do a lot of capacity work, but can't tackle the toughest problems.
- **DOD MOD (Larry Davis): “Benchmarking For HPC Systems Acquisition”**
 - We have a benchmark suite we've been using. In the next round of acquisitions, we have a directive to pay more attention to what users want. Users tend to want more of what they already have.
 - We don't foresee any one system handling all of our problems well. We use all “HPC systems classes”: distributed memory, shared memory, parallel vector

- (which really means high-bandwidth). We have just started looking at a fourth category: PC clusters.
- We are stressing the use of benchmark data vs. theoretical peak performance:
 - Important to use representative benchmarks, not peak
 - Actual implementation of benchmark evaluation is difficult
 - Prediction of performance on new systems often substituted for actual benchmark data
 - Price/performance still will be a criterion (\$/peak). Management looks at this and at our standings on the Top500 and other lists.
 - These lists have given us a lot of trouble because they give management bragging rights.
 - We've formed a benchmarking team with representatives from all four major centers (MSRCs). The team will manage benchmark development for the next round of computational capabilities requirements.
 - We will follow closely the User Forum activity in new performance metrics and benchmarking.
 - Our benchmark suite is likely to be broad and extensive. It will include complete application codes, kernels, and synthetic benchmarks.
 - Maybe this industry needs two classes of metrics:
 - **extensive** (related to the size of systems) and
 - **intensive** (shows the characteristics of systems).
 - (Cray): The last DOD MOD round was two years ago. The reason for peak procurements today is related to the time and expense of full-blown benchmarking for procurements.
 - (Ford): We continually update our benchmarks.
 - (DOD MOD): We run both in-house and third-party codes.
 - (Boeing): We have waited a long time for parallel code development and have been disappointed.
 - (DOD MOD): Many of our applications now run well on big parallel systems of up to 500-1000 processors. Another fraction needs very fast, very high-bandwidth systems.
 - (Comment): Price alone is not enough -- we need to include the cost of use, including the cost of making applications work on parallel machines.
 - **Candace Culhane, NSA: "A New Performance Metric"**
 - We have 12 unclassified HPC benchmarks that we give to vendors.
 - Sustained memory bandwidth is important to us
 - Our metric is Giga Updates Per Second (GUPS). This uses a standard distributed memory programming model.
 - 3 different classes of tests (vendor can choose), plus an additional test that characterizes memory performance.
 - Vendors had 4 weeks to prepare responses, then an additional 3 weeks to refine responses.
 - The high-end vendors are capable of building systems that can run these benchmarks. All vendors responded with current products.
 - The results are very interesting in showing major drop-offs when cache is exceeded, and then when running outside of main memory. A very telling picture/profile is created.
 - Vendors haven't delivered much real progress in 5 years.
 - **Bill Kramer, NERSC: "The NERSC EWP Test (Effective System Performance)"**
 - Concept: simulate a day in the life of an MPP
 - Test: to evaluate real systems in the real world.
 - Goals: evaluate existing systems and systems in development -- sustained system performance is better than peak or Linpack.
 - Performance: how much science can be done for a given quantum of CPU time.
 - Test includes full configuration run (at least ½ of the CPUs).

- We have results to date for 512-PE IBM, T3E, Compaq, and soon SGI. Our contract with IBM includes measured performance improvements over time. The IBM runs the tests in less time, but the T3E shows better utilization (current test doesn't include startup times, etc.). 90-95% utilization on T3E, 80-85% on SP.
- (Boeing): We tried this approach. It's hard to keep tests current with actual codes being run.
- (NERSC): The test is new, so not an issue yet.
- **Jose Munoz, DOE/ASCI**
 - ASCI has benchmarks that have been publicly available but won't be used in the future.
 - We are in a quandary: if the benchmarks are too much like our applications, that's a security issue; if not close enough, they're not useful.
 - I agree with others who said a one-size-fits-all benchmark won't work.
- **David Filkin, Dupont**
 - In recent procurements, we've been predicting the requirements of our key applications, and systems requirements to support this.
 - When we ask management for money, we don't talk about hardware or architectures only, we talk about the value of the applications and the cost of providing them.
 - In our own evaluations, within our own group, however, benchmarks are important.
 - The other measure we use is "delivered CPU cycles per month" (delivered to users).
- **Vendor Updates:**
 - **Mike Henesey, SRC:**
 - We now have powered on 32 processors (Intel Xeon). The product goal is to scale in a flat memory mode to 512. Targeting availability year-end 2000 timeframe. Reconfigurable logic, for interconnect and for programmable logic.
 - We are looking for users to give us codes to profile.
 - Privately financed today. We'll probably go after private placement sometime in the next 6 months.
 - Have been partnering with Oak Ridge, but others have expressed interest too.
 - (Question): What programming model? (Answer): We have Solaris running and are doing MPI work. The aim is to soon run cache coherent to 32 CPUs, going to 64.
 - (Question): Which workloads will do well? (Answer): We should do well on high-bandwidth, low-latency work. Also on work exploiting our reconfigurable logic.
 - **Ron Matlock, Fujitsu America (Houston, TX):**
 - VPP5000 (vector systems) has 9.6 GF CPUs, and scaling up to 512 CPUs. Customers include ECMWF, Toyota, Honda, Meteo France. They're running NASTRAN, CFD.
 - PrimeTower2000 (scalar) has 1.35 GF CPUs. 128-way SMP (64-way UMA). SPARC 64 CPUs.
 - HPCClusters@Fujitsu. 176 nodes. Fujitsu SL100 cluster. 156 GF, 88 GB memory, 12 TB disk. Solaris.
 - **Richard Kaufmann, Compaq:**
 - Lawrence Livermore: 507 GF peak. Supports MPI, SHMEM, and UPC. 6 microsecond from any PE to any other.
 - AlphaServer GS320. Up to 32 CPUs.
 - DSICL Beowulf. Ethernet or Mirronet (Livermore is Quadrics). EV67 CPU. 40 CPUs/node (chassis). Will run Linux.

- Next: EV7 (Alpha 21364). Glueless SMP, >1 GHz, ~13 GB/second/CPU, and 3 GB/second IO. We increased memory bandwidth and decreased latency.
 - EV7 available end of 2001, various systems a few months to 1 year afterward.
 - **Pete Ungaro, IBM:**
 - ASCI White is big milestone for us. 512 nodes, 16-CPU nodes. 12.88 TF peak, 160 TB disk managed by single parallel file system. Our first big copper machine. By end 2000, we'll have 8 systems with 1 TF peak or more.
 - Power4 update: >1 GHz. Booted both AIX and Linux. 4 contracts at 1 TF peak or more (3 Europe, 1 U.S.).
 - Doing lot of work in industrial markets: auto, aero, chem/pharm.
 - We think that a new performance metric is a good idea. Encourage PAG to be user-based, no vendors.
 - **Bill Minto, SGI.**
 - Warren Pratt is new head of engineering. Two R&D efforts: Mountain View and Chippewa Falls/Eagan.
 - About to announce Origin 2000 successor (Origin 3000, aka SN1). Pushing modularity to exploit availability of new technologies. Scales to 512 CPUs. Run in SSI mode to 512 CPUs, or users can run in partitioned mode.
 - \$100 million in advance orders, including systems to 512 CPUs.
 - MIPS/IRIX now. Intel/Linux coming, based on the same "fabric" of modular components.
- **Erich Strohmaier: Background on Top500**
 - Original intent of Top500 was not to spark competition among vendors.
 - Historically, supercomputer performance has shown exponential growth on Moore's Law curve. That can't be maintained indefinitely, but we'll see it for at least the next decade.
 - 1993 introduced the first Top500, an updated list of the most powerful computers in the world, defined as "supercomputers."
 - Now the emphasis has shifted from defining "supercomputers" to defining "powerful."
 - Linpack was the only practical choice in 1993. Main purpose was to show that a system could stay up long enough to complete the test. We know at least 500 vectors were installed worldwide in '93, ergo Top500.
- **Erich Strohmaier: "Benchmark For High-End Scientific and Technical Computing"**
 - To develop a benchmark, you first need to decide the objective. Do you want a procurement benchmark, an architecture benchmark, an applications benchmark, a throughput benchmark?
 - I'd like to focus on a benchmark that has scalability, meaning it is useful over a longer period of time. That means it must be (a) synthetic and (b) simple.
 - Hard to predict how architectures will change, or what new architectures will emerge -- ergo, don't base the benchmark on architecture.
 - Applications also change, creating a need to change the benchmark, so don't base the benchmark on applications either. For example, the NAS Parallel Benchmarks (NPB) were based on vectorizing codes on the fly.
 - You need to watch out for people who try to devise a method to circumvent the benchmark (bend or break the rules).
 - Also, problem size must scale or else, given Moore's Law, the benchmark will be done in cache at some point.
 - A simple model of a "system" that includes four elements: CPU, memory, disk, and external network. This model is simple enough to last for a long time.
 - Linpack is a pure CPU benchmark today.
 - Positive: scalable over wide range, although approaching its limit.

- Negatives: not sensitive to memory speed, although it depends on memory size; over-optimized compared to real applications.
 - NPB started as a set of 8 different benchmarks.
 - Positive: Data structures more like real life, although most aren't hard enough.
 - Negatives: Large overlap in what they show—can be boiled down to 2 numbers without losing discriminative power (same story with Spec95); only partially scalable—no automatic scaling and different problem sizes are hard to compare.
 - Scalability requirements for a new benchmark:
 - Simple enough to be usable and maintained
 - Complex enough to reflect the influence of all system attributes of interest
 - Scalable to utilize resources in variety of system sizes
 - How to design benchmarks:
 - Not good to test each attribute (CPU, memory, etc.) in isolation. This is simple, but does not reflect real life and is open to over-optimization.
 - Simulated applications are harder to design and understand, but provide more realistic performance.
 - Don't expect a benchmark that completely eliminates the need to benchmark some of your own codes. The standard benchmark is so you can get beyond handing vendors 30 codes to benchmark.
 - How many benchmarks are needed in the suite?
 - 2 to 3 well chosen ones will do the job.
 - Scalability Requirements:
 - To be useful for 10 years, the benchmark must span a range of 100,000 in system size (to cover Moore's Law).
 - It must use the whole system, independent of system size.
 - Runtime must remain reasonable.
 - Key features:
 - Small number of benchmarks
 - Synthetic
 - But with data structures similar to real applications
 - Pencil and paper benchmark descriptions
 - Reference implementations
 - Scalable problem sizes
 - Performance attributes independent of size
 - Reasonable runtimes
 - Discussion of Strohmaier Presentation:
 - (DOD MOD): Sounds like we're creating a new list that's less evil than the Top500 list.
 - (Earl): Yes, but alongside this standard benchmark we'll also create the matrix suite of real-world benchmarks as proposed in the IDC bulletin.
 - (Erich): This could be an add-on to the Top500, with more attributes.
- **Chris Willard, IDC: "Clusters, SuperClusters and Internet Computers: The Ultimate Computers?" (Proposal for an IDC study)**
 - We want to look at the impact of these things on current computing strategies:
 - End-user clustering strategies and costs
 - Vendor clustering strategies and costs
 - Internet computing strategies and costs
 - (NERSC): Security is a major obstacle to Internet computing. Also, how many home PC users will let people put 'garbage' on their PCs?
 - [SETI@home](#) was not-for-profit, but now there are also Parabon and two other for-profits doing Internet computing.
 - (Several comments): Include the ASP [Applications Service Provider] model in the study. There are many ASP server farms today claiming tight security.

- For this study, “cluster” means mainly roll-your-own systems assembled at customer sites, not factory-made systems. “Supercluster” means very big examples of these (1000+ CPUs).
- (Question): What’s the threshold? (Answer): A \$20K-\$25K system qualifies as a cluster.
- (Comment): As part of study, look at how this all fits into HPC. (Answer): We plan to do that.
- (Comment): I’m concerned management will say, “Internet computing is doing great things. Why do you need to spend \$10 million for a supercomputer?”
- (NSA): We and the auto industry have strong time-to-solution requirements that [SETI@home](#) doesn’t. SETI hasn’t found an answer yet - we need an answer in X amount of time.
- **Stan Posey, SGI (vendor presentation on user success stories: “Breakthroughs in Grand-Scale HPC for Research & Industry”**
 - Presented industry observations and a number of user success stories.

TUESDAY STEERING COMMITTEE MEETING: 7-24-00

- (Question): How will the PAG work? (Earl): Primarily by email, but also meeting in September (Richmond, VA) and at SC2000.
- **Earl will email Erich’s slides to PAG members.**
- (Question): How will a final benchmark be ‘adopted’—by IDC, by PAG, by the User Forum?
- **Earl will draft an ‘adoption procedure.’**
- (Earl): Re the role of IDC vs. the Steering Committee—in general, not just vis-à-vis the performance issue—the Steering Committee will function like a board of directors to set direction, and IDC will function like operational staff to implement the direction.
- (Boeing): Linpack has a singular focus on peak megaflops. Through the IDC benchmark, we can communicate our requirements to vendors. We could generate a list of, say, 20 variables (e.g., IO, memory bandwidth, CPU speed) that lets user organizations assess the strengths and weaknesses of each vendor system. Also of interest to us is the amount each vendor is investing in R&D for the future, e.g. for Cray, Inc., we’d want to know how much the company is spending on the ‘Cray’ side versus the Tera side.
- (Debra, IDC): HPC has an important impact on society but is a small dollar market. It’s a rounding error for the largest vendors. Benchmarks are important to reward HPC focused R&D. If this isn’t rewarded, it won’t happen and peak will increasingly dominate as a design goal.
- (Boeing): I object to single-focus metrics. (Earl): We could use a single number, but it would be based on multiple different tests, e.g. the peak times bandwidth numbers that I presented yesterday. We need two things:
 1. A standard benchmark to provide air cover with management — at least a better list than the Top500.
 2. A more realistic matrix suite of different actual benchmarks to assist with procurements.
- (Group decision): No more vendor presentations of customer success stories.
- (Group decision): Earl will review/delete Top100 stories submitted by vendors to the user forum web site and will ask vendors to have users re-submit them.
- A list of changes and suggestions were discussed that will be incorporated in future meetings, e.g. more room space, follow the agenda schedule, Earl to moderate the presentations, etc.

END