Report of the NAS Reinventing Team

Ronald Deiss, Richard Gunderson, Jeff Hultquist, Thomas Lasinski,
James McCabe, Leigh Ann Tanner, Harry Waddell, Howard Walter,
Sisira Weeratunga and Thomas Woodrow

Report NAS-94-003, April 1994

NAS Systems Division
NASA Ames Research Center
Mail Stop 258-6
Moffett Field, CA 94035-1000

Introduction

In the last decade, advanced computing and communications technologies have become increasingly important factors in ensuring America's continued leadership in the aeronautics industry. At the forefront of this trend, NASA created the Numerical Aerodynamic Simulation (NAS) Program in 1984 to focus resources on solving critical problems in aeroscience, space technology, and related fields using the power of the most advanced supercomputers available. Since then, NAS has pioneered many supercomputing technologies and techniques that have become industry standards, and at the same time provided a powerful research tool for the national aeroscience community.

In the last five years, however, rapid changes in technology, government, and the world marketplace have profoundly affected both aeronautics and supercomputing. Since the end of the Cold War, decreasing budgets and increasing competition from foreign aerospace companies have put pressure on the U.S. aerospace industry, causing a ripple effect throughout the aeronautics community. The emphasis within NASA has shifted from open-ended research-oriented programs to projects with more clearly defined performance and technology transfer objectives. During the same time frame, other supercomputer centers have approached or equaled NAS' high-speed computing capability, and advances in workstation technology have placed significant computational power on the desktop. NAS no longer has the leadership position in supercomputing or pathfinding that it once did. In this new environment, it does not make sense to continue operating on the same set of assumptions on which NAS was formed a decade ago. In short, NAS must “reinvent” itself in order to continue to fulfill its mission of providing NASA with a unique, leading-edge computational resource.

This report describes the results of the first phase of a reinventing process that was initiated in October 1993 by the NAS Division Chief. A Process Action Team of ten staff members, representing all areas of NAS, was directed to begin with a “clean sheet of paper,” and design an effective and innovative computational aeroscience resource for the next decade. During this reinventing effort, a series of fundamental insights and observations about NAS' role in aeronautics and high-speed computing emerged. Based on these realizations, the team made a number of important recommendations that call for dramatic change in the way NAS does business. Taken collectively, these recommendations form the blueprint for a revitalized NAS, designed to meet the challenges of high-speed computing into the next century.

New Realities

Historic changes in society and technology are irrevocably altering the market realities for both aeronautics and supercomputing. These changes have created an environment far different than the one in which NAS was founded ten years ago. As the NAS Process Action Team (referred to as the Reinventing team) began its work, two key questions quickly emerged:

- What are the current and projected “market conditions”, within which NAS must succeed?
- How can NAS best contribute to the NASA aeronautics effort, given these conditions?

The Reinventing team’s responses to these two questions created the framework for the reinventing process and the foundation for a new operational paradigm.

Foreign competition in aeronautics, reduced defense spending, and dwindling research budgets motivated NASA to become more efficient and customer-oriented as a research organization. Aeronautics research efforts are now focusing sharply on programs directed toward enhancing the global competitiveness of the U.S. aeronautics industry, rather than on programs simply aimed at basic research. Consequently, NASA resource management has shifted from a policy of entitlement, which allowed extensive freedom to pursue research of interest, to one of project accountability. Now, for many programs, research plans are constructed at the field centers expressly to meet customer-oriented objectives established by NASA management. Although NAS is recognized as one of the more successful programs within the NASA Aeronautics organization, with this new environment, the original NAS goals of providing “a national computational capability” and “a strong research tool for the Office of Aeronautics” no longer seem sufficiently clear and compelling. *The Reinventing team agreed that NAS should be re-oriented toward solving “real” problems that will have the greatest positive impact on the competitive position of our customer, the U.S. aeronautics industry.*

While political and economic forces have changed the way the aeronautics industry does business, the maturation of the supercomputing industry has significantly eroded NAS’s leadership position in high-speed computing. When NAS was first conceived in the late 1970s, the idea of supercomputing in a production environment was still very new. Supercomputers were rare, as was the expertise required to use them effectively. When NAS went online in the mid-1980s, one of its major goals was to make the use of supercomputers in aeronautics routine by making them readily available to scientists. NAS accomplished this goal by pioneering many techniques that have become standards for integrating supercomputers into a production environment, such as: networking to other computers with common operating systems, using scientific workstations to visualize datasets, and developing transparent methods of handling data transfer and storage.

The high-speed computing landscape looks very different than it did ten years ago. Supercomputers are much more common and easier to use. In 1985, NAS installed the first production gigabyte-memory machine in the world; now there are at least 75 in service around the world. At the same time that supercomputer hardware and software have matured, other supercomputer centers have emulated the NAS paradigm. Where NAS was once a unique computational resource, most U.S. supercomputing centers currently offer similar capabilities, delivered in a similar manner. In a way, NAS has become a victim of its own success. As a consequence, it is currently achieving very little differentiation in a crowded market where the barriers to entry are rapidly dropping. *Based on this realization, it was clear to the Reinventing team that NAS must move decisively to re-establish its uniqueness in the marketplace in order to avoid redundancy and obsolescence. Just as NAS originally recognized a significant need and sought to fill that need, it must do the same in the current market.*

When NAS was founded, supercomputers were several orders of magnitude more powerful than minicomputers. In fact, most computational aeroscience codes could only be run on supercomputers, as minicomputers lacked the necessary memory and processing power. The prohibitive cost of purchasing and maintaining these supercomputers, however, dictated that they be shared among many users. As a practical consequence, scientists learned to design codes that were tailored to run in a reasonable amount of time on their “share” of the supercomputer. In short, as supercomputing centers have evolved, supercomputer cycles have become commodities that are divided more or less evenly among large numbers of scientists, who then apply them to problems of similar scope.

During the same period of time that supercomputer use has become routine, the cost/benefit equations for supercomputers and workstations have changed dramatically. In recent years, the speed of workstations has increased much more rapidly than that of supercomputers, while the cost per unit of workstation computational speed has dropped significantly. For instance, the performance of the new IBM RS6000-590 RISC processor on the NAS Parallel Benchmarks is about one-sixth that of a single processor CRAY C90. At the same time, a reasonably configured IBM RS6000-590 workstation is about one-twentieth the cost of a single processor C90.[1]

Because of the competition for time on the NAS C90, most users get overnight turnaround on a one-hour job. The same job would take about six hours on the IBM workstation and would also be run overnight. This means that any researcher who can find $120 thousand for a 512-MB IBM RS6000-590, and whose jobs fit in that size memory, will have no need for NAS. In fact, a recent study by the NAS High Speed Processor group found that nearly all of the jobs currently running on the C90 could be run on this type of workstation. The present method of project selection, which is outside the control of NAS, divides the resource among so many users (about 1500) that it effectively neutralizes the large computational capability of the C90. *The Reinventing team concluded that continued support of this customer set will only further erode NAS' viability as a center for solving aeronautics problems of national importance, eventually leading to downsizing, consolidation, or closure.*

---

1. D. Bailey, E. Barszcz, L. Dagum and H. Simon, “NAS Parallel Benchmark Results,” NAS RNR Technical Report RNR-94-006.

In order for NAS to successfully reinvent itself, it must formulate a new operational paradigm that is customer-focused, unique to the aeronautics community, and makes the most effective use of its computational resources for solving aeroscience problems of national importance.

A New Paradigm

As the Reinventing team undertook the challenge of creating an new operational paradigm for NAS, more difficult questions arose.

- Do our customers still need us?
- What kind of return on investment should be expected from a large, expensive national resource like NAS?
- Why does NAS need a 16-processor C90 with gigabytes of memory?
- If NAS does need the C90, how can its computational power be used most effectively?

In the course of answering these questions, the team considered several proposals for new operational paradigms. Most of these proposals which were quickly discarded after failing to satisfy the established criteria of customer-focus, uniqueness, and effective use of resources to solve important aeronautics problems.

For example, one proposal which suggested that NAS should become a center for distributed computing was dismissed early on. Not only did this proposal fail to focus on customer needs and effective use of existing resources, but several other organizations are already making significant contributions to distributed or network computing. It was also suggested that NAS become the sole NASA center for supercomputing. This approach is attractive as a cost-saving measure, yet it merely emphasizes the shortcomings of the traditional supercomputer center model. Competition for resources would only increase, further diluting the effectiveness of the C90. Others have suggested that NAS concentrate on research and development in support of scientific computing, becoming a center for software and hardware evaluation and integration. A focus of this kind would lack uniqueness and also fail to apply computing resources directly to the solution of aeronautics problems.

In the end, only one proposal definitively answered the difficult questions and satisfied the necessary criteria. Although it is the most radical of the concepts proposed, it also holds the greatest potential benefits for NASA and the U.S. aerospace industry. *The Reinventing team proposes that NAS become* the *center for solving applied aeroscience and engineering problems by working directly with customers to focus the entire NAS resource on selected projects that will extend the limits of scientific computing and engineering.*

Under this new paradigm, NAS will provide its customers with a unique and powerful end-to-end computational system (including high-speed processors, networks, mass storage, and visualization) aimed at solving “critical path problems,” that is, problems where the system is expected to yield at least an order-of-magnitude improvement in performance, thereby reducing or eliminating an obstacle in the aircraft design or manufacturing process. In order to effectively solve these critical path problems, the entire NAS resource would be allocated in one-month increments to individual customer projects. NAS would then support about a dozen projects per year instead of the current project load of approximately 500 projects. Consequently, greater attention could be given to the precise requirements of each customer. For instance, a NAS team would be assigned to each project to help expedite the creation of special hardware configurations and to satisfy requests for unique software. Team members, serving for the duration of a project, would be drawn from the various technology groups within NAS to provide the appropriate skill mix.

Two major types of “critical path” projects have been identified by the Reinventing team:

- Radical solutions to scientific or engineering problems through high-fidelity computational models, which use more physics and greater resolution than are currently possible.
- Radical use of computer resources to demonstrate the feasibility of solving time-critical engineering problems, such as shortening segments of the design cycle by an order of magnitude or more.

Aeronautics firms are, of course, very interested in reducing the time involved in the aircraft design cycle, since it directly affects both the time-to-market and the end cost of products. While NAS is not meant to be a design center for aeronautics, it can be used to demonstrate how a focused computing capability will significantly compress various aspects of the design process. At the same time, it can help develop computational tools that can later be used by the aerospace companies on their own systems. A typical project might involve using the NAS capability to find a way to reduce the design time for a specific vehicle component from a year to a month or even less. The success of such projects would improve the United States competitive position in the world market, while encouraging greater investment in computing by aeronautics firms.

Today, supercomputer users only think in terms of making incremental improvements to their models because of limitations placed on the computer time, memory, and disk space available to them. Under the traditional supercomputer center paradigm, the user's share of the resources usually remains approximately the same over time, so the computing capacity available to the individual only increases meaningfully if the whole system is upgraded.

In contrast, the Reinventing team is proposing to immediately make available to leading scientists a level of computational capability that would not be available to them for many years in a shared system. By concentrating the enormous power of the NAS resource in the hands of a few scientists working on a small number of problems, truly significant advances in the solution of critical scientific and engineering problems will be possible.

Much of the technical description and details are discussed in the Appendices. Appendix A describes a prototypical large problem and its associated CPU, I/O and graphics requirements, Appendix B discusses some of the risks that such a focused research effort entails, and Appendix C lists some sample guidelines for selecting projects.

Reengineering NAS

With the new focus on solving critical path problems, the Reinventing team began the second phase of its mission. The “blank piece of paper” they began with now had a new goal written on it, but no clear path toward accomplishing that goal was yet defined. The team realized that a radical change in goals would require an equally radical change in work processes in order to bring the organization into alignment with those new goals. Consequently, the team initiated a comprehensive reengineering process, taking into account the internal realities described by the new paradigm as well as budget and staffing restrictions.

Due to the success of a popular book on the subject, “reengineering” has become the management current “buzzword.” Yet beneath the hype, reengineering provides powerful techniques, which the Reinventing team applied in designing more effective work structures and processes for NAS. Reengineering is defined as “the fundamental rethinking and radical redesign of business processes to achieve dramatic improvements in critical, contemporary measures of performance, such as cost, quality, service and speed.” In this case, the Reinventing team was faced with the task of designing a new organization that would provide the expertise and technology to directly support customers in solving critical path problems, and at the same time working within the constraints of flat to decreasing budget and staffing levels over the next few years. In other words, like many business organizations in the 1990s, NAS must learn how to do more with less. The Reinventing team defined two critical changes that NAS must make.

First, NAS must set more rigorous measures for selecting research and development activities. Under the umbrella of a balanced system concept, NAS continues to invest in many different technology areas. The great breadth of investment that NAS has made has severely limited the amount of resources that have been invested in any given area. *This lack of focus has created a NAS that no longer has a clear leadership position in any area*. With this realization in mind, the team set out to develop a process that would generate an R&D portfolio for NAS that is better suited to the newly proposed paradigm, and that produces maximum returns on NAS’s constrained R&D budget.

Second, NAS must design a new organization that applies its resources toward the solution of critical path problems with maximal efficacy, minimal waste of resource, and negligible overhead. The existing organization is composed of Applied Research, Development, and Computational Services branches, none of which individually have the resources needed to create the project teams called for by the new paradigm. Clearly, personnel and resources would be required from all parts of NAS. One way to facilitate the creation of teams would be to transform NAS into a matrix organization where personnel could simultaneously work on a problem team while continuing to pursue their current assignments. This idea was rejected because it was feared that conflicting priorities on the part of the staff would detract from the efficacy of the problem teams. What is needed is an organizational structure that facilitates the creation of dedicated teams and allows the basic R&D needed to solve the next generation of critical path problems to continue.

New Organization

After the Reinventing team analyzed a number of reorganization options, one stood out as the logical way in which to proceed. *The Reinventing team recommends that NAS reorganize into two branches, with one dedicated to solving customer problems and the other charged with developing the technology needed to solve the next generation of problems.* Members of each branch would be dedicated to their task, eliminating the coordination problems endemic to matrix organizations.

The Customer Problem Branch would be composed of a number of critical path problem teams, perhaps four or five at any given moment, with each team managing a different stage in the solution of a current problem. Once a problem team completes its job, members of that team would return to their respective positions in the Technology Branch. Each member's experience will help keep NAS researchers and developers focused on the needs of aeronautics, thereby enhancing the customer focus of NAS and providing invaluable insight into the direction that should be taken in preparing for the solution of future problems.

All of the research, development, and system support functions served by the current branches would be merged into one large Technology Branch. The work of this branch would be fulfilled by a set of teams with responsibility for a section of the NAS system as spelled out in a Statement of Work (SOW) created by the team. Furthermore, many of the Reinventing team members felt that the teams should have an "end-to-end" responsibility for the systems that they develop and manage. Compared to the current system where research, development, and implementation are handled by different groups, each being actively managed by a civil servant, this new approach has several advantages.

First, the "end-to-end" responsibility for a product or project will encourage developers to be more aware of the customer needs and service implications that a given project may pose. If, for instance, a team is charged with developing and then operating a system, they are much more likely to make sure that all the tools required to efficiently service and maintain that system are in place, thereby leading to increased operational efficiencies.

Second, each team will have relative freedom to operate within the guidelines of the SOW, allowing for greater self-direction and, consequently, resulting in a decrease in required supervisorial overhead as necessitated by the current realities.

Third, consolidating the functions of the current branches will result in increased efficiency in implementing new systems by eliminating much of the administrative delay experienced when projects are transitioned from one "sphere of influence" to another.

Much of the detail of how the teams would be organized, and the pros and cons of various options appears in Appendix D. For instance, some members of the Reinventing team were troubled by some of the implications of teams having an end-to-end responsibility. *Clearly, the actual breakdown of functions within the Technology Branch must be left up to the NAS branch chiefs and the division chief.*

New Processes

For NAS to succeed in its new mission, it is of paramount importance that NAS makes the correct investment decision. Therefore, the team attempted to develop a process by which these budgeting decisions could be made more effectively. One of the keys to unlocking this resource allocation puzzle lay in the SOW that each team would prepare. By integrating the high-level strategic planning carried on by upper management with the knowledge and experience of the staff, especially those with recent experience on critical path problem teams, it becomes possible to make optimal decisions.

The Work Selection Process

The proposed Work Selection Process has five parts: Vision and Strategy, Writing the NAS Statement of Work, Reviewing the Statement of Work, Normalization and Ordering, and Resource Allocation.

NAS needs a vision of what it wants to achieve. The current vision is broadly stated in our mission SOW and is only updated every few years. The Reinventing team is proposing a fundamental change. *To correctly manage the NAS resources, it is essential that upper management meet frequently (perhaps on a weekly basis) to refine the principles, goals, and objectives of NAS.* This vision and strategy must be continually updated and annunciated by senior management. Clearly, management will need to draw heavily on the technical people at NAS if they are to have any hope of developing a vision and strategy that will keep pace with the rate of technical advancement that the foreseeable future holds. Some of the types of decisions and commitments that would need to be made are:

- Define, maintain, and provide timely updates to the NAS vision.
- Establish and monitor the pathfinding and production balance.
- Determine rough resource allocations between the different subsystems.
- Determine the number of NAS Solution Teams.
- Give guidelines on technical balance; for example, cluster computing versus vector or parallel architectures or openness of information versus security concerns.

Every six months, teams or individuals would use the NAS vision to guide them in creating or updating a SOW which describes the contribution and role of each team member and the technical objectives of the team for the next six months. SOWs may also be prepared for which there is no proposed staffing as a way of identifying needs that have not been met, or as a mechanism for introducing work that was stimulated by an outside source such as the User Interface Group (UIG). All SOWs and their subsequent scores must be made available to the NAS staff, preferably via some mechanism such as a World Wide Web server.

Each SOW will then be scored by at least three people using a set of evaluation criteria, which are enumerated in Appendix E. It is suggested that one reviewer be a customer and one other be outside the proposer's team. Any NAS staff member may choose to prepare an unsolicited review of any proposal, and this input will be considered in the overall score of the proposal if the review is publicly available and includes the reviewer's name. Depending on the nature of the project (for example, service, development, or research) the criteria may be interpreted differently. For instance, when evaluating the importance of a project to customers, a service project could be judged on the usefulness of its deliverables, while a fundamental R&D project could be judged on the potential impact it may have on some key technology area. It is proposed that the following serve as a set of evaluation criteria:

- Importance to customers or potential impact on aerospace
- Technical merit and Innovation
- Required resources
- Individual/team qualifications
- Schedule and deliverables
- Technology transfer or implementation plan.

Reviewers must score the SOWs on each criteria and explain how the score was determined. The more costly the project, the more carefully its proposal will be reviewed. Proposals for very large projects may be reviewed outside NAS, while it may be more appropriate to implement a streamlined review process for small projects. If submitter(s) feels that the proposal has been judged unfairly, there should be one week in which they may call for a meeting with reviewers to defend the proposal.

After the SOWs are reviewed, a substantial amount of "post-processing" needs to be done before any decisions can be made. The scores need to be normalized by comparing scores of various proposals to each other in an attempt to eliminate different scoring "tendencies" that reviewers may have. In the event of a large difference of opinion, the reviewers should be approached and asked to clarify their position. After normalization has occurred, *a rough ordering of the proposals by score can be made based on technical merit.* Certainly, there are some unresolved issues with respect to ordering the SOWs. In Appendix E, it is suggested that the criteria be weighted based on the research, development, or support orientation of the proposed project in an effort to find a balance between high-risk research and low-risk support and implementation projects. The people charged with doing the normalization and ordering may choose to adopt a different weighting structure or decide on some other measure with which to order proposals.

Final determination of which projects to fund and what personnel allocations to make are arrived at in the Resource Allocation step. Resource Allocations could be made by plotting the proposals by cost and score, by examining histograms relating cost, score, and staffing requirements, or by some other means. The "clear" winners should be accepted: the borderline cases should require further examination before being finally accepted or rejected. Then, on a case-by-case basis, attempt to unite unassigned people, or even reassign people, to deserving projects. If any remaining proposals have significant merit but cannot be funded or staffed, they should be added to a "To Do" list available to NAS technical staff, management, customers, and stakeholders in hopes that it will provide impetus for new funding or directions for new research when other work has been completed.

In the course of evaluating the effectiveness of the Work Selection Process the Reinventing team realized that if NAS is to have any flexibility in responding to new work, it must try to limit the level of long-term budgetary commitment that it makes. Any new proposals to make a long-term commitment must be made while bearing in mind the consequent restraints implied.

Conclusion

Current realities dictate that NAS must transform itself to regain its former prominence. Constrained manpower and budgets, a changing NASA culture, and the homogeneous landscape of supercomputing are all compelling factors that force the Reinventing team to conclude that NAS must change. A new paradigm for doing business has been proposed that will provide the maximum benefit to our aerospace customers. To address the new realities, a new organization better suited to solving critical path problems has been designed and better, more efficient work processes have been proposed. It is the team's opinion that if all these changes are adopted, NAS will better serve the needs of aerospace into the next century.

Appendix A: Technical Requirements for Solving Large Problems

The following describes the requirements for solving a relatively small problem when compared with the critical path problems of even the very near future. This problem was proposed by an aerospace company. The objective is to show that simulations could be used to study a vibration problem in a rotorcraft design. Under the existing paradigm, this is very difficult to do from a technical and administrative point of view and would take months to accomplish.

Example Problem Description

An example problem was described by Barszcz, Weeratunga, and Meakin in a proposal written in August 1993. The proposal is to perform a large parallel CFD problem on the C90 to stress the complete NPSN, run an application to demonstrate the capability to the NAS user community and produce a large benchmark for MPP vendors to match. The problem is a time-accurate simulation of the Tilt Rotor in forward flight.

CPU Time

Revs: 300
Steps/Rev: 1500
Y-MP time/step: 60s
Ratio C90/1 - Y-MP/1: 2
Parallel Efficiency: 70 percent
C90 Processors: 16
Total Time Required: 355 hours or ~14 days

Data Storage

Steps/Dump: 15
Total Files: 30,000
Var/Grid Pt.: 5
Bytes/Var: 4
Grid Size: 1.3 M points
File Size: 26 MB
Rate: 26 MB/40s or .65 MB/s
Total Storage Required: 780 GB

Bottlenecks in Current System and High-level Requirements

There are several bottlenecks or problems in the current system:

1) insufficient disk to hold the dataset (on the HSP or TAVS)
2) moving data between HSP and the TAVS or MSS
3) insufficient CPU power to run complete analysis passes daily
4) insufficient memory to perform interactive analysis on a meaningful percent of the solution

Given the bottlenecks, there are several requirements we suggest are necessary for a future system.

1) The ability to run daily passes through the complete dataset.
   a) minimal requirement: one pass a day
   b) realistic requirement: two to four passes a day
   c) optimal requirement: one pass every 30 minutes (based on a 10-frame/sec rate)

2) The ability to save output from the daily passes and replay it at another time.

3) The ability to back up or restore an entire dataset to removable media (tape, optical) in several hours.

Detailed Analysis of Requirements

Given the example problem described above and the high-level requirements, here is a more detailed analysis which drives some of the component performance requirements.

In the example problem, a total of 450,000 time steps will be generated over 14 days (1 of every 15th time step will be saved, for a total of 30,000 files). The analyst must therefore be able to keep up with 2,143 files a day. In order to drive the remaining requirements, we will assume that the analyst is using 100 streaklines to interpret the dataset.

In the following analysis we look at the requirement to keep up with the data created by the flow solver and the worst-case requirement to run through the entire dataset. These two cases are called out explicitly in each section.

CPU Requirements

One assumption made through the following scenario is that for any run through the data, only the latest N (e.g. 500) particles are saved from each streakline emitter. This is because the total number of particles is cumulative, one added per step, and looking at thousands of steps may obscure the image and add needless calculation.

Daily runs:

100 emitters * 500 particles/emitter * 1 particle/emitter/step <= 50,000 particles/step
2000 steps * ~50,000 particles/step = 10,000,000 total particles calculated

Assuming 1,000 operations per particle per step.

50,000 particles/step * 1000 operations/particle = 50 Mops/step
2,000 steps * 50 Mops/Step = 100 Gops for the daily run

Requirements for full runs:

30,000 steps * ~50,000 particles/step = 150,000,000 total particles calculated

50,000 particles/step * 1000 operations/particle = 50 Mops/step (million ops/step)
30,000 steps * 50 Mops/step = 1500 Gops (billion ops) for the full run

Calculation Operations	Time	Performance (MFlop/s)	Rate (steps/sec)
Daily Run	8 hours	3.5	.07
2,000 steps	2 hours	13.9	.28
100 Gops	.5 hour	56.0	1.10
Full Run	8 hours	52.1	1.04
30,000 steps	2 hours	208.3	4.20
1500 Gops	.5 hours	833.0	16.70

As a point of reference, UFAT currently takes about 10-14 seconds per step on the existing Convex C3240 to do a problem of about this size.

I/O Requirements

Input:

Given 1.4 M nodes/step, a new grid at each step and one vector and two scalar quantities of interest at each step, we have the following input data:

1.4 M nodes * (4 + 3 + 2[1]) * 4 bytes/node = **50 MB/step input**

Input totals:
daily run: 50 MB * 2,000 steps = 100 GB
full run: 50 MB *30,000 steps = 1.5 TB

Output:

10,000 particles * 8[2] * 4 bytes/particle = **.3 MB/step output**

1. The components of the term are (grid: x, y, z, iblank; vector, two scalars).
2. The eight values in a .GRA file are: rgb code, rgb components, position code, position components.

Output totals:
daily run: .3 MB * 2000 = .6 GB
full run: .3 MB * 30,000 = 18 GB

I/O Requirements:
Assuming that the input and output can be done independently, these are driven entirely by the input requirement.

Input	Time	Rate (MB/s)
Daily Run	8 hours	3.5
100 GB	2 hours	13.9
	.5 hours	5.0
Full Run	8 hours	52.1
1.5 TB	2 hours	208.3
	.5 hours	833.3

Graphics Requirements

Restricting the solution to 50,000 particles (or colored points) per step means that this is well within the performance of current graphics technology. This is the case even if we decide to later view at a rate up to 5 to 10 times the calculation rate.

The rates in the table below are driven by the calculation rate. The table does not reflect the geometry related to the airfoil, which is assumed to be less significant in most cases.

	Frame Rate (frames/s)	Graphics Rate (points/s)
Daily Run	.07	3500
	.28	14000
	1.10	55000
Full Run	1.04	52000
	4.20	210000
	16.70	835000

Appendix B: Risk Assessment

No other production supercomputer center has tried to operate under a model like the one proposed by the Reinventing team; this presents NAS with a real pathfinding opportunity if it acts promptly. Other centers are also aware of the changes in the marketplace, and they too must eventually address these changes. For now, NAS has the opportunity to pioneer this approach to supercomputing and in doing so, provide a significant competitive advantage to its customers. Unfortunately, a change in direction as radical as the one being proposed also carries a certain amount of risk. NAS will once again be pathfinding in uncharted territory, and will inevitably make mistakes as it searches for the most effective approaches to operating under this new paradigm. After considering the risks involved, the Reinventing team concluded that the potential rewards easily outweigh the risks, but that these risks cannot be ignored. Consequently, several strategies were developed for managing the risks and maximizing the chances for success.

The Reinventing team discussed the nature of critical path problems at length and made some initial estimates of system requirements for solving them. Based on this data, it was generally agreed that NAS could technically handle such problems. Some initial system designs and costs were identified. One of the by-products of this discussion is that NAS's choice of hardware and/or software technologies to significantly invest in should be based on which technologies best support the timely solution of customer problems.

Given the unheard of scale and scope of these efforts, many unanticipated difficulties in such calculations will arise. Failure will initially be high. Because of unexpected numerical instabilities, for example, a user might find his problem converging much too slowly, after having spent a significant amount of the allotted time.

Great efforts engender great expectations and great scrutiny. The stakeholders, NASA Headquarters and Congress, expect a good return on their investment. As we concentrate our resources on fewer customer problems, the stakeholders will expect more significant results. NAS must become even more accountable. As fewer projects (perhaps a dozen as opposed to the current 500) are run, NAS's "political" support base could be narrowed. Fairness in access to the resource will be critical. We must convince the aeronautics community that this is the wise way to use the resource and that it will be fairly allocated.

Our users must learn to think in terms of radical computational solutions to their problems. They must think on a very different scale in defining problems. We must commit to using the resource in this manner so that they can assess the possibilities and propose appropriate problems. That resource may initially be under-utilized as users come up to speed in proposing problems commensurate with that resource.

Appendix C: Selection of Customer Projects

Under the new paradigm, special attention should be given to the task of selecting customer projects. Given the small number of projects, allocation and scheduling of the NAS resource will be of critical importance.

NASA must develop a list of appropriate problems and selection criteria. This must be done in close cooperation with industry and academia representatives of the aeronautics community. As NAS is a national resource, the actual selection of projects will be done by NASA personnel. It must be perfectly clear, however, that the selection criteria and appropriateness of the problems have the full support of the broader aeronautics community.

Results of calculations done at NAS should benefit, at least in part, the entire aeronautics community. Proprietary calculations are not to be excluded; however, the details of such calculations should at least be available to NASA personnel so they can use the information to improve NAS capabilities.

NAS should play an active role in the selection process to ensure that computing resources are used appropriately.

Appendix D: Organizational Details

Several alternative organizational alternatives for the breakdown of the Technology Branch were considered; most were similar to the four alternative organizational structures presented below. Some of the organizations split the functions of NAS along “product time to market” or *temporal* lines in much the same way as the current organization splits the NAS functions up as Research, Development, and Service. In other organizations considered, the Technology Branch was split by the products and services produced at NAS, or along *functional* lines. A very simple way of describing the models is as follows:

- S.A. Model: Systems and Applications functional split.
- L.S. Model: Long-term vs. Short-term temporal split.
- Subsystem Model: Functional split breaking NAS into subsystem teams with end-to-end responsibility.
- R.D.S. Model: Places current NAS research, development, and services as large teams under the Technology Branch.

Each organization has some advantages and disadvantages, which are briefly summarized in the following table, using the current organization as a baseline for comparison.

Criteria	Current	S.A.	L.S.	Subsystem	R.D.S.
Solve Large Problems	Lacks Focal Point	Focal Point	Focal Point	Focal Point	Focal Point
Ability to make good decisions	Difficult, turf battles	Clear since minimal interfaces	Potential turf battles	Long Decision time, highest technical content	One more interface than L.S., hence more turf battles
Efficiency of Life/Death Process	Difficult, institutional, entitlement mindset	Work selection process same for all new orgs	Work selection process same for all new orgs	Work selection process same for all new orgs	Work selection process same for all new orgs
Balance Production and Pathfinding	Focused R, D, and S	Potential Problem	Potential Problem	Potential Problem	Focused R, D, and S
Prod/Project Accountability	Disconnect	Strong	Disconnect	Strong	Disconnect
Staff Development	Limited, easy for R and D, harder for S	Easier for S, harder for R and D	Less different than existing for S, harder for R and D	Easier for S, harder for R and D	Same as existing

In general, functional organizations provide a clearer focus toward producing products, but pose the danger that long-term research efforts may be preempted by daily operational needs. But, functional organizations allow the staff members assigned to support functions to be more readily involved in research and development, providing an opportunity for growth that might otherwise be unavailable.

Appendix E: Proposal Selection Guidelines

Because of the wide range of projects that are undertaken at NAS, the Reinventing team has proposed a common set of evaluation criteria, along with some suggestions or guidelines as to how they should be applied in the context of service, development, and research activities. These characterizations are not meant to imply that we need three branches, each with its own set of guidelines. We intend merely to tailor each proposal document to the questions most relevant to the task at hand.

Every proposal will contain a project description of one to five pages, describing customer services to be performed, systems to be purchased or built, or what research topic to investigate. This description is followed by an additional page or two, further describing the project with respect to these "evaluation factors"

- Importance to customers
- Technical merit and innovation
- Required resources
- Individual/team qualifications
- Schedule and deliverables
- Technology transfer or implementation plan

Reviewers will be encouraged to be verbose in their discussion of the merits of each proposal. The written comments will be important for properly evaluating a complex or controversial task. Reviewers will also be asked to assign numerical scores to each of the six proposal sections. The team suggests that the following weights be applied to these scores; note that the point distribution varies depending on the primary focus of the proposed task.

Criterion	Service	Development	Research
Importance	40	25	30
Technical Merit	10	25	25
Resources	15	10	10
Qualifications	5	5	15
Schedule	20	10	10
Tech Transfer	10	25	10

To a large degree, the proposals may be treated identically, regardless of their individual focus on service, development, or research. This "generic view" of the proposal is described below. Each question is augmented by some discussion of the differences in emphasis among the three categories. The text below is intended to guide the thoughts of the authors and reviewers. A point-by-point response to every question is not needed.

Importance to Customers

Explain the contribution of this proposed work and how it will support NAS customers, either immediately or in the future. If possible, identify by name the specific customers this is targeted for. Discuss the future benefits, both short- and long-term.

- Who are the customers, either internal or external?
- How will the activity serve the NAS goals?
- Is it called for by the NAS Vision and Strategy?
- What would be the benefit if this task were funded?
- What is the probability of user acceptance?

Service projects should also answer the following:

- Describe the benefit to be obtained from any new service.
- Identify any NAS customers currently using this service.
- Provide metrics of how this service has been used in the past.

Development projects should also answer the following:

- Identify the perceived need this project is intended to satisfy.

Research projects should also answer the following:

- Identify the categories of customers most likely to benefit from the knowledge generated from this investigation. Explain the nature and scope of its potential impact on their activities.
- Explain how and why the knowledge generated from this research project might support the more broadly defined NAS missions, as well as its more near-term goals.

Technical Merit and Innovation

Discuss the innovative elements of the effort. Since NAS has limited resources, it is important that we choose our battles wisely.

- Are there appropriate alternative sources for the proposed work?

Service projects should also answer the following:

- Are continuous improvement methods being applied to this effort?
- How is this going to improve the quality of service?
- Does this work offer a significant advance of our service?

Development projects should also answer the following:

- Is this a pathfinding activity or an incremental improvement?
- What is the technical probability that the project will succeed?
- Relate the time-to-complete vs. potential impact; that is, is there a limited “shelf-life” beyond which this technology will be no longer useful?

Research projects should also answer the following:

- What are the technical motivations for investigating this topic?
- What outside work has been done on this question?
- How will this work advance the state of the art?

Required Resources

Identify any special equipment, external grants, travel, or training expenses which would be required for this project. Estimate the cost of any procurements, and when these are expected to occur, excluding personnel requirements, which are covered in the next section.

- Could this project be successful if partially funded?
- Identify any other tasks that would be impacted if this proposal were not funded at the requested level.
- Identify any possible enhancements and the added value to potential customers if additional resources could be made available.

Service projects should also answer the following:

- Estimate the long-term costs of providing this service.

Development projects should also answer the following:

- Describe both development and sustaining costs.

Research projects should also answer the following:

- Why is this worth studying?
- What benefit might eventually accrue from this investment?

Individual/Team Qualifications

Identify all personnel associated with the project; describe the individual responsibilities of each team member. Identify the relevant capabilities of these people, especially any past experience in projects of a similar nature.

> **Note:** This criterion would not apply to so-called “unstaffed” project proposals. The Reinventing team suggests that for such proposals, the weights assigned to both criteria (1) and (2) be raised to 40 percent each.

- Who are the team members?
- What is the experience of the team members?
- Why are they the right people for the task?
- Have team members participated in similar efforts?
- Identify any applicable experience or training.

Development projects should also answer the following:

- Have you ever built anything similar to the proposed item?

Research projects should also answer the following:

- What prior related research have you conducted?
- Have you published any papers related to this topic?

Schedule and Deliverables

What is to be delivered and when? Also, what is the team's recent performance with regard to the overall quality of work, completion of deliverables within the allotted budget, timely cancellation of failing projects, meeting of scheduled milestones, and efforts to transfer technology and knowledge to relevant customers?

- Describe what you intend to deliver.
- Provide a schedule with concrete milestones or deliverables.
- Is this schedule realistic?
- Identify milestones that are critical to the success of the project.
- How accurate has the team been on past schedules?

Service projects should also answer the following:

- Provide an implementation plan that identifies risks and contingency plans to ensure success. Describe how the users will be made aware of any significant changes to their working environment.

Development projects should also answer the following:

- Are all the deliverables clearly defined and broken down to fit the time lines of any periodic design reviews?

Research projects should also answer the following:

- Can a sharp conclusion, positive or negative, be made from this research?
- Are the deliverables sufficient to achieve effective dissemination of knowledge to all potential customers?
- Will any prototype hardware or software be created?

Technology Transfer

Describe how you will deliver this product or service to end users.

- Will the work directly change the NAS environment, or will this resulting technology and knowledge be delivered indirectly?

Service projects should also answer the following:

- Identify any plans to export the results of this work beyond NAS; for example, via workshops, conferences, or papers.

Development projects should also answer the following:

- How will you get the result of this work into customer hands?
- Is this effort intended for one customer or for a group of customers?
- Will any software be sent to COSMIC?

Research projects should also answer the following:

- What plans, if any, exist regarding publication of this work?

Appendix F: The Bottom Line for Staff Members

This part of the Reinventing team report addresses some questions and concerns that staff members may have about the "new" NAS that the team is proposing.

First, the new Work Selection Process will *not* lead to anyone losing their job. It *will* probably lead to reassignment to different projects as the NAS tactical goals are developed, and in some cases, people may choose to leave if they are unable to do the work that they are interested in. The extra freedom to "write your own ticket" provided by the new Work Selection Process should more than offset this. Any staff member can suggest a new approach for doing something and then pursue it, provided that they can show promise and relevance to aerospace.

Second, all staff members will have the opportunity to work on a critical path problem team. In the course of such work, they will have the opportunity to work with representatives of the aerospace industry, thereby gaining experience and making professional connections beyond NAS. These new connections should help staff members better serve NAS customers and benefit their professional lives beyond NAS.

Other questions and side-effects of the proposed reinvented NAS certainly exist, but they are probably too numerous to list here. If you have any questions or concerns, please address them to a Reinventing team member, the NAS reinventing newsgroup, or the NAS System Division Chief.