- Vendors from all industry segments continue to integrate and combine parallel-processing technologies, exploding the number of choices for users and creating new challenges in benchmarking, operating systems and applications development.

- The commercial use of parallel processing for data-intensive applications will be a defining shift in the market.

- Market consolidation/shakeout will force early and well-known participants into niche markets to survive.

- Traditional vector supercomputing will be tied to legacy applications, but lower-cost designs will slow the migration to parallel systems due to application availability.

- Storage systems and data locality will emerge as both roadblocks and enablers of client/server computing.

- Government initiatives will no longer significantly impact the massively parallel processing (MPP) market's development.

Reader Notes

1. **How will the relentless improvement in underlying system technologies manifest itself in user choices for parallel and vector systems?**
2. **What is the evolving role of traditional vector supercomputers?**
3. **How will parallel systems impact commercial data-processing solutions in the 1990s?**
4. **Which high-performance computing (HPC) vendors will adapt to the recent market changes? How will they adapt? Which will prosper and which will not?**
5. In what applications will parallel-processing systems dominate alternative computing platforms?
6. **How will users deploy HPC systems within the existing IT infrastructure?**

The HPC market is undergoing broad changes, as users absorb the changes in technology that new CMOS and RISC microprocessors have delivered from a number of vendors. These advances are changing views of traditional vector supercomputing, the development of MPP, and the commercial use of parallel systems. This scenario will review the changes in technologies and the impact on users and vendors in the next several years.

How will the relentless improvement in underlying system technologies manifest itself in user choices for parallel and vector systems?

Reader Notes

Number of Processors Required to Reach 1 GIPS (Billion Instructions per Second)

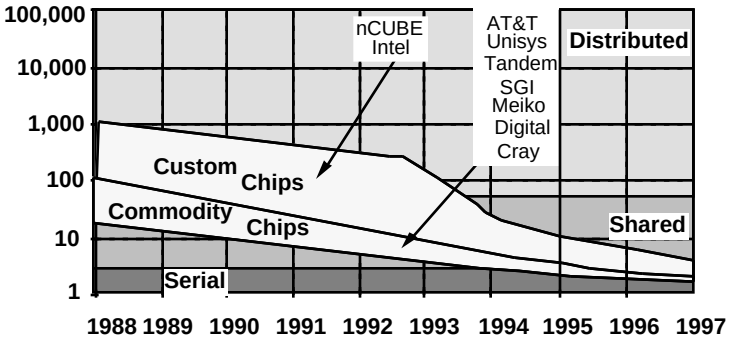

Average Microprocessor Clock Speed (MHz)

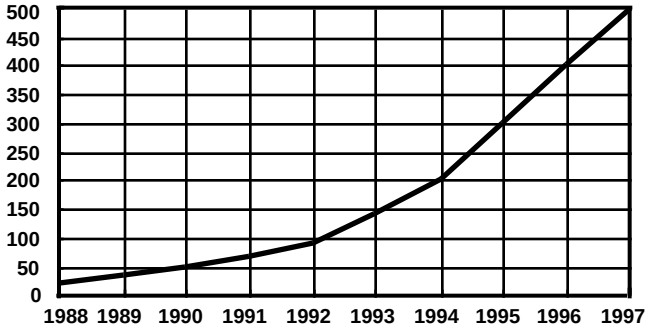

Source: Gartner Group

During the past few years, parallel-processing vendors have adopted commodity RISC technology in ever-increasing numbers. Desktop computing is so prevalent that few HPC vendors can afford the kind of massive microprocessor development efforts put in by workstation and PC vendors. Modern commodity RISC microprocessors are increasing in speed and capability at a predictable rate, and the trend toward faster processors is likely to continue for the foreseeable future. This relentless increase in processor speed has caused a fundamental change in the character of the HPC world. As processors become more powerful, the number of processors required to reach a given level of performance continues to decrease, eliminating the need to address massive parallelism, with all its characteristic difficult technical issues such as load balancing, interprocessor communications and parallel programming methods. In addition, specialized vector supercomputers themselves are in danger of becoming obsolete. The need for parallelism itself may eventually be called into question — a conclusion which is inherently flawed.

The need for increased computing and communications power for new commercial data-intensive applications will exceed the rate of RISC processor improvement (0.7 probability).

Reader Notes

Processors Required in a Typical Commercial Data-Intensive Application (1994 Relative)

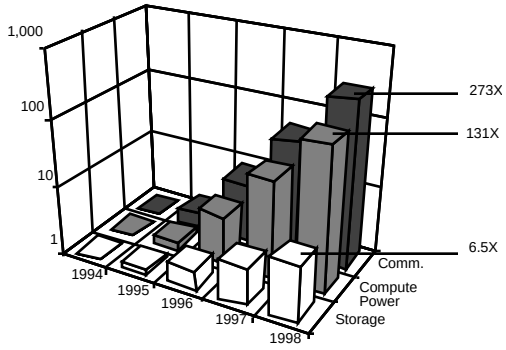

Assumptions:

Growth Rate Estimates	1994	1995	1996	1997	Overall
# of records	20%	20%	20%	20%	107%
data per record	0%	40%	50%	50%	215%
# of users	20%	50%	60%	40%	303%
# of queries per user	10%	30%	30%	30%	142%
# of records accessed per query	0%	40%	75%	75%	329%
RISC processor power	52%	52%	52%	52%	434%

Performance Required in a Typical Commercial Application (1994 Relative)

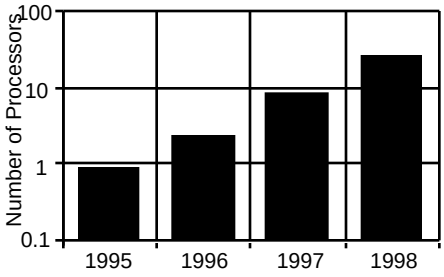

Source: Gartner Group

Key Issue: How will the relentless improvement in underlying systems technologies manifest itself in user choices for parallel and vector systems?

While increases in RISC processor performance continue to drive down the number of processors needed in a system, new applications and expanded uses for older applications drive the number of processors up. We can estimate the growth in performance requirements for a typical commercial MPP application by separately predicting the changes in different parts of the application, and looking at which elements affect database size, computing-power requirements and communications requirements. Even taking into account the fact that RISC processor speed nearly doubles every 20 months, computing requirements for commercial applications will require more than an order of magnitude more processors than are sufficient today. Parallel OS design and hardware infrastructure for the management of communication resources will eclipse processor speed as the prime determinant of price/performance.

Software and hardware advances will lead to increased application availability and the emergence of hybrid systems by 1996 (0.7 probability).

Reader Notes

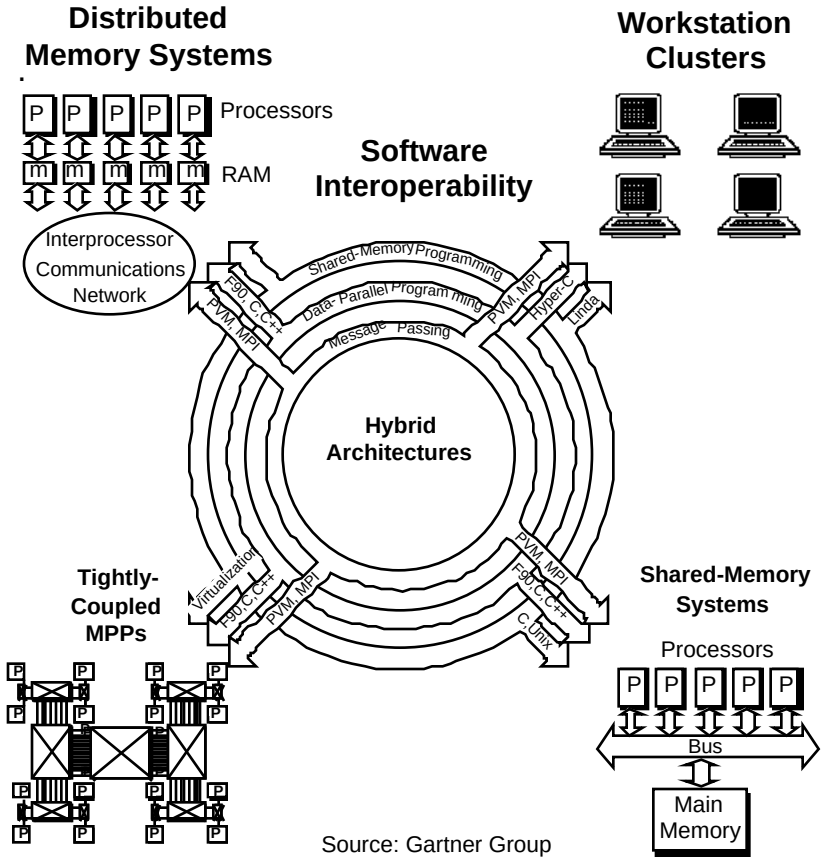

Key Issue: How will the relentless improvement in underlying system technologies manifest itself in user choices for parallel and vector systems?

Interoperability techniques are playing a greater role in parallel processing. Standardized message-passing APIs are allowing simultaneous access to a wide range of architectures. Virtual shared memory is being supported via both hardware and software, and sophisticated caching strategies are being used to automatically manage the movement of data. The vendor's desire to access the widest possible range of applications, markets and users is leading to an explosion of software interoperability and the hardware needed to support it. Out of the explosion of different architectures in the early 1990s, and the subsequent scramble for interoperability, has also emerged a fundamental realization: No single paradigm or architecture has been able to provide both performance and ease of programming over a wide range of scales. The correct architecture is that which will scale to the requirements of an application during its life expectancy at the lowest possible cost.

The need to rapidly create applications that are scalable, multiarchitectural and configuration-independent will lead to the emergence of third-party parallel applications management systems (PAMS) in 1996 (0.7 probability).

Reader Notes

Traditional Parallel Runtime Environment

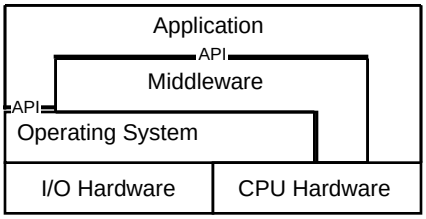

PAMS-based Parallel Runtime Environment

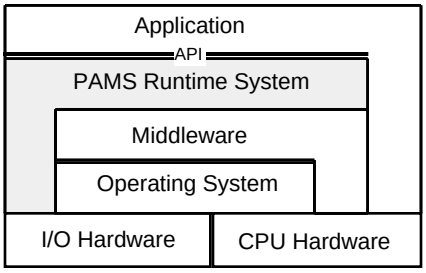

Source: Gartner Group

Key Issue: How will the relentless improvement in underlying system technologies manifest itself in user choices for parallel and vector systems?

Parallel applications software will require more capability per dollar, and incorporate more knowledge per man-year of development time, than ever before. The industry is moving toward an era where the the expectation level for the capabilities of competitive parallel applications will include: rapid development; linear scalability; the ability to run effectively on uniprocessors, symmetric multiprocessing systems, MPPs, networks of workstations and hybrid systems; and the ability to adapt dynamically to changing loads, communication patterns, database sizes and forms of connectivity. This challenge, combined with the migration of parallel expertise from failed hardware companies to software startups and systems integrators, will result in a new layer of software from a new type of vendor called PAMS, which will improve the applications development environment and provide greater application portability than is possible today.

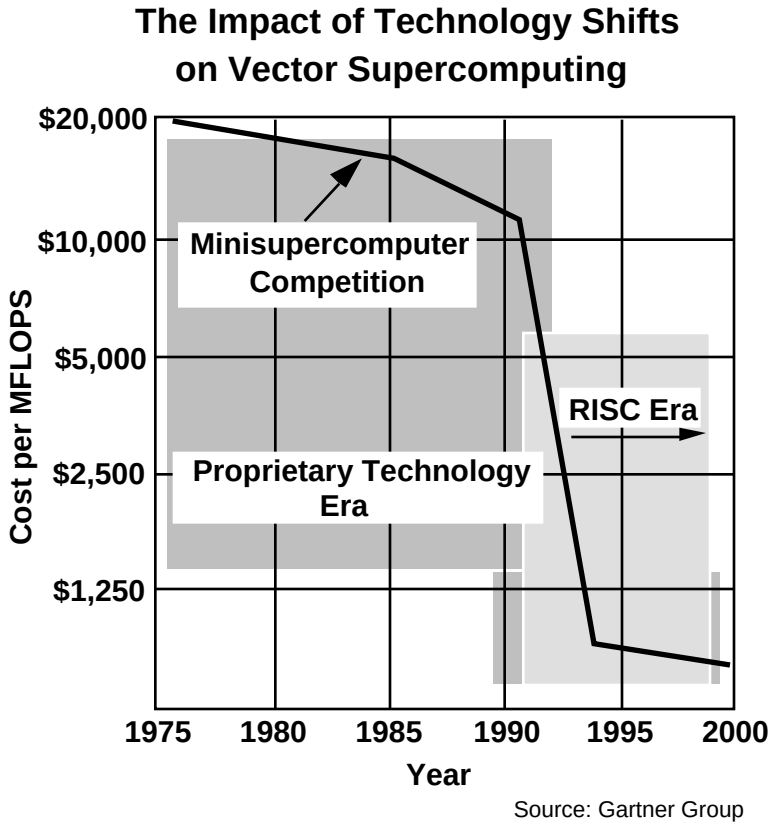

Source: Gartner Group

Proprietary supercomputing technologies from companies like Cray Research, Fujitsu and Convex provided four- to five-year product life cycles and only a gradual reduction in costs due to competition. It was not until the beginning of the 1990s, when advanced RISC architectures came into the market, that price/performance began to radically change. As the chart shows, there has been a hundredfold decrease in the cost per MFLOP during the past two decades, with the majority coming in the past several years. The potential of proprietary CMOS-based systems is clearly demonstrated by Fujitsu's VP-300 and NEC's SX-4. With the improvements in both performance and price/performance, the vector supercomputing market will have a new lease on life: as small, networked servers that take advantage of the application base built over the past two decades. How long this phenomenon will last is open to question by prospective users as well as their vendors.

Technology advances will have enveloped the processing requirements of most vector applications by the end of 1995 (0.7 probability).

Reader Notes

The Impact of Technology on Application Processing

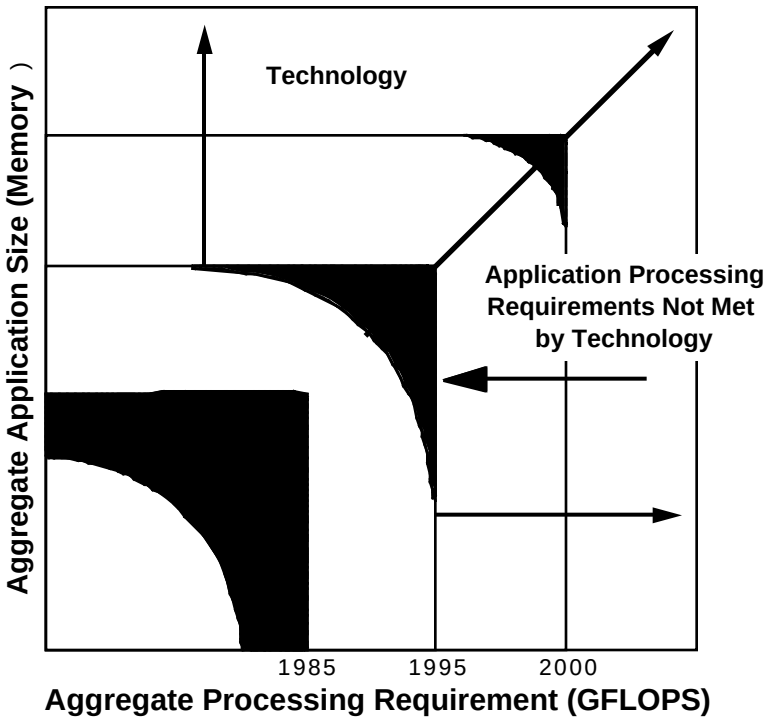

Source: Gartner Group

Key Issue: What is the evolving role of traditional vector supercomputers?

Most of the applications that run on large, specialized supercomputers today have their origins in the early 1980s. In the chart above, the horizontal axis represents aggregate processing requirements in GFLOPS, the vertical axis represents application size in memory and required bandwidth, and the shaded area represents the relative requirements for traditional supercomputers. The rapid developments in RISC microprocessor and memory technologies have overtaken the growth of most applications in the past decade. Until there is a major development in application requirements, traditional supercomputers will process those legacy applications that have been designed for a particular production environment. Even the rapid improvements in price/performance experienced in the past two years will provide only a brief respite from the relentless march of RISC developments.

How will parallel systems impact commercial data-processing solutions in the 1990s?

Reader Notes

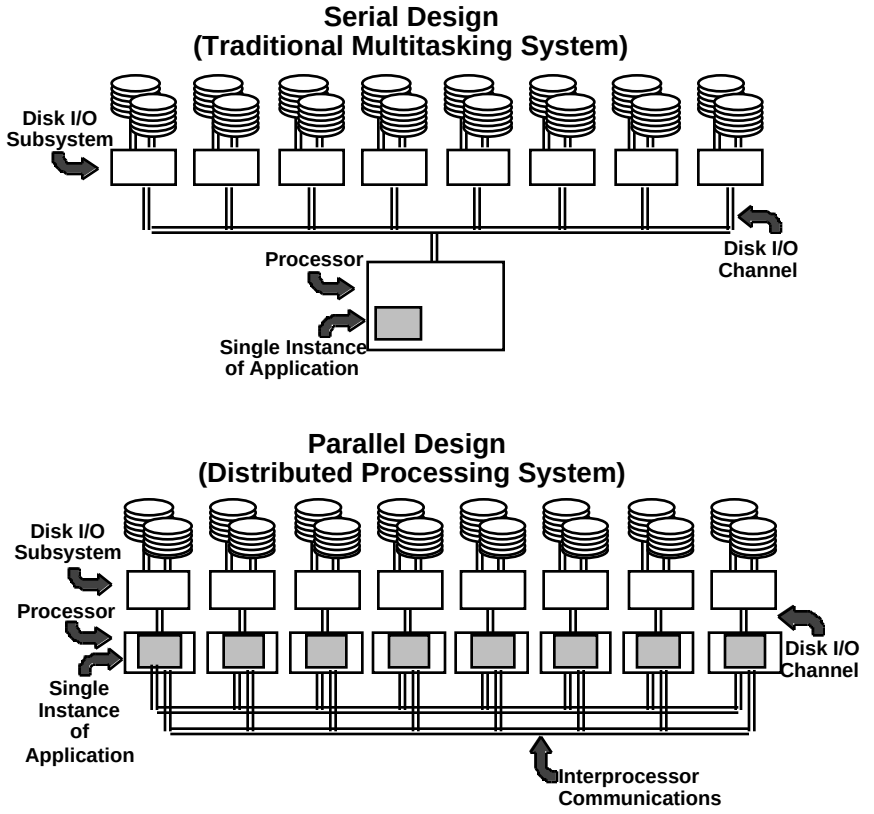

Source: Gartner Group

Traditional general-purpose commercial and scientific systems were designed to process as much work as possible through the system to optimize the price/performance. The operating systems were engineered to process large numbers of applications in a robust and reliable manner. These traditional systems were not designed to bring parallelism to the individual user application. Individual applications are processed in a serial fashion and are limited in their ability to move large quantities of data through the application itself. Parallel systems can support data bandwidths of 100 to 200 (or more) Megabytes of data per second to the application, as well as hundreds of MIPS, providing an order-of-magnitude improvement in performance.

Parallel systems will have a greater impact on commercial, data-intensive applications than on scientific applications, due to the availability of parallel relational-database technology to insulate the end user from the complexities of parallel programming. In scientific applications, modest levels of parallelism will be the rule until new applications are developed, which will require the performance characteristics of hundreds of processors.

By 1996, advances in symmetric multiprocessing (SMP) systems will enable greater application scalability than previously possible (0.7 probability).

Reader Notes

Scalability

Performance vs. Capacity/Cost

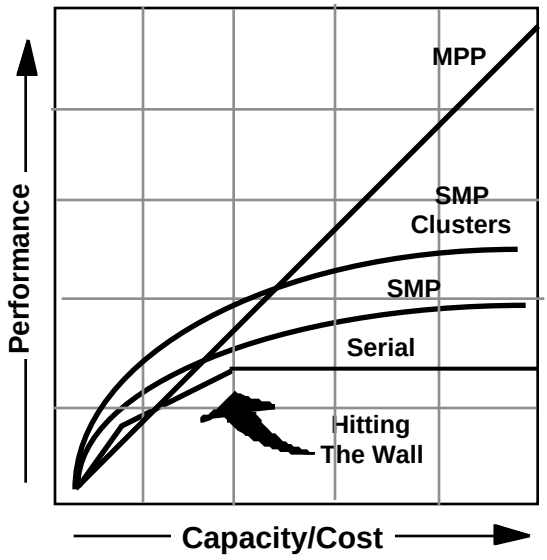

Source: Gartner Group

Key Issue: How will parallel systems impact commercial data-processing solutions in the 1990s?

The primary difference in SMP and MPP systems is their ability to scale to the requirements of the application. Only a few years ago, SMP systems were limited to a few processors. Today, advances in hardware architectures and operating systems are supporting systems of up to 64 processors. These advances are expected to continue, which will enable users to deploy SMP systems in support of extremely large applications. At the same time, advances in MPP systems will continue, and these systems will continue to provide the highest degree of scalability for extremely large applications.

Which HPC vendors will adapt to the recent market changes? How will they adapt? Which will prosper and which will not?

Reader Notes

The Quick, the Dead and the Hopeful

Status	Vector Super-computers	Vector Mainframes	Mini-Super-computers	Array Processors	Scientific Parallel	Commercial Parallel
The Quick (42)	Convex Cray Research Fujitsu Hitachi NEC SGI	Bull Fujitsu Hitachi IBM NEC Unisys	Convex Cray Research	Analogic CSPI Mercury Numerix Sky	Convex Cray Research Cray Research Superservers Digital Alpha Farm IBM Intel MasPar Meiko nCUBE SGI	AT&T GIS Cray Research Superservers Digital Encore HP IBM Meiko nCUBE Pyramid/SNI Sequent SGI Tandem White Cross
The Dead (44)	Ametek CDC Chopp Cray Comp. Denelcor ETA IBM SSI	Amdahl CDC Digital Key Trilogy	Alliant AM Super Astronautics Celerity Culler Digital ELXSI FPS Gould Prisma Saxpy SCS Supertek	FPS Starr	ACRI AMT BBN Cydrome E & S FPS Goodyear ICL DAP Kendall Square Multiflow Myrias Suprenum Thinking Machines	ACRI Kendall Square Thinking Machines
The Hopeful (8)					Cambridge	Amdahl Convex dataCACHE ICL Intel MasPar Unisys

Source: Gartner Group

The above chart contains lots of names — too many, in our opinion. With a lack of third-party applications available for a number of vendors, the industrial market will not be able to support all of them. The winners will be those vendors that bring robust operating environments and rich suites of applications to the market. Government initiatives and purchases sustained a small number of HPC vendors, but reductions in government funding have placed in jeopardy those vendors that depended on that segment of the market. The private sector will determine the overall health of the industry. Industrial users are less enamored of peak "machoflops" than of balanced, general-purpose systems that include competitive price/performance, a broad variety of user-friendly application programs, a comprehensive suite of high-quality programming tools, good customer support, and a sound technology growth strategy. Those vendors offering only the latest "hot" technology with limited distribution capabilities will not have a long run in the marketplace, as was clearly demonstrated by the short life span of most mini-supercomputer vendors in the late 1980s, and of many scientific parallel vendors in the 1990s.

Innovative CMOS and RISC designs will eclipse traditional proprietary systems from Cray Research and Convex by 1996 (0.7 probability).

Reader Notes

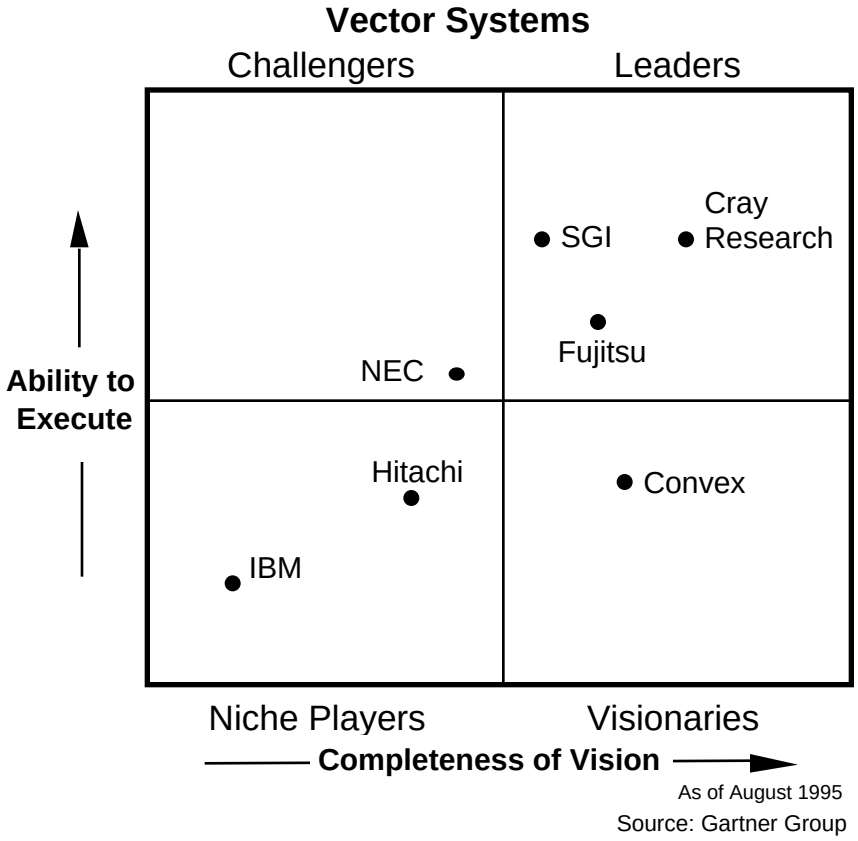

Key Issue: Which HPC vendors will adapt to the recent market changes? How will they adapt? Which will prosper and which will not?

For the next few years, the principal advantage of traditional vector HPC systems over their challengers will be relative ease of use. This may take the form of a library of proven applications programs, a portfolio of well-polished programming and system management tools, or a combination of these. It will almost certainly include the ability to run users' legacy applications that have been developed and optimized to particular environments. New and less-expensive vector systems are coming to market that utilize innovative CMOS and commodity RISC processors, offering users far-less-expensive alternatives to traditional systems.

Cray Research will find its lead challenged by vendors such as Silicon Graphics and Fujitsu initially, and then by a larger number of SMP vendors over time as the new technologies finish enveloping the current application base.

As the commercial data-intensive market develops, successful vendors will have to develop robust systems as well as strong distribution capabilities and alliances (0.8 probability).

Reader Notes

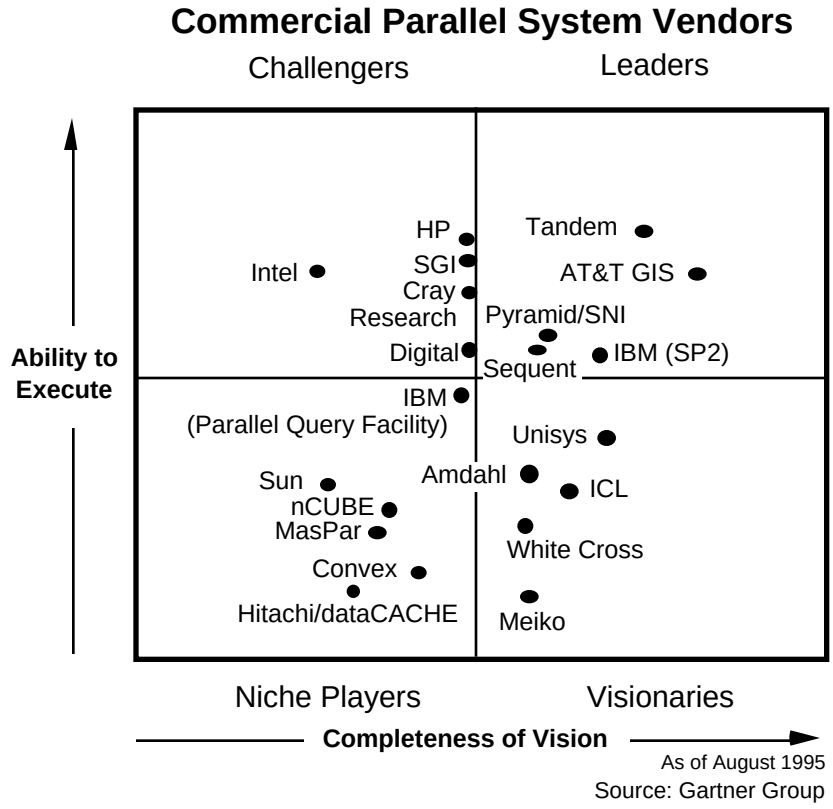

Key Issue: Which HPC vendors will adapt to the recent market changes? How will they adapt? Which will prosper and which will not?

With the rapid advances in parallel relational-database technology, large numbers of vendors, both SMP and MPP, have entered the rapidly developing commercial market. Smaller, visionary vendors that pioneered in the scientific market are finding it extremely difficult to compete against large and well-established companies such as AT&T GIS, Digital, HP, IBM, SGI and Tandem. Those vendors that remain hardware-oriented will find it difficult to survive in a conservative IT market with high expectations on robustness and application availability.

The next year will see a change in positioning, as more vendors demonstrate the ability to support and process terabyte-level databases, a key factor in determining leadership.

Successful parallel-system vendors will have financial staying power, good marketing, and a clear focus on application availability (0.8 probability).

Reader Notes

Scientific Parallel System Vendors

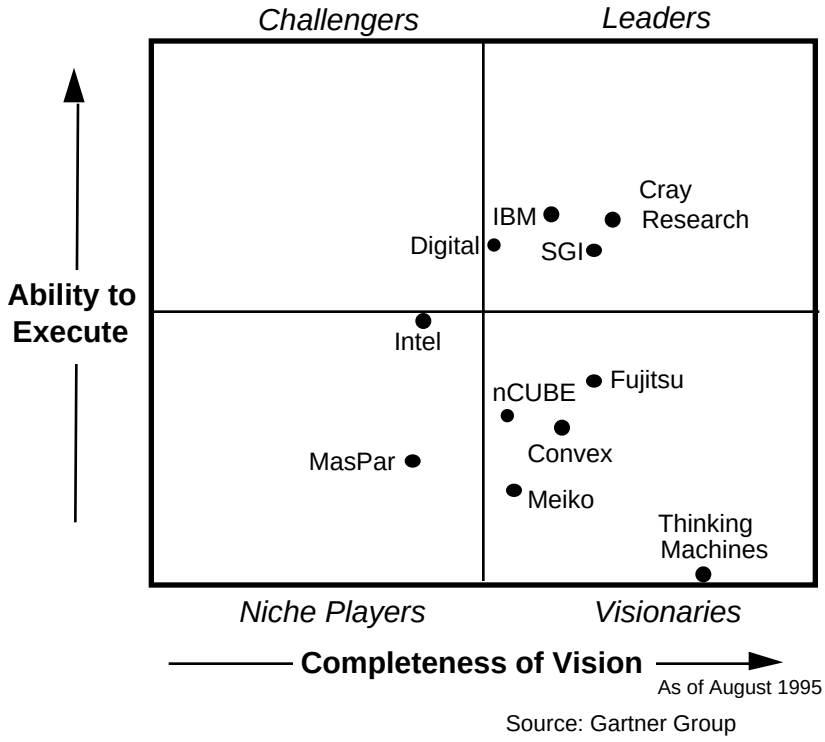

Source: Gartner Group

Key Issue: Which HPC vendors will adapt to the recent market changes? How will they adapt? Which will prosper and which will not?

One of the defining shifts in the scientific parallel market over the past several years has been the dominance of industry's use of parallel systems as opposed to government-funded procurements. With industry now accepting these systems, the more-generalized multipurpose parallel computers have been the systems of choice. In addition, major vendors such as SGI and Digital have shifted their focus from technology to applications availability and distribution capability. Smaller vendors are being forced to concentrate on niches where they can bring sufficient value in order to survive.

The role of governments will not be significant in deciding who wins and who loses. Those decisions will be made in the commercial and industrial market, forcing some vendors to become footnotes in computer history books.

How will users deploy HPC systems within the existing IT infrastructure?

Reader Notes

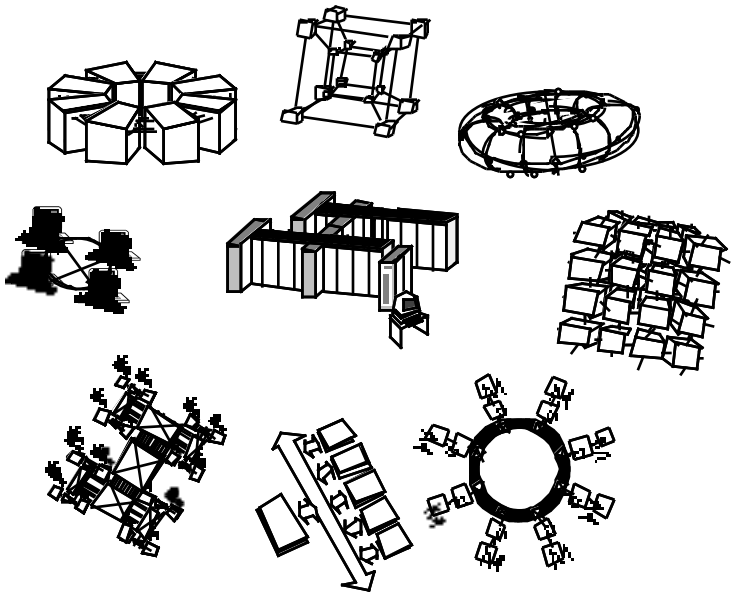

Since the early 1980s, HPC systems have been seen as specialized systems serving a relatively small number of users. These systems have generally been placed in separate organizations and physical facilities. With ongoing pressures to reduce costs through reorganization and through a reduction in the duplication of resources — and with the rapidly growing use of HPC systems by commercial users — the need to integrate these systems into the existing IT environment cannot be avoided. Will these systems replace or complement existing systems? How should users view the challenge of integrating these technologies?

The challenge for vendors and IT executives is to better integrate HPC systems into the existing environment to permit broader use in advanced applications.

HPC systems will be integrated into existing IT environments in a client/server paradigm using 100-Mbyte/sec. backbones (0.8 probability).

Reader Notes

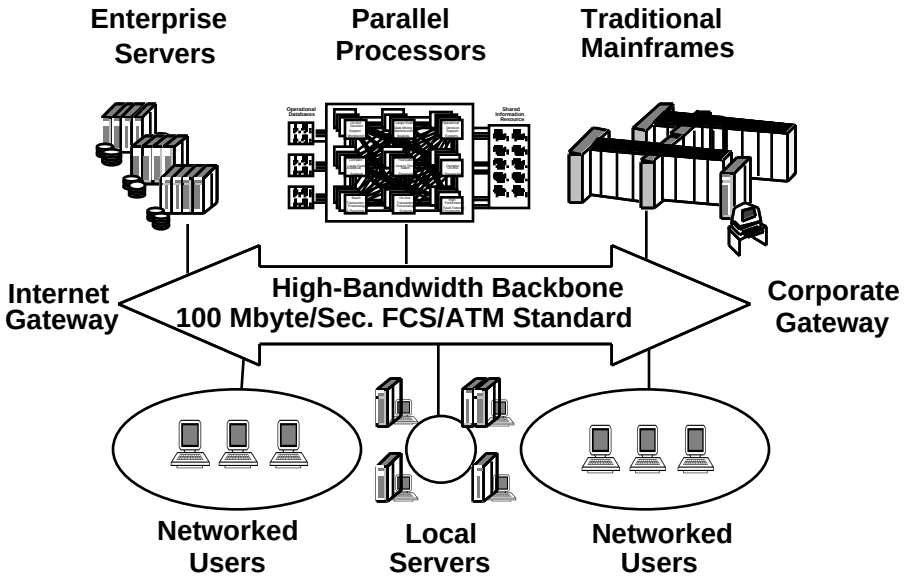

Source: Gartner Group

Key Issue: How will users deploy HPC systems within the existing IT Infrastructure?

With the rapid emergence of parallel-processing systems into the commercial data-intensive market, the challenge is integrating disparate system architectures within the existing environment. Traditionally, vendors have offered various channel connections to provide access to the legacy systems. With the channel connections and appropriate software, data can migrate between systems as needed.

In the future, high-bandwidth backbone networks combined with client/server application implementations will emerge as the primary interconnection strategy of choice. This evolution will enable users to take advantage of a broader choice of systems than is possible today, better matching the use to the technology.

- The need for increased computing and communications power for new commercial data-intensive applications will exceed the rate of RISC processor improvement.
- Software and hardware advances will lead to increased application availability and the emergence of hybrid systems by 1996.
- The need to create applications that are both scalable and configuration-independent will lead to the emergence of third-party parallel applications management systems in 1996.
- Advances in SMP systems will enable greater application scalability than previously possible by 1996.
- Technology advances will have enveloped the processing requirements of most vector applications by the end of 1995.
- Innovative CMOS and RISC designs will eclipse traditional proprietary systems from Cray Research and Convex by 1996.
- Successful parallel vendors will have financial staying power, good marketing, and a clear focus on application availability.