



Creating the Possible

Architectures for Electronic Business

intel®



*November 1999
Intel Corporation*

Abstract

The Internet is changing the way enterprises conduct business. Virtually everyone agrees that this change is instrumental to business processes as well. With the Internet dynamic comes opportunity and a whole new world of possibilities. Information technology organizations will be challenged to create architectures and infrastructures that can take full advantage of the opportunities provided. Yet today no one can articulate what will be possible, and in fact, this approach may not be the best way to consider the problem.

This paper discusses the emergence of the IA-64 architecture and its importance to electronic business (e-Business), in terms of its capabilities combined with Intel's vision of an agile, flexible IT organization that can "create the possible" in this rapidly changing environment and capitalize on all the opportunities presented by the changing Internet environment.

Electronic Business via the Internet: the Market

Technology is clearly changing, and it continues to change the world. Over its history, information technology (IT) has become a core business regime. At first, IT helped automate existing business processes. Today, IT helps with the evolution of richer software, faster hardware and economical, high-speed networks to enable truly global enterprises.

The growth and success of the Internet and its applications have changed our view of information and interaction. Originally a venue for electronic mail and file sharing, the Internet has provided significantly increased value to businesses and consumers.

Table of Contents

Abstract	2
Electronic Business via the Internet: the Market	2
The Unknowable: Business in the Internet Age	3
Creating the Possible	3
The Firehose, the Fortress and Chain Reactions: Elements of the New Architecture	4
<i>The Firehose: Cooling Off Bandwidth Hot Spots</i>	4
<i>The Fortress: Keeping the Keys Safe</i>	5
<i>Chain Reactions: Objects May be Closer than They Appear</i>	6
The Engine Inside the Internet Economy	7

Yet much more value still lies ahead. The notion of using the Internet for business-to-business transactions is now a thin sliver of the overall market, but it is expected to demonstrate the most rapid growth.

Figure 1 shows that today marks the beginning of an enormous expansion.

Recent updates to these numbers indicate that the \$1 trillion mark may be surpassed as early as 2002.

This trend means not only a significant change in the way business is conducted, but also that IT will be essential in this market shift. Businesses will be challenged to establish competitive

advantage with the advent of e-Business and this is clearly dependent on unconstrained IT infrastructures. Success in this market will be directly linked to how well IT organizations can respond to this challenge. The industry has witnessed instantly-created markets, communities of interest, and trade based more on the model of auctions and bazaars than based on catalogs and showrooms in today's marketplace. The marketplace of tomorrow holds many more opportunities for business.

The Internet's inherent ability to instantly link individuals and concerns has great immediate value. As a universal mechanism for one-on-one communications, where every conceivable

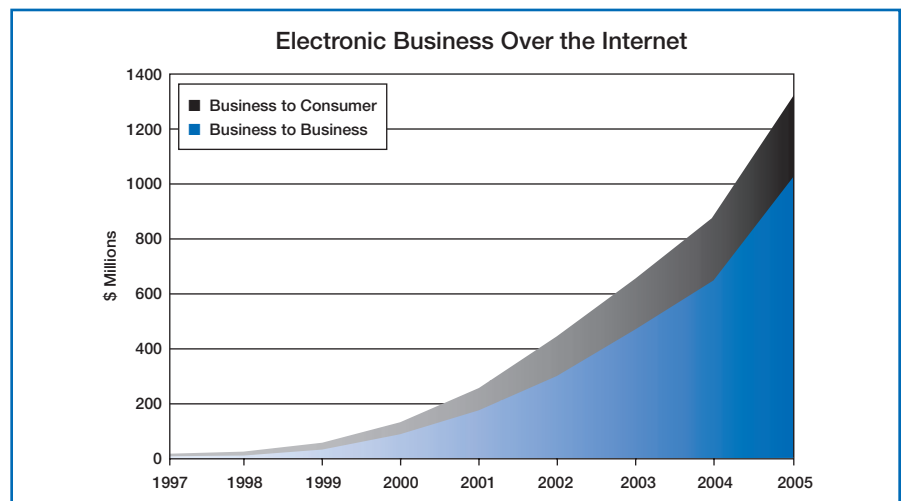


Figure 1: The Emerging Digital Economy, U.S. Department of Commerce, 1997

organization can provide its unique presence and value proposition, the Internet represents a businessperson's paradise of choice and opportunity.

The specifics of the business-to-business marketplace are yet to become clear, but as Intel Chairman Andy Grove recently said: "96 percent of e-Business servers have yet to be deployed."

In other words, less than 4 percent of the servers that are expected to be necessary to service the Internet have been deployed to date. This prospect raises several questions:

- How can IT best respond to the challenge and create value in this trillion-dollar market?
- How do architectures such as IA-64 fit into this new e-Business environment?
- How does IA-64's architecture accelerate change by creating opportunities of capacity, performance and value?

Put simply:

IA-64 architecture features * Industry Solutions
= e-Business success

The Unknowable: Business in the Internet Age

The ability to move bits, not atoms¹ from one place to another more economically and quickly has presaged each inflection point in information technology.

Today the Internet offers ubiquitous, reliable and responsive access to information. The challenge is that little is known about the specifics of what will occur, except that certain events will happen, and they will happen with unprecedented speed.

The business outlook that this fosters could be called the Unknowable:

¹The term "bits, not atoms" is widely credited to Nicholas Negroponte.

- User populaces whose size, needs and demands on internal systems are unpredictable
- Rapidly evolving partnerships, alliances and communities of interest
- Trading networks formed to take advantage of spot markets, bid on contracts for commodity goods and all manners of intermediaries
- New and valuable Internet-based services (killer applications)

This leads to new and unprecedented demands on IT architectures:

- The need for agile systems that can be expanded and reconfigured quickly to adapt to changing needs
- The requirement for impeccable reliability of systems and for reducing or eliminating scheduled downtime
- The need for logical security, risk management and safeguarding of information from prying eyes
- The requirement to deploy systems that can recall data very quickly and turn it into information.

Creating the Possible

The successful e-Business will take advantage of market shifts and turn uncertainty into a core competence. Successful e-Business will be enabled by new system architecture that can reliably and economically scale and adapt as the enterprise benefits from the Internet as a business venue.

These new IT architectures will present themselves to business partners as a library or catalog of published business processes and capabilities "at the edge" of the business' electronic enterprise. Similarly, the successful e-Business will subscribe to business processes provided by its partners.

The publication of new processes and the subscription to partners' processes will allow for and encourage enormous dynamism within this growing market. Certainly, no two customers or partners are likely to see or use these tools in the same way. Thus, this new model will spawn the mass customization not just of products, but of business relationships also.

In a practical sense, customers will have more and more power to devise the kind of solution they desire, by mining the wealth of information available regarding potential solutions to their particular problem and combining it with the published business processes that most apply to them. The ultimate relationship is unique but not expensive to any of the parties involved.

From an IT standpoint, this means creating the possible and organizing the mechanisms by which these possibilities can be created and used quickly and securely. Such published services must be unimpeachably reliable and responsive, because the business will suffer instantaneously if they are not.

The Intel® IA-64 Architecture will be ideal at the center of IT, as it represents the next generation in processor design, a new design philosophy and a set of principles.

Unprecedented speed and predictability:

- The IA-64 architecture represents a new approach, based on the Explicitly Parallel Instruction Computing (EPIC) design philosophy that provides the scalability and design durability well into the 21st century
- The architecture's unique approach to marrying compilers with the chip provides for welcome flexibility and adaptation to particular application scenarios, without requiring overt effort from the programmer

Reliability in the core of the microprocessor:

- Error correction on data and memory paths
- Power-up self-testing
- Monitoring of main memory for multiple-bit memory errors
- Integration with the operating system for logging of and possible recovery from errors that typically brought a system down in the past

Thus, while it is impossible to know all the ramifications of e-Business, it is possible to set a course that promises a broad range of software, systems vendors and architectural topology support. The course can be charted with Intel's IA-64, based on the EPIC architecture.

The Firehose, the Fortress and Chain Reactions: Elements of the New Architecture

The demands that the Internet can place on critical business functions are already well documented. The failure of several highly publicized electronic business systems due to wildly growing workloads seems to be a weekly or daily news item.

Initially, the failure of an Internet service meant simply that access to information was lost. However, as businesses rely increasingly on the Internet as part of their financial strategy, system failures due to surging, unpredictable workloads are unacceptable, expensive and a real threat to the life of the business.

Similarly, systems that do not inherently contain a high level of integrity checking, down to the processor core, cannot be considered "ultimately reliable."

Given that the Internet will place great and unpredictable demands on the IT infrastructure, several things are certain:

- There's no such thing as excess capacity anymore.
- The economy of such capacity is important. Paying a premium for proprietary systems is no longer acceptable and justified.
- Reducing the number of dissimilar, proprietary architectures and platforms will increase agility and reduce support and acquisition costs.
- Software must be available on the chosen platform, including operating systems, applications and middleware "glue."

Intel's IA-64 Architecture offers a way to address each of the following issues outlined in Table 1.

The Firehose: Cooling Off Bandwidth Hot Spots

Waiting for information is irritating. To a connected customer, too much waiting around may lead to going somewhere else. Competition is just a click away.

Hot spots—where bandwidth requirements exceed that which is available or economical—inevitably occur when creating systems architectures. Extinguishing such hot spots involves specialized software or hardware solutions.

Smart Caches, Fast Directories and Big Databases

Being able to quickly and reliably deliver information is key to success on the Internet. In some cases, waiting around for slow, secondary storage mechanisms such as disk storage can exacerbate bottlenecks. Thus, it's necessary to keep the most frequently accessed information for a particular service in memory.

Capacity	<ul style="list-style-type: none"> ■ The EPIC design is one that will get faster through instruction-level predication and advanced speculation techniques. ■ As each generation of EPIC processor arrives with more capabilities, existing applications will have the opportunity to further extend their scalability and performance capabilities. ■ In multiprocessor systems, EPIC enhances processor scaling by reducing unnecessary memory and inter-processor contention. ■ IA-64 is designed to handle large amounts of high-speed, economical memory.
Economy	<ul style="list-style-type: none"> ■ Intel's volume economics for IA-64 will deliver excellent value. ■ Intel will support a wide variety of systems vendors by providing integrated system building blocks.
Reliability	<ul style="list-style-type: none"> ■ Internal error-checking and control will provide systematic, end-to-end systems reliability. ■ Enterprise-class operating system vendors have committed to provide solutions based on IA-64.
Architectural Similarity	<ul style="list-style-type: none"> ■ IA-64 offers broad system vendor choice, from workstations to massive servers.
Software Availability	<ul style="list-style-type: none"> ■ Intel has made a commitment to provide investment capital, software support (porting and optimization) centers and individual support to independent software vendors to promote the availability of optimized and reliable software. ■ IA-64 will enjoy a broad choice of solutions in terms of operating systems, applications and supporting software.

Table 1: The IA-64 architecture addresses the demands of the Internet

This requirement may mean many gigabytes of high-speed memory systems connected to equally fast processors accessing it, plus software groomed to provide very quick access. For example:

- Internet caches to reduce pressure on the core Internet connection
- Directories to catalog and articulate corporate resources, business methods, authorization, security catalogs, personal profiles and other business-critical information
- Databases managing very large caches to keep pertinent information close at hand

The characteristics of systems that can best be part of these “firehoses” are:

- Fast processor speed
- Large memory with high bandwidth and low latency
- Ability to interface seamlessly with the surrounding network

The design of the IA-64 architecture includes unique features intended to keep the instruction streams executing, which is key to accessing and returning information very quickly:

- **Predication.** Predication is the ability to identify certain condition inline, rather than relying on inefficient, processor-stalling test and branch situations. Instead of requiring a branch in situations, the test for a condition and setting of a variable occur inline. This optimizes “if-then-else” scenarios, as well as “case” statements, among others. This architecture feature, along with the massive register resources contained within EPIC, minimizes or eliminates costly memory accesses. Yet, where predication is not a good approach, standard test and branch can be used, at the compiler’s option. One classic problem with RISC architectures is that they are groomed for straight-line

execution and branches, which are common in business code such as databases. This quickly drains away the processor’s ability to do productive work. EPIC avoids this type of problem.

- **Speculation.** Although speculation has been common on the chip itself, it has not always provided a benefit. On-chip speculation has been used by RISC architectures to keep the instruction stream moving along by looking at alternative strategies and making guesses about uncommon vs. common branch scenarios. The EPIC philosophy places responsibility for speculation on the compiler, providing unprecedented ability to groom code sequences automatically for particular scenarios. In multiprocessor systems, for example, fine control over speculation can reduce or eliminate unnecessary inter-processor memory contention, the main reason why multiprocessors do not “scale” given a particular set of software. IA-64 processors and this fine, automatic control allow the scenario for scalability.
- **Parallel instruction execution.** The way IA-64 instructions are packaged enhances opportunities for parallel execution. Instructions are not only packaged several to a single word, but also grouped together so instructions can be executed in parallel. This approach has immediate benefits because the IA-64 architecture already contains multiple internal resources such as execution and mathematical units. Subsequent generations of IA-64 processors, which may have substantially more such units, will be able to run even more instructions in parallel.

These features help make systems built with IA-64 processors responsive and fast. In such branch-intensive or otherwise pathological workloads, the performance improvements

should be more noticeable. This will occur not just in the example of the “firehose”, but also in many other real-life application environments.

The Fortress: Keeping the Keys Safe

Central to administering security is a secure fortress to contain vital information such as private keys, certificates and other information that authenticates an enterprise as being precisely what it is.

Not necessarily a specialized system, but highly secured and audited, the fortress manages the security infrastructure of the enterprise in terms of Public Key Infrastructure.

These fortress systems must be highly responsive when replying to authorized requests for the information they keep. They must also be highly reliable, with end-to-end integrity checking and redundancy to promote component survival.

Secure connections today have a reputation for being slow. A server that can handle hundreds of simultaneous connections is typically brought to its knees when encryption is brought into play, particularly the time required to set up a connection. Similarly, the additional mathematical burden of processing the clear data stream into an encrypted one creates more overhead just to pump data out on the network or receive it from another trusted source.

As encryption-cracking techniques become more sophisticated, reliance on algorithmic methods of random number generation is considered risky and prone to attack. At present, the only way to clear this bottleneck is to follow an expensive path: proprietary encryption processors based on custom logic that may not be able to scale or become economical over time.

Systems that can clear this bottleneck and provide robust security must:

- Provide for world-class mathematical throughput
- Scale, including support for multiprocessor capabilities
- Contain extensive internal error-checking and correcting to eliminate the chance of errors
- Generate random numbers that are truly random as keys are computed

The IA-64 architecture is constructed with significant mathematical capability in mind:

- World-class mathematical performance.
- The ability to calculate equations such as $a = b * c + 10$ with a single instruction.
- The ability to execute mathematical instructions in parallel across multiple mathematical units. With the EPIC philosophy's notion of creating instruction groups that can be executed in parallel, subsequent generations of IA-64 will further extend the architectural parallelism and performance capabilities.
- Massive internal resources, including 128 floating-point registers and a similar number of general-purpose (integer) registers.
- An innovative random number generator that doesn't rely on an algorithmic, "pseudo random" approach.

Chain Reactions: Objects May be Closer Than They Appear

As Nathan Myhrvold² has said, "Software is a gas: it expands into the space given it."

Every advance in processor and systems capacity has been consumed by richer, more complex software. Software development and

programming environments are no longer built strictly around optimizing the number of bytes of memory and the number of processor cycles required to solve a particular problem (nor should they be). Object-oriented systems are certainly the latest evidence of this.

Clearly, objects will play a central role within the IT enterprise in the future. Over the past several years, object programming has proven to be extremely valuable in increasing programmer productivity. Although object programming was initially confined to complex tasks such as user interfaces, the wide acceptance of object frameworks and languages (particularly Java[®]) as application platforms has greatly improved programmer productivity on server-based applications as well.

The notion of interfacing and extending business processes via an object framework lends itself to flexibility both within and without the enterprise. Instead of having many disparate application islands, IT can publish a directory of business objects and methods.

With appropriate security, users inside the enterprise and partners outside can see the portions of the directory that are relevant to them and, in essence, subscribe to them. This approach allows for mass customization of the electronic relationship between the enterprise and its partners.

In fact, this notion can alter the idea of a linear "supply chain" or "value chain." Namely, instead of a chain of supply or a chain of values, an enterprise can embrace a matrix or lattice of supply, with many enterprises teaming up to provide one another with a constantly changing series of interactions and value creation.

This model has profound meaning for consumers as well. As a user's desktop computer becomes an equal part of this lattice of trade, he or she can act at an enterprise level in creating individual, customized products instead of accepting one-size-fits-all solutions for services such as travel, finance and general information.

Where, then, does the concept of the "client computer" end and the concept of the "server computer" begin? Soon, the line may be hard to distinguish; in fact, it may become irrelevant as large, capacious computers, all with fast connections to the Internet, interact.

This trend is likely to create significant uncertainty for IT architectures in the following ways:

- The user populace may not be entirely known.
- User demands may surge at any moment.
- Long-lived transactions (and consumed resources) may be implied by complex business interactions.
- Modeling resource consumption for a new object type may be impossible. Even the best-designed object environments can be highly non-deterministic.
- As demanded by a rapidly changing e-Business landscape, new business objects will be built and published with increasing frequency.

Although this level of uncertainty creates a seemingly impossible problem for IT, the successful e-Business enterprise will quickly realize a third dimension to Metcalfe's³ law. Specifically, the utility of the network is not just the square of the number of connected devices; it is that times the number of available

²Nathan Myhrvold is the Vice President of Microsoft Research.

³Robert Metcalfe suggested that the utility of a network is the square of the number of connected devices.

services and exposed business methods. This updated law reflects the third dimension of the Internet and highlights the area where true e-Business opportunity exists.

Object programming has clearly delivered the benefits expected of it, but it provides some significant challenges to the IT enterprise. Distributed, object-oriented, event-driven applications are making way for the distributed, object-oriented, event-driven enterprise. In practice, this resembles a chain reaction, with one event causing one or many more, ending finally with the information or transaction that the business partner desires.

Similarly, unlike prior generations of IT architectures, which were built on well-reasoned and highly tested environments, the object enterprise will not permit the kinds of capacity testing previously done before production. In fact, the "time-to-market" requirements make testing essentially impossible.

The number of events and object invocations can't always be precisely known at the initiation of a request, and they depend on an equally unpredictable initial state of the system. Thus, satisfying a simple request may require significant and perhaps long-lasting consumption of system resources.

In particular, languages such as Java, C++ and associated object frameworks benefit from:

- Large numbers of internal registers to manage data-flow-oriented processing
- Large address spaces to manage the instantaneous creation, management and consumption of object states
- Fast processing speed to support "virtual engines"

Further, these environments have suffered because their penchant for causing an explosion of object invocations causes significant memory thrashing, and their exception-handling generally requires significant amounts of testing and branching.

Many factors make the IA-64 architecture the natural choice for such an environment:

- IA-64 contains large register resources tailor-made for keeping object state on the chip rather than requiring constant references to memory. RISC architectures are usually register-poor in coping with object environments, requiring unnecessary access to memory.
- IA-64 supports a large, flat address space to contain the state of millions of active objects.
- IA-64 can process stack-oriented operations on-chip by allocating registers as a stack (with automatic spillover and retrieval from memory to eliminate the chance of a stack overflow). Because the stack does not reside in main memory, deep recursion can be managed on-chip.
- IA-64's ability to cope with the "throw and catch" method of exception-handling, possibly requiring many layers of delegation, is unique and powerful in object environments, particularly Java. RISC processors must rely on old-fashioned, costly test and branch.
- Predication permits rapid identification and avoidance of exceptions without requiring processor stalling tests and branches. RISC processors predicate based on fixed logic on the chip, whereas IA-64 permits speculation at the appropriate time, place and aggressiveness. The IA-64 predication model can be tuned for specific needs simply by changing the instruction sequence.
- Speculation under an exception-handling regime (that is, being able to guess what truly is an exception vs. just a potential

change in control) means that the compiler and processor can be very aggressive in taking otherwise risky chances on speculative execution

The preceding characteristics make IA-64 the architecture of choice in highly object-oriented, distributed environments—environments that, as business relationships change, will infer more rich and complex business processes. That, plus the object environment's penchant to be non-deterministic, means that there will be no such thing as excess capacity.

For business people preparing for an unpredictable future, the prudent course may be to invest not only in the technology designed to cope with these environments, but the technology that will, consistent with Intel's business philosophy, also be the economical solution.

The Engine Inside the Internet Economy

Clearly, the enterprises that will thrive in the e-Business age will be those that can create the possible, and continue to create the possible by adapting, adopting and producing value quickly. Parallel to these requirements will be the need to deploy highly reliable, scalable and agile systems built on widely available components or solutions.

Intel has built the IA-64 architecture and all its surrounding components to focus on the center of the enterprise, addressing core IT requirements while providing for unprecedented scalability, well into the 21st century.

In summary, the best chance for e-Business success will be enterprises that work to create the possible by actively extending their enterprise onto the Internet. These enterprises will endeavor to "know the Unknowable"—to unleash growing and demanding user populaces with limitless

Transaction-processing Performance	<ul style="list-style-type: none"> ▪ Scalability in multiprocessor configurations ▪ On-chip parallel execution of instructions
Business Intelligence and Massive Database Support	<ul style="list-style-type: none"> ▪ Large memory addressability to process huge data sets ▪ World-class mathematical (analytical) capabilities
Ideal for Object-oriented Enterprises	<ul style="list-style-type: none"> ▪ Register model ideal for managing object workloads ▪ Large memory for managing vast numbers of long-lasting objects
Software Catalog	<ul style="list-style-type: none"> ▪ Broad, enterprise-class operating system support ▪ Application software solutions
Reliability and Availability	<ul style="list-style-type: none"> ▪ Health and welfare checking ▪ Internal error-checking and correcting of internal data and memory paths ▪ Extensive work with systems providers to produce end-to-end reliability. ▪ Integration with major operating systems to provide self-healing, fault-isolation and repair
Security	<ul style="list-style-type: none"> ▪ Mathematical capability for quick setup and processing of secured connections

Table 2: The IA-64 architecture addresses core IT requirements

sharing of information, to create business value at the edge of the enterprise—while providing the reliable computational capacity to keep the business impeccably responsive under all conditions.

Intel is uniquely positioned to provide this capacity, reliability and choice through its groundbreaking IA-64 architecture, its many systems alliances and the strong support that it enjoys from the software industry.

Intel is propelling this inflection point in the computer industry with IA-64 to capitalize on the many opportunities presented by the changing Internet dynamics and to provide a wide array of possibilities for business. It truly positions IA-64 to become the Engine Inside the Internet Economy.

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® IA-64 Architecture may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725 or by visiting Intel's web site at <http://www.intel.com>.

Copyright © Intel Corporation 1999. All rights reserved.
*Other names and brands are the property of the respective owners.

