

FAME

Flexible Architecture
for Multiple Environments

Statement
of Direction



- *Bull architecture*
- *for large open servers*
- *in a networked world*
-
-

Table of contents

Executive Summary	4
New challenges and Bull strategy	5
<i>New enterprise challenges and information technology</i>	5
<i>Bull's strategy for NovaScale, its new generation of large servers</i>	6
<i>Bull and Intel</i>	7
Bull's target markets	8
<i>Large business applications on open systems</i>	8
<i>GCOS 7 and GCOS 8 applications.....</i>	8
<i>Scientific and technical applications.....</i>	8
The FAME architecture	10
<i>Bull's innovation and architecture principles</i>	10
<i>Bull's added value has been applied in five critical areas.....</i>	10
<i>Main technical characteristics of FAME.....</i>	12
<i>FSS: FAME Scalability Switch</i>	12
<i>I/O subsystem.....</i>	12
<i>Storage Area Network infrastructure.....</i>	13

Executive Summary

Bull has initiated the development of a new generation of high-end open servers that are designed to meet the new strategic needs of companies and public organizations. These challenges result from an evolution of their business processes, impacted by customer and citizen requirements in terms of reactivity and personalization. This evolution leads to entrust information systems with new missions in a context marked by an intensive use of communications and the Internet.

Bull is leveraging all its experience and know-how developed over the years in large strategic servers (GCOS 7, GCOS 8 and Unix), and at the same time applying its ability to anticipate and use the most forward-looking industry technological innovations. In addition, this new generation of servers, marketed under the NovaScale™ brand name, exhibits increased capabilities for high performance computing that match the exacting scientific and technical computing requirements.

To design and develop the NovaScale server range, Bull launched in 1998 a strategic program, which has led to the development of a new system architecture named FAME (**F**lexible **A**rchitecture for **M**ultiple **E**nvironments). This architecture is based on a new family of 64-bit processors: the Itanium® Processor Family (IPF), within the framework of a very close cooperation with Intel. More precisely, the first FAME-based systems are built on the second generation of this family of processors. Bull thus benefits from the advanced technology brought by the Itanium® 2 processor and its successors in order to position itself in the long term, with its range of large open servers.

Main characteristics of FAME

- **An architecture for large multi-processor systems based on optimal use of market standard components.**
The FAME architecture, based on standard components supplied in high volume by Intel, introduces a breakthrough by significantly reducing the price/performance ratio. On this foundation, Bull has developed and integrated a number of high performance features in order to design SMP type (*Symmetrical Multi-*

Processing) systems, which can include from 8 to 32 processors.

This use of standard components equally guarantees access to a huge range of applications.

- **Highly available architecture.**

Since they are integrated into the enterprise's most strategic environment, these servers must be able to function continuously without any interruption of service. In the FAME architecture, this capability is guaranteed not only by component redundancy, but also by the possibility to isolate and repair a defective element without interrupting system operations.

- **Scalable architecture for multiple environments.**

Thanks to its dynamic partitioning features, the FAME architecture allows the consolidation on the same server, of several applications running under different operating systems, such as, Windows, Linux, GCOS 7 and GCOS 8. In the Internet world, characterized by random peaks, this capability allows the system to dynamically adapt resources to the load requirement without having to be stopped.

These FAME-based servers are designed to run the most demanding transactional and decisional applications: e-business transactions, datawarehouse, datamining, ERP, as well as satisfy the need to consolidate several applications on one system.

Furthermore, applications running under GCOS environments will run either autonomously or simultaneously with Linux or Windows applications on the same system.

In addition, drawing on the exceptional floating-point performance of the Itanium® family of processors, these servers will also serve the scientific and technical applications markets: climatic environment modeling, molecular modeling for discovering new drugs, industrial objects design, etc.

New challenges and Bull strategy

New enterprise challenges and information technology

In the business world, IT infrastructures are currently facing two growing and complementary challenges:

- The need, for enterprises and public organizations, to put new functions into operation in order to renovate and grow the business: getting closer to customers, citizens and partners, optimizing and integrating processes, increasing flexibility for rapid changes.
- The wealth of technologies linked to communications and Internet, or the accelerated digitization of different sorts of information (image, voice, and multimedia), which fulfil the above requirements.

In the technical and scientific world, progress is increasingly linked to the capability to carry out simulations, or analysis or visualization of very large amounts of data, in ever broader application domains: better understanding of living processes and environment, new materials, new design and production processes, etc.

These evolutions are very demanding for large servers. Responding requires that they:

- Absorb important and unpredictable peak transactional processing loads.
- Have sufficient power to extract the relevant information from large amounts of data that have themselves become much more complex due to their increasingly multimedia or semi-structured nature.
- Provide top computing power to carry out very large-scale scientific computations.
- Support very large directories recording the information related to the hundred of thousands end-users who may access on-line services through the Internet.
- Operate round-the-clock, 24 hours a day.
- Integrate into totally heterogeneous environments.

An optimized combination of the following characteristics answers these challenges:

- Guarantee over the long term the use of processors that are always the most powerful on the market, both for scientific and business applications.
- Use a large number of processors running in parallel on an architecture optimized to minimize access times to memory.
- Provide very large central memories, allowing the processing of large volumes of data.
- Have the fastest I/O throughput.
- Integrate specific, redundant hardware features that make it possible to automatically isolate components that may fail or go down, also making it possible to repair them without stopping the system.
- Be able to integrate themselves within wider computing complexes, via very high-speed networks, either locally in clusters, or remotely over large distances, in grid computing over the Internet.

Further economic imperatives such as improving TCO (Total Cost of Ownership) must be included along with the above operational features:

- Excellent price/performance ratio.
- The capacity to dynamically adjust the hardware resources dedicated to each application working simultaneously on the same system, according to their load. These applications can be running on different operating systems.
- An easy-to-use administration system to monitor all the operations of the platform and to optimize use of the resources dedicated to each task.
- The availability of a very large application catalogue.

Bull's strategy for NovaScale, its new generation of large servers

Capitalizing on its recognized experience in large enterprise systems, Bull has decided to design a new line of servers capable of satisfying the needs outlined on the previous page, while anticipating their future evolution.

Thus, the FAME architecture builds upon the expertise accumulated by Bull over several decades, such as:

- Large SMP servers and clusters architectures, both in open systems and large mainframes.
- Processor technology and VLSI design.
- Highly available infrastructures.
- System administration software.
- Secure systems.
- Heterogeneous networks.

The FAME architecture, at the heart of Bull's vision for tomorrow's large servers, is based on:

- Using the advanced technology of the Itanium® 2 processor and its successors, within the framework of a close cooperation with Intel. Thanks to this choice, Bull strongly positions itself in the long term on the market.
- The systematic use of standard boards and circuits developed for volume distribution makes it possible to improve the price/performance ratio, while attracting the availability of a large number of applications, thanks to standard operating systems.
- The support for four types of environments that can simultaneously co-exist on the same server: Microsoft® Windows Server 2003, Linux® and Bull mainframe operating systems, GCOS 7 and GCOS 8.
- Very large memory capacity, from 8 to 256 GB, and beyond.
- The dynamic partitioning of hardware, progressively introduced in servers based on the FAME architecture.
- Diagonal power scaling up to 32 processors, in the full-term, in SMP mode together with clustering configuration capabilities.
- Highly robust: clustering and high availability technology, dynamic reconfiguration of system and sub-system components.

- The I/O sub-system that has an embedded storage area network infrastructure (SAN).

The systems based on the FAME architecture complement the current Bull offer: IA32-based and Unix servers, while opening up new opportunities in the scientific and technical markets.

Bull and Intel

The choice of Intel's Itanium® Processor Family for the FAME architecture

Designed to equip powerful servers, the Intel Itanium® Processor Family is based on a 64-bit architecture, which applies very innovative concepts for executing programs. In order to achieve very high performance, the founding concept closely allies the compilers and the silicon. It has the specific benefit of not relying only on frequency increase as a factor for improving performance. The evolutions planned in the coming years for the Itanium® architecture thus provide headroom for increased performance.

Furthermore, beyond the business applications domain, this new family of processors delivers high performance computing, thus allowing to target equally the scientific and technical markets.

For these reasons and taking into account Intel's technology leadership role on the worldwide market as well as its capabilities to invest over the long term to continuously improve the performance and features of the Itanium® processors, Bull selected this family to equip its NovaScale server range.

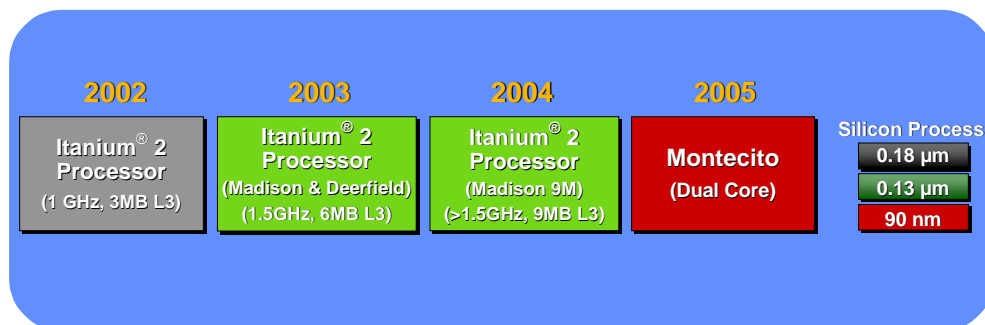
Bull and Intel: a powerful cooperation

Bull and Intel value each other's technologies and share the same vision of the market; therefore they established a very close partnership right from the start of the FAME architecture.

This cooperation ensures:

- For Bull, advanced access to Intel developments, use some of their components and design validation in line with Intel specifications.
- For Intel, to test their components within multiprocessor architecture and to integrate into their roadmaps the different needs related to the world of large servers.

Throughout this cooperation, there are regular exchanges of engineers between the laboratories of the two companies and some Bull prototypes are available for tests at Intel labs.



Intel® Itanium® Processor Family Roadmap

Bull's target markets

Servers based on the FAME architecture target three main markets in the coming years:

- Large business and enterprise applications on open systems.
- Applications in the GCOS 7 and GCOS 8 environments.
- Scientific and technical applications.

Large business applications on open systems

This market is a major target for Bull, because of its in-depth understanding of the needs of large organizations in both the private and the public sectors. It includes:

- Business Intelligence, in particular datawarehousing and datamining.
- Large enterprise applications: ERP, CRM, SCM, etc.
- Large database servers for Internet transactions.
- Consolidation of servers disseminated across the enterprise that contributes to IT systems optimization.
- Large business sector applications, such as billing for operator's online services, on-line reservations and banking transactions.

To implement the corresponding applications, the servers based on FAME architecture, support open operating systems such Microsoft® Windows Server 2003 and Linux®. Bull's expertise in Windows enterprise solutions (Windows Server 2003 Datacenter Edition, high availability, fault tolerance, centralized server architecture with thin clients), makes it possible to offer powerful solutions that use 64-bit Microsoft SQL Server and the .NET framework.

Bull also leverages its involvement in Open Source initiatives, such as the Atlas project (Linux) or the ObjectWeb Consortium J2EE architecture, to offer top-range open source infrastructures that are reliable and easy to manage.

This offer will be progressively enriched with market applications as they become available on Intel® Itanium® 2 architecture.

GCOS 7 and GCOS 8 applications

Bull provides its GCOS 7 and GCOS 8 customers with a long-term roadmap and a clear evolution path, which will enable them to take advantage of FAME architecture without having to change or recompile their applications. Investments and specific developments have also been carried out to ensure binary compatibility on the Intel platforms of the GCOS 7 and GCOS 8 operating systems.

In May 2001, the Group had already launched Bull DPS7000 XTA, the result of the Diane project for GCOS 7 systems. The availability of GCOS 7 on Intel 32-bit platforms guarantees total binary compatibility for applications. Many customers have chosen this solution and have deployed it easily and successfully in their production environments, where it was transparent for users and provided significant performance gains.

GCOS 7 64-bit prototypes on Intel® Itanium® 2 architecture have already been demonstrated.

The Helios project has an approach similar to Diane, and is in development and testing for the GCOS 8 environment in Bull's laboratories in Phoenix, USA, and in Les Clayes, near Paris, France.

Scientific and technical applications

The FAME architecture developed by Bull based on the Intel® Itanium® 2 processors (and later on its successors) is particularly suited for applications that need high performance computing power. Some examples include defense, energy, scientific research, automobile industry, seismology, aeronautics, etc.

The Itanium® 2 processor makes it possible to smooth the differences between the scientific and business IT architectures. The 64-bit chips integrate floating-point processors with vectorial instructions. Furthermore, they benefit from important improvements in memory speed.

The shared memory SMP architecture constitutes the most efficient solution for multi-threaded applications. Nevertheless, the most demanding applications in terms of computing power exceed

the capacities of the current largest available SMPs. This has led the industry to combine the shared memory multiprocessor-programming model with message passing programming. It is therefore planned that for these types of applications, servers based on FAME architecture will be interconnected by a low latency and wide bandwidth link, to provide high-performance clusters to meet the most demanding power requirements.

In this framework, Bull has signed several partnerships in these domains, in order to offer complete solutions for enterprises and public organizations.

The FAME architecture

Bull's innovation and architecture principles

FAME is based on a NUMA (*Non Uniform Memory Access*) architecture optimized by Bull, which makes it possible to build a large SMP system using QBB (Quad Brick Block) Intel quadriprocessor boards with the Itanium® 2 Processor Family and memory.

In addition to the QBB modules, the building-up of the system also rests on the modular input/output units IOB (Input Output Boxes).

This approach by component allows to expand the modularity even further by combining processors, memory and I/O. This allows either the construction of elementary systems being able to be used as entry-level, or the construction of a federative architecture of a top-of-the-range system, by replication and interconnection.

This construction is based on the implementation by Bull of very high-speed interconnections and on the development of a very sophisticated chip: the FSS (FAME Scalability Switch), which ensures that each processor, has access to the I/O and has a coherent vision of the global memory that can reach 256 GB. This memory is formed by uniting the memories associated with each quadriprocessor board.

Bull's added value has been applied in five critical areas

- **The interconnection of components at very high speeds and low latency.** In this context Bull has developed the FSS at the heart of the FAME architecture. Using 0.18 micron CMOS technology with copper interconnection, the FSS, a square chip with 18-mm sides, includes 60 million transistors and 1520 input/output pins. These characteristics make it one of the most advanced chips.

The FSS ensures information coherence across the whole system, whether the information is in memory, or copied in the cache of the processors. The FSS thus provides a large symmetrical multiprocessor system, with a global addressing space.

The FSS also optimizes the traffic between all the processors and synchronizes all the

internal communications of the server. Thanks to the strong optimization of the circuits and the protocols used, the memory access times are practically uniform, varying only by a factor of one to three, whereas on the first NUMA architectures, this factor varied from one to 15.


- **The built-in Platform Administration and Maintenance (PAM)** software suite is also a significant Bull's added value. The PAM manages the system including the operation of the partitioning. It plays an important role in guaranteeing the availability of all modules.

The PAM is also a proactive administration tool since all the pertinent events are not only recorded but also can generate automatic corrective actions and calls to the support centers. In particular, it will enable, by using the FSS, efficient management of the dynamic partitioning of the system (processors, memories, input/output), as well as the storage infrastructure.

- **Dynamic partitioning.** A system based on the FAME architecture has the capability to be partitioned into several independent domains (a maximum of four). A domain is a coherent set of resources managed by one of the operating systems. This arrangement ensures true physical isolation of the different subsystems, which makes it possible to have different environments on the same system, without one disturbing the other, particularly in the event of a failure of one of them. There are no conflicts to manage, nor resources to share; this strengthens flexibility, performance and operational security.

These features also provide the "power on demand" function, which makes it possible to meet unpredictable load peaks, such as with an Internet application, by adding, on demand, supplementary processors.

- **The high availability** has been carefully fine-tuned. It associates:
 - **Data integrity:** all the data paths are protected from end-to-end by error-detecting and auto-correcting codes.
 - **Redundancy** of all system components, so there is no single point of failure.
 - **Hot swappable units** of all components without stopping the system: these include logistic components (power supplies, ventilators), peripheral access boards, and even QBB boards including CPUs and memory, together with the operating systems that allow this.
- **Packaging density.** Particular attention has been applied to packaging, in order to provide exceptional global performance of the system while offering easy maintenance. This is a real technical feat, especially by housing in one third of a cubic meter, a module of four QBB boards, that is 16 processors with their associated memory, as well as two I/O boards offering 22 PCI-X slots, as well as all the logistics components (ventilators, power supply).



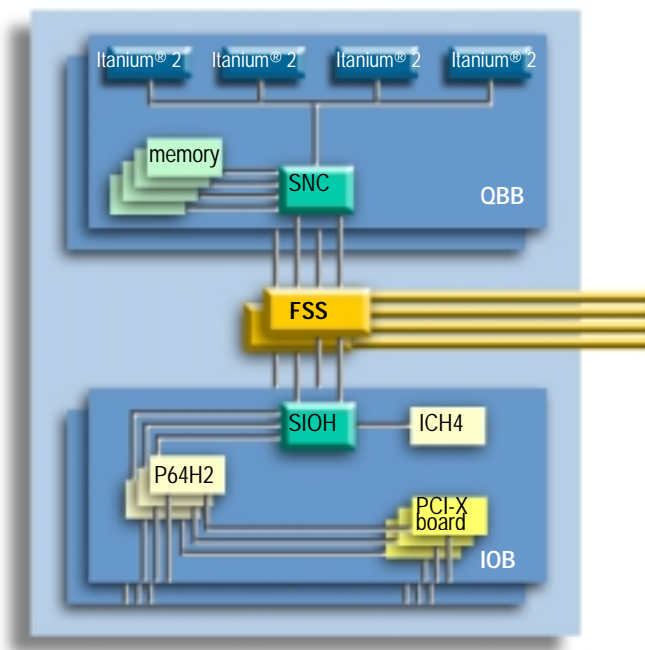
FAME architecture with 32 processors

Main technical characteristics of FAME

FSS: FAME Scalability Switch

The key piece of the FAME architecture is therefore the FSS interconnect processor, designed by Bull to support the cumulative performance of up to 8 QBB quadriprocessor boards. A 32-processor system consists of modules, each of which has two FSS for fault tolerance and bandwidth, up to four QBB and two IOB. Interconnection between the modules is achieved by making the FSS of each module communicate via four XSP (eXtended Scalability Port) communication channels, with total throughput over 25 GB/s using a frequency of 2.5 GHz.

The memory bandwidth is provided by aggregating the SNC memory controllers (Scalability Node Controller), situated on each quadriprocessor block. Each SNC has a memory bandwidth of 6.4 GB/s and can deliver up to 25 GB/s on all its interfaces. In its maximum configuration, the FAME architecture provides a memory bandwidth in excess of 50 GB/s.



FAME Architecture

Special attention has been paid to optimizing memory access latencies, by using an efficient filter for cache coherence traffic. As a result, the inter-QBB flow is significantly reduced. The “NUMA factor”, generally the weak point in systems that are constructed using “building blocks” is also significantly improved. Thus, by achieving a NUMA factor of 1:2:3 (respectively for memory access to the same QBB, to another QBB on the same module, and finally to a QBB on another module), the FAME architecture offers a high performance level, even with software that is not specifically tuned for NUMA.

I/O subsystem

The I/O subsystem of FAME has been designed to optimize access to data and to the network. Using standard market components and powerful switching features, it will also be able to easily integrate emerging technologies.

The system has been sized to deliver an I/O performance that matches its computing power:

- 1/4 of the bandwidth switching of the FSS is dedicated to the I/O.
- Fulfilling its mission as an architecture for large enterprise systems, FAME implements “fat pipes” with a number of fast PCI-X buses (8 per IOB) that enable the high-speed controllers (e.g. FCS/Fiber Channel Standard at 2 Gb/s) work without interference and with improved error containment.
- Using a SAN (Storage Area Network) infrastructure, usage of multiple access paths guarantees efficient load balancing of the I/O subsystem.
- Furthermore, this I/O subsystem provides both excellent scalability and performance: 6 GB/s of peak I/O bandwidth, 2 GB/s of sustained throughput, and 250 to 300,000 I/O operations per second.

Storage Area Network infrastructure

All the peripheral and the communication devices are connected to the IOB via the PCI-X boards.

The FAME architecture includes a SAN network that can be shared between several servers.

The centralized administration ensures coherent reconfiguration of the domains and of their storage access paths. Thus, when a set of resources (for example: two QBB and one IOB) are moved from one domain to another, it is necessary to carry out synchronous reconfiguration of the interconnection network, by using the FSS and the Fiber Channel network that access the storage units. This operation is performed without any manual intervention on the system. Moreover, a single console can manage the domains of several servers based on the FAME architecture; the resulting configuration is then called an “extended FAME”.

© Bull S.A. March 2003 - Bull acknowledges the rights of proprietors of trademarks mentioned herein. Bull reserves the right to modify this document at any time without notice. Some offers or parts of offers described in this document may not be available in your country. Please consult your local Bull correspondent for information regarding the offers which may be available in your country. Intel and Itanium are registered trademarks of Intel Corp. or its subsidiaries in the United States and other countries NovaScale is a trademark of Bull S.A.

Bull SA - 68 route de Versailles - 78430 Louveciennes - France
UK : Great West Road, Brentford, Middlesex TW8 9QH
USA : 300 Concord Road, Billerica, MA 01821
www.bull.com

