

White Paper

# A Smarter Approach to Multi-Core:

## Freescale's Next-Generation Communications Platform

---

## Overview

---

Multi-core processing is establishing itself as the solution of choice in a number of electronic industries, yet the technology itself is still in its infancy. As Freescale moves forward with our strategy for exploiting the outstanding capabilities of the PowerQUICC® family of communications processors, we are developing multi-core platforms to be more efficient, flexible and easier to integrate into networking applications than any other solution on the market.

Networks are moving into a new era. No longer is it sufficient to just reach out and touch someone. As network consolidation and convergence have driven costs down, they have also enabled the advent of true online collaboration. This will be an era of networking where every connection matters: friend-to-friend, supplier-to-consumer, peer-to-peer. Just making the connection anywhere, anytime and on any device is still paramount, but the experience of the interaction is the new focus. The user experience will become one of immediate content delivery with context, assured security and privacy, and bundled, expanding services. How can we deliver this seamless interactive experience that grows with user preferences, behaviors and capabilities?

Several homogeneous multi-core solutions are already seeding the marketplace attempting to add processing power to the problem, but they also come with growing pains. Multi-core systems are only as effective as the software's ability to take advantage of parallelism, and the pure processing potential of multi-core platforms today is not yet tapped. A related challenge is that the majority of the networking installed base is still operating on a mixture of single-core processors, ASICs and DSPs. Developers have the enormous task of effectively migrating millions of tested, proven and fielded lines of code to multi-core architectures before the full benefits of any multi-core system-on-chip (SoC) solution can be realized.

Freescale's next-generation Multi-core Communications Platform is an innovative evolution of our established PowerQUICC architecture designed to enable a new era of networking where the reliability, security and quality of service for every connection matters. This white paper introduces you to the technical architecture of this platform and its comprehensive approach to solving the programmability problem.

# Contents

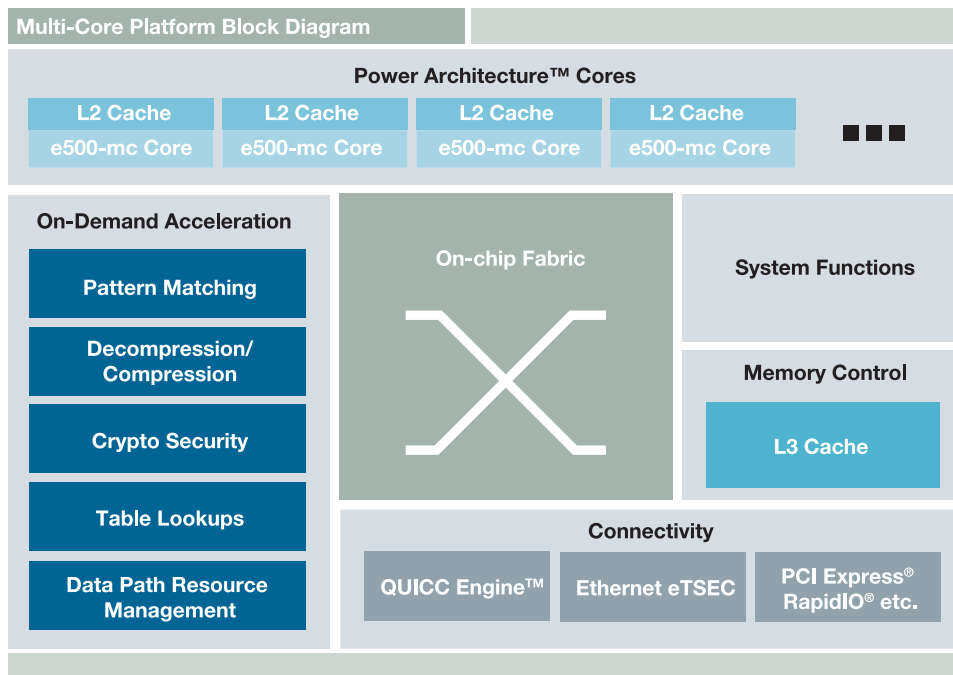
---

The Freescale Approach to Next-Generation Multi-Core .....	4
Scalable On-Chip Fabric.....	4
Enhanced Power Architecture™ Core .....	5
“Best of Both Worlds” Cache Hierarchy.....	5
On-Demand Application Acceleration .....	6
A Few Words about Data Path Resource Management.....	6
Unlocking Multi-Core’s Potential.....	7
Resource Partitioning and Virtualization .....	8
Debugging and Performance Monitoring .....	8
Hybrid Simulation Environment .....	9
45 Nanometer (nm) Process Technology Edge.....	10
Heritage of Multiprocessing Innovation .....	10
Summary: Engineering the Balanced Multi-Core Platform.....	11

# The Freescale Approach to Next-Generation Multi-Core

Freescale’s multi-core communications platform represents a balanced approach to multi-core system-on-chip (SoC) design. It introduces an advanced on-chip connectivity fabric that is able to support up to 32 cores and beyond. The platform implements enhanced cores built on Power Architecture™ technology, each with private Level 2 (L2) cache, also known as backside cache. In addition, the platform extends our proven on-demand application acceleration capabilities with data path resource management functions. The multi-core platform is targeted to be implemented in 45 nm silicon-on-insulator (SOI) process technology with a development path to 32 nm SOI technology.

While the multi-core platform is designed with aggressive performance targets, ease of use has also figured prominently in our platform definition. One of the significant obstacles in multi-core implementations today is programming efficiency and debugging. Freescale’s industry-leading ecosystem of tools and operating system vendors is mobilized to simplify multi-core development and debugging. Especially noteworthy is our collaboration with Virtutech to deliver a new hybrid simulation environment that allows programmers to migrate application software and evaluate performance in advance of prototype hardware.



## Scalable On-Chip Fabric

The multi-core platform will employ our highly scalable and modular on-chip fabric, the result of multi-year research and development, which enables cache-coherent, concurrent, low-latency connectivity among cores. Unlike a shared bus as interconnecting medium among cores, memory and peripherals, the on-chip fabric helps to reduce the bus arbitration and contention issues that other multi-core architectures face as more traffic is introduced into the system. It behaves like a mesh, allowing concurrent traffic to enter and exit the system from any point within the fabric rather than through a single point.

Inherently scalable, the fabric is designed to sustain multiple, fully-coherent transactions every cycle and easily expand to accommodate more cores. Our fabric also supports the option for heterogeneous clustering, allowing Freescale’s full portfolio of Power Architecture cores, which spans a wide range of power and performance design points, to be mixed and matched in a product with full coherency among the cores.

## Enhanced Power Architecture Core

An enhanced Power Architecture core, based on the familiar e500 core, is targeted for the first multi-core platform implementation. The e500-mc core's frequency in a wide multi-core product will be targeted to 1.5 GHz. This allows the platform to retain the PowerQUICC platform's industry-leading power-to-performance characteristics with the highest instruction-per-cycle (IPC) and highest frequency for a given watt per area. The e500-mc cores are also designed to offload repetitive and computing intensive operations to high-performance acceleration blocks, increasing the number of processing cycles for higher throughput or new services and applications.

Each e500-mc core in the platform will have its own L2 backside cache. Backside cache is connected to the CPU through a direct channel, enabling extremely high application performance. It allows the cache to match the half speed of the CPU, resulting in latency improvements well over 50 percent of "shared-bus/shared-cache" architectures. L2 backside cache also enables tuning the contents of the cache between instruction and data, according to different application needs, easing partitioning and improving performance by drastically reducing CPU stalls. In addition, the L2 backside cache reduces traffic on the on-chip fabric and main memory, which reduces latencies and improves bandwidth for other users of the fabric and system memory.

Because Power Architecture technology is the leading processor architecture for the networking and communications market, by continuing to leverage these high-performance cores in our multi-core platform, Freescale enables a smoother migration path from existing single-core implementations to dual-core and multi-core platforms.

## "Best of Both Worlds" Cache Hierarchy

Recognizing the limitations of existing processors that rely on a shared cache model, we took a new approach by incorporating a three-tiered cache hierarchy into the multi-core platform. Level 1 (L1) cache is retained on the core. As previously mentioned, L2 cache is attached to the cores as a backside implementation that can significantly improve performance. However, there are some tasks for which a shared cache is desirable, such as inter-processor communication and operating on shared data structures. For those instances, we are also providing a multi-megabyte Level 3 (L3) cache. This high-bandwidth, shared cache maximizes hit-rates while providing fast memory access for input/output (I/O) and accelerator blocks. For example, I/O blocks can use the L3 cache to provide the same header allocation capabilities offered by the front-side L2 cache in current PowerQUICC III devices.

The on-chip fabric works in concert with the caching hierarchy to enable cache-coherent and concurrent accesses. The innovative backside cache implementation combined with the fabric is designed to enable data replication, modified intervention and full hardware coherence tracking.

## On-Demand Application Acceleration

Freescale's proven, flexible networking and application performance acceleration technologies take the multi-core architecture to a new level. We have been delivering application acceleration in current devices, such as our MPC8572E PowerQUICC III processor, and we continue this shared-resource approach with our next-generation multi-core platform. On-demand application acceleration offers performance advantages over pure core processing cycles, enables lower power implementations and reduces silicon area thus reducing cost. Freescale's on-demand, high-performance acceleration technologies include:

- Pattern matching for deep packet inspection and full content processing
- Decompression/compression to unpack data for inspection and pack it for delivery
- Crypto security for confidentiality, integrity and authentication
- Table lookups for packet parsing and flow classification
- Data path resource management to efficiently allocate on-chip resources

### A Few Words about Data Path Resource Management

The multi-core platform introduces an advanced on-demand acceleration function, data path resource management. This function handles intra-chip message passing to support the most flexible and optimal use of the available resources. This resource management function supports:

- Passing data between cores, hardware accelerators and network interfaces
- Intelligent load spreading of packets across pools of resources including cores and hardware accelerators based on resource utilization and flow classification
- Virtualization of shared resources such as hardware accelerators and network interfaces
- Low overhead, low latency software reception and data transmission
- Mechanisms to ensure packet ordering and sequencing

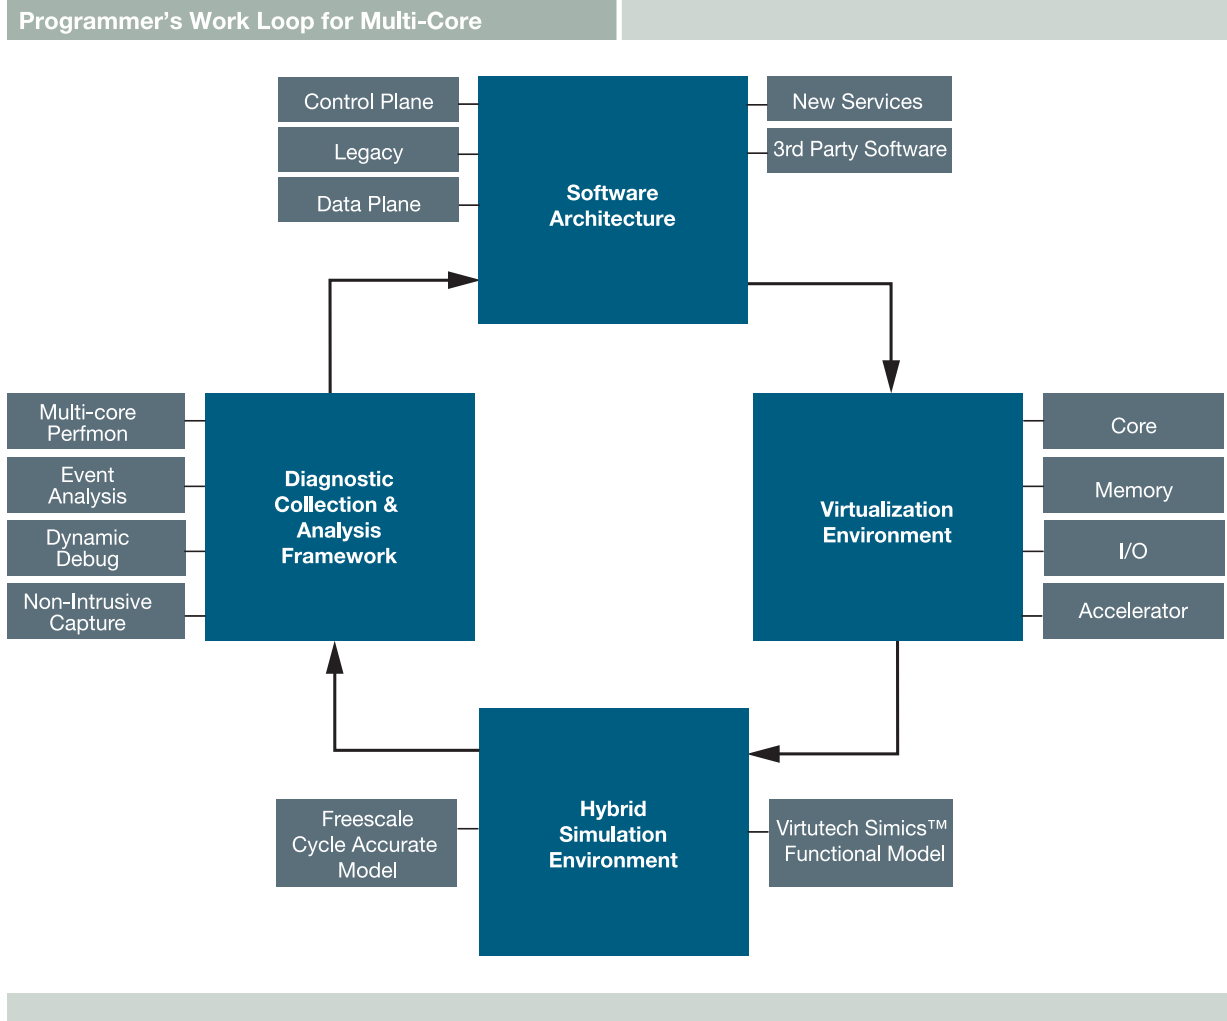
To enable the highest performance possible, the multi-core platform provides hardware support for memory buffer allocation by resources. This helps to significantly reduce the software overhead to manage the buffers for multiple, high-performance hardware accelerators and network interfaces.

# Unlocking Multi-Core's Potential

Multi-core platforms require software engineers to spend significantly more time thinking about software architecture. Exploiting the performance potential of multi-core processors means embracing parallel processing, which can be a challenge given the long and successful history of single core systems that are largely self synchronizing. Networking applications offer coarse grained parallelism in the form of packet processing, and the interactions between a networking data path and the control plane are sufficiently decoupled to create an additional level of parallelism.

While this immediate parallelism is easy to envision, things get interesting when the performance requirements of a data path flow exceed a single CPU's capabilities, or when a single core can't provide sufficient control plane responsiveness. Load balancing and mixed asymmetric/symmetric multi-processing environments on the same device are challenges that Freescale's multi-core platform is designed to address.

While software architects are thinking about distribution of tasks, the processing densities offered by multi-core platforms will cause hardware architects to think about consolidation and re-partitioning functions that have been distributed across discrete CPUs or modules. These decisions will interact strongly with the introduction of new services and capabilities in the system. For both software and hardware architects, there is a need for a great deal of flexibility in a multi-core processor and for good mechanisms to help facilitate experimentation with future architectures.



## Resource Partitioning and Virtualization

Consolidation of discrete CPUs into a single multi-core SoC and potential repartitioning of legacy software on those cores introduces many opportunities for unintended resource contentions to arise. Hardware partitioning mechanisms that are both robust and intuitive to the user accelerate code migration and system verification and improve system performance by relying on more efficient synchronization and locking methods.

The multi-core platform provides hardware mechanisms that help to ensure cores only access the resources (memory, peripherals, etc.) that they are designated to access. For example, some resources may be dedicated to a particular core, or set of cores, and access is restricted to only those cores. It also provides an I/O memory management unit (MMU) to enforce access controls on bus mastering peripherals.

In addition, the platform's data path accelerators inherently virtualize access to network interfaces by decoupling them from the cores. Any core can send or receive packets on any network interface, or any accelerator, without the need for inter-core synchronization in software.

A hypervisor is a trusted layer of software that manages resources and access to them. It dedicates resources to cores, as appropriate, when the system is initialized and manages access to shared resources during system runtime. The multi-core platform permits high-speed peripherals to be dedicated to cores (or sets of cores running a symmetric multi-processing operating system) with minimal changes to legacy software. Run-time hypervisor software can then enable shared access to resources that are not dedicated.

## Debugging and Performance Monitoring

While perhaps obvious, it's important not to underestimate the challenge of debugging multi-core processors. Leveraging the additional capacity of the extra cores requires parts of the application to execute in parallel, compelling synchronization between tasks to protect shared resources against concurrent accesses. Many problems related to concurrency, such as race conditions, tend to be sensitive to timing and workload and become more pronounced in multi-core platforms. Previously existing concurrency problems on single-core processors may be more easily triggered when migrating code to multi-core processors due to the interactions between tasks running on different cores. Diagnosing such problems adds the complexity of simultaneously tracing these interactions between the tasks as they access platform resources.

Debugging a SoC is not limited to the cores. SoCs are complex devices with multiple hardware accelerator blocks, memory, caches and network peripherals. As such, debugging solutions should strive to provide greater system visibility to all of these elements. Increased integration of cores and SoC peripherals itself reduces the ability to use traditional debug tools, such as logic analyzers, to view interactions that no longer cross externally-visible interfaces. Freescale is engineering the necessary hooks in our multi-core platform to enable advanced debugging of the platform, and we are also working in tandem with industry leading vendors to assure the availability of tools that can take advantage of these features. These hooks include, among others, trace capabilities on the cores, watchpoint triggers, cross triggering capabilities, performance monitoring capabilities on the cores and SoC and the Power ISA defined debug features. Freescale also plays a leading role in the development of industry debug standards.

Obtaining the highest levels of performance requires detailed, real-time information about how the silicon is performing. This need is even more pronounced on multi-core devices. In order to provide no-overhead cycle-granularity performance information, Freescale's multi-core platform provides several different features. Each core has a private set of performance monitors allowing it to examine core events including branch mispredictions, instruction mix, L1 and backside L2 cache accesses, MMU misses and interrupt latencies. The platform provides a shared performance monitor set responsible for analyzing system-wide events including DDR page misses, bus utilization, RMON statistics and transaction counts. These types of profiling and performance tuning features will ensure appropriate code efficiencies for our multi-core platform.



## Hybrid Simulation Environment

Freescale is collaborating with Virtutech to develop a full system simulation model of our multi-core platform. Specifically, the companies are developing a hybrid simulation environment that combines Freescale's proven cycle-accurate modeling technology with Virtutech's established functional modeling technology, called Simics™ simulator, that enables ease of software development, performance prediction and optimization of customer applications for the Freescale multi-core platform.

Using the hybrid simulation environment, which allows easy switching between functional and cycle-accurate models, developers can migrate and partition operating systems, middleware and applications onto the virtualized multi-core platform for development, debugging and benchmarking—even prior to silicon availability. The environment also enables safe and easy experimentation with partitioning, parallelizing and optimizing systems and applications. Software developers can perform “what if” scenarios and tune the performance for specific situations without real-world hardware constraints.

The hybrid simulator provides a programmer's view of the hardware, and features:

- A fast, functional model of the Freescale's multi-core platform
- A detailed cycle-accurate model of Freescale's multi-core platform
- A comprehensive package with infrastructure and tools for software development, code partitioning and debugging, profiling and visualization
- Visibility into system state both architectural and microarchitectural including caches, registers and pipelines
- Run-time control of execution software including breakpointing, stepping and reverse execution
- Ability to boot multiple operating systems

A major advantage of a hybrid simulator is its ability to dynamically switch back and forth from a high-speed functional mode to a more detailed cycle-accurate mode. This allows software developers to quickly boot an operating system and execute code at critical points and then switch to the more detailed cycle-accurate mode to analyze specific areas of interest—no more waiting days for results.

As a development platform for multi-core systems, the hybrid simulation environment is designed to enable an extensive amount of flexibility and experimentation in a non-invasive environment—no instrumentation is needed in the operating system or application. Software developers are able to decrease bring-up time for the target system all while improving the overall quality of their code.


A full system simulation model of Freescale's MPC8572E dual-core PowerQUICC III processor is available today for developers and can serve as a lead in to the model of the next-generation multi-core platform.

## 45 Nanometer Process Technology Edge

In January 2007, Freescale began our collaboration in the IBM technology alliance for joint semiconductor research and development. The agreement includes complementary metal-oxide semiconductor (CMOS) and silicon-on-insulator (SOI) technologies as well as advanced semiconductor research and design enablement starting at the 45 nm generation. In addition to leveraging owned capacity in internal fabs and its existing relationships with leading foundry manufacturers, Freescale will have access to the combined manufacturing capacity of IBM's Common Platform™ partners. This includes IBM and Chartered Semiconductor fabs, both with proven capabilities in high-performance SOI wafer processing.

45 nm SOI is a rapidly maturing technology, expected to reach full process certification in 2008. The 45 nm SOI technology includes a range of high-performance transistor offerings and static random access memory (SRAM) bitcells that provide an excellent balance of performance and low power. Some key enablers of this technology are 193 nm immersion lithography for scaling and reducing sources of device variation, porous low-k ( $k=2.4$ ) dielectric for minimized back-end wiring delay, and advanced strain techniques for enhanced transistor performance. Achieving the highest level of process technology performance in 45 nm SOI is a key building block for multi-core success within a usable power envelope.

### 45 nm Edge: 50% Reduced Die and Power from 90 nm



## Heritage of Multiprocessing Innovation

Freescale introduced our first PowerQUICC communication processors in 1995, and since then have been refining and expanding the families with advanced core development and integration techniques to enhance embedded performance within customer power envelopes. The extensive protocol interworking capability of our communications processors along with application accelerators have consistently addressed our customers' application needs. In 2006, we launched the MPC8572E PowerQUICC III dual-core processor with advanced content processing, including a pattern matching function. This is a ground-breaking product for its level of integration: dual Power Architecture e500 cores and flexible application acceleration including crypto security and pattern matching.

### PowerQUICC Family of Communications Processors



We have more than ten years of PowerQUICC development in heterogenous multi-processing products, with specific lessons learned from the dual-core MPC8572E and MPC8641D built on Power Architecture technology and the industry leading quad-core DSP (MPC8144E). This extensive experience, combined with input from various customers and clients, has allowed Freescale to enter a new era of SoC innovation and processing intelligence as we disclose our next-generation communications platform, optimally engineered to support multiple Power Architecture cores.

## Summary: Engineering the Balanced Multi-Core Platform

---

Tomorrow's networking needs can no longer be met by increasing the operating frequencies on single-core architectures. Thermal management challenges are overwhelming the performance improvements achievable by increasing CPU frequency. The answer, however, is not simply adding cores to a die. As many current implementations show, more does not necessarily mean better. There may be contention for bus bandwidth and memories, scalability problems and perhaps even worse, unused processing cycles due to lack of programming visibility.

Freescale's next generation multi-core communications platform is not only designed to provide superior performance and energy efficiency, but also to help make the transition to multi-core processors as quick and as painless as possible with an industry leading enablement ecosystem.

The multi-core platform will change the multi-core landscape with its:

- Scalable on-chip fabric
- Enhanced Power Architecture e500-mc cores
- Three-level cache hierarchy
- Proven on-demand application acceleration
- Advanced 45 nm process technology

We invite you to learn more and share your experiences in an open and interactive manner that Freescale has always fostered with our user community.

## How to Reach Us:

### Home Page:

[www.freescale.com](http://www.freescale.com)

### Power Architecture Information:

[www.freescale.com/powerarchitecture](http://www.freescale.com/powerarchitecture)

### e-mail:

[support@freescale.com](mailto:support@freescale.com)

### USA/Europe or Locations Not Listed:

Freescale Semiconductor  
Technical Information Center, CH370  
1300 N. Alma School Road  
Chandler, Arizona 85224  
1-800-521-6274  
480-768-2130  
[support@freescale.com](mailto:support@freescale.com)

### Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH  
Technical Information Center  
Schatzbogen 7  
81829 Muenchen, Germany  
+44 1296 380 456 (English)  
+46 8 52200080 (English)  
+49 89 92103 559 (German)  
+33 1 69 35 48 48 (French)  
[support@freescale.com](mailto:support@freescale.com)

### Japan:

Freescale Semiconductor Japan Ltd.  
Headquarters  
ARCO Tower 15F  
1-8-1, Shimo-Meguro, Meguro-ku,  
Tokyo 153-0064, Japan  
0120 191014  
+81 3 5437 9125  
[support.japan@freescale.com](mailto:support.japan@freescale.com)

### Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd.  
Technical Information Center  
2 Dai King Street  
Tai Po Industrial Estate,  
Tai Po, N.T., Hong Kong  
+800 2666 8080  
[support.asia@freescale.com](mailto:support.asia@freescale.com)

### For Literature Requests Only:

Freescale Semiconductor  
Literature Distribution Center  
P.O. Box 5405  
Denver, Colorado 80217  
1-800-441-2447  
303-675-2140  
Fax: 303-675-2150  
[LDCForFreescaleSemiconductor@hibbertgroup.com](mailto:LDCForFreescaleSemiconductor@hibbertgroup.com)

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright license granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.



Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org. © Freescale Semiconductor, Inc. 2007  
Document Number: MULTICOREFTFWP  
REV 1

