

Developing Portable HPC Applications using Intel® QuickAssist Technology Accelerator Abstraction Layer Software (AAL)

An Introduction to AAL

Joseph Grecco
Sr. Staff Systems Architect
Digital Enterprise Group
Intel Corporation
Joe.Grecco@intel.com



Objective

- Introduce Intel® QuickAssist Technology Accelerator Abstraction Layer (AAL) Software
 - High level overview
 - What is it solving and how?
- Take away:
 - Is AAL something you would like to learn more about?

Problem Statement

Need for acceleration may be met by specialized acceleration technology, but...

- High cost of entry
 - \$\$ for hardware, development time...
- Non-standard software development environment
 - Technology/vendor specific interfaces
- High total cost of ownership
 - Typically low-level interfaces; device driver, ISA...
- Technology changing rapidly
 - Difficult to adopt new technology



Solution

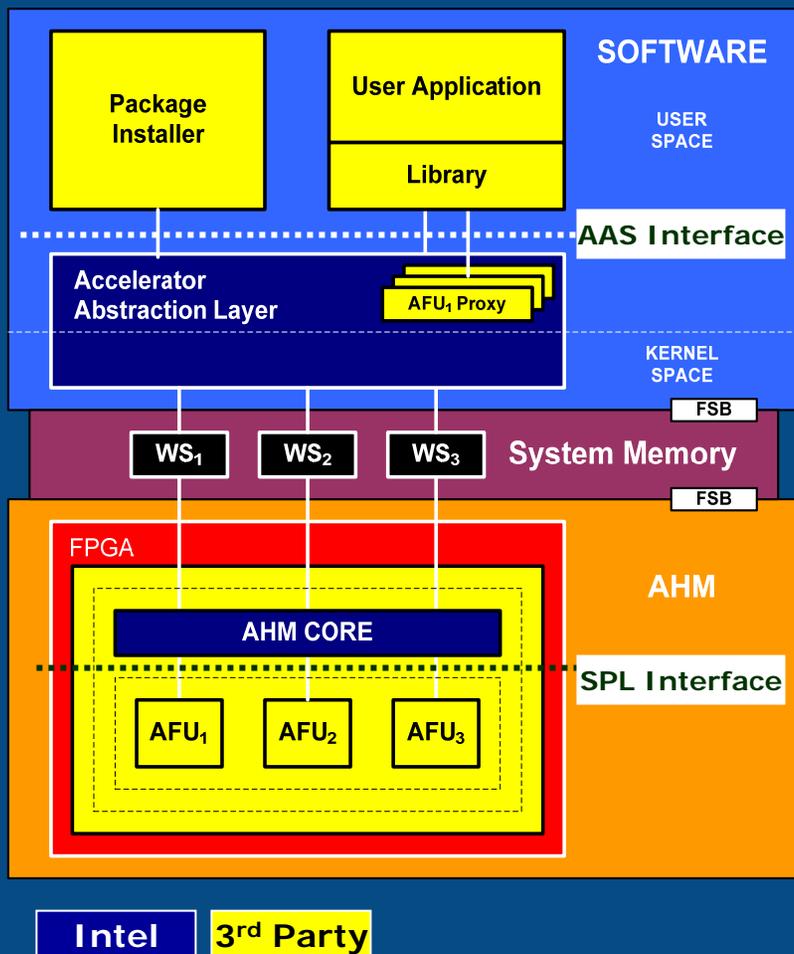
Provide enabling technology that:

- Reduces cost of entry
 - Increase choice of accelerator technologies
 - Enable best of breed
 - Simplify development
- Reduces total cost of ownership
 - Enable efficient use of platform resources
 - Build in reliability and extensibility
- Improves platform agility
 - Enable developers to keep pace with changing technology

**Intel's FPGA accelerator platform with AAL
enables this solution**



FPGA Accelerator Platform Components



- **Accelerator Hardware Modules (AHM)** contain FPGA and local memory
- **Accelerated Function Units (AFU)** represent an FPGA accelerated algorithm
- Intel-supplied **AHM Core** logic implements FSB protocol, low-level AHM device and internal system interfaces (System Protocol Layer - SPL)
- **Workspace** is system memory that is managed by AAL
 - Pinned and mapped into application's virtual address space
 - Data exchanged using messaging and shared workspace
- **Accelerator Abstraction Layer (AAL)** provides platform level services (AAS) for integrating accelerators into Intel platforms

AAS and SPL interfaces ensure portability of application and AFU

Intel® QuickAssist Technology Accelerator Abstraction Layer (AAL) SW

- Low level software layer
 - Uniform set of platform services
- Independent of accelerator technology
 - Tightly coupled accelerators: FSB and Intel® QuickPath Interconnect
 - Loosely coupled accelerators: PCI Express* Technology
- Independent of accelerated workloads
 - Data and Task parallel
 - Very Large to small data sets, streaming data
- Designed to be compatible with existing frameworks



AAL Architecture Overview

- Service Oriented
 - Publish/Request Service Model
 - Client Decoupled from Implementation
 - Location, Quantity, Taxonomy
 - Increased Reuse
 - Resource aggregation and sharing
 - Dynamic reconfiguration/reprovisioning
 - Dynamic Interface Binding
 - Accelerator resources “Virtualized”
 - Reduced compile/link level dependencies

**Service Oriented Architecture provides
Platform Agility**



AAL Architecture Overview

- Object Oriented
 - Interface oriented design
 - Abstraction from implementation
 - Extensibility
 - Higher degree of reuse
 - Simplified development

**Object Oriented Architecture provides
Portability and Simplicity**

AAL Architecture Overview

- Infrastructure Enables Common Services
 - Configuration Management:
 - Resource, Policy, Workspace
 - Communications:
 - Registration, Request, Binding, Transport
- Extensible
 - New services can easily be added
 - No impact on existing services
 - Many services can be replaced or enhanced

AAL Architecture Overview

- Workspace Memory Management
 - Allocation, Pinning, Security
- Resource and Policy Management
 - Enables sharing between Apps/Users
- Resource Management
 - Enables dynamic “runtime” reprovisioning and reallocation
- Policy Management
 - Enables sophisticated policies for resource sharing
 - Open interfaces allow 3rd party to define policies, properties

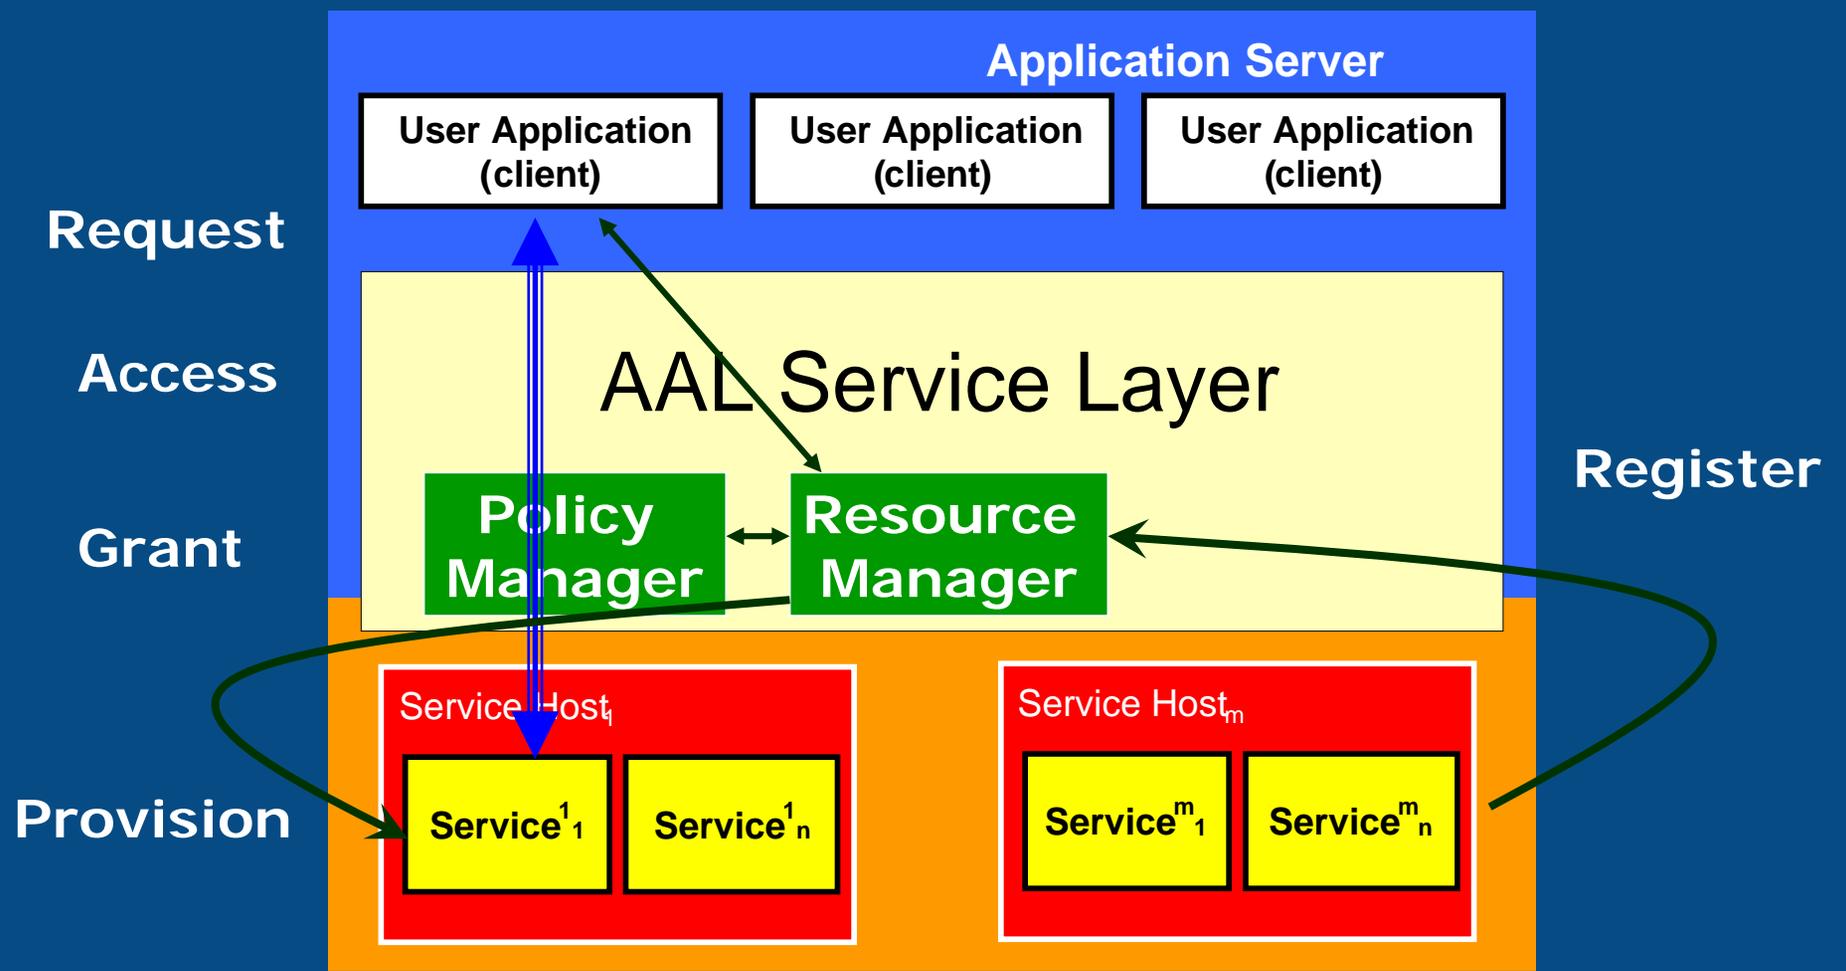


AAL Architecture Overview

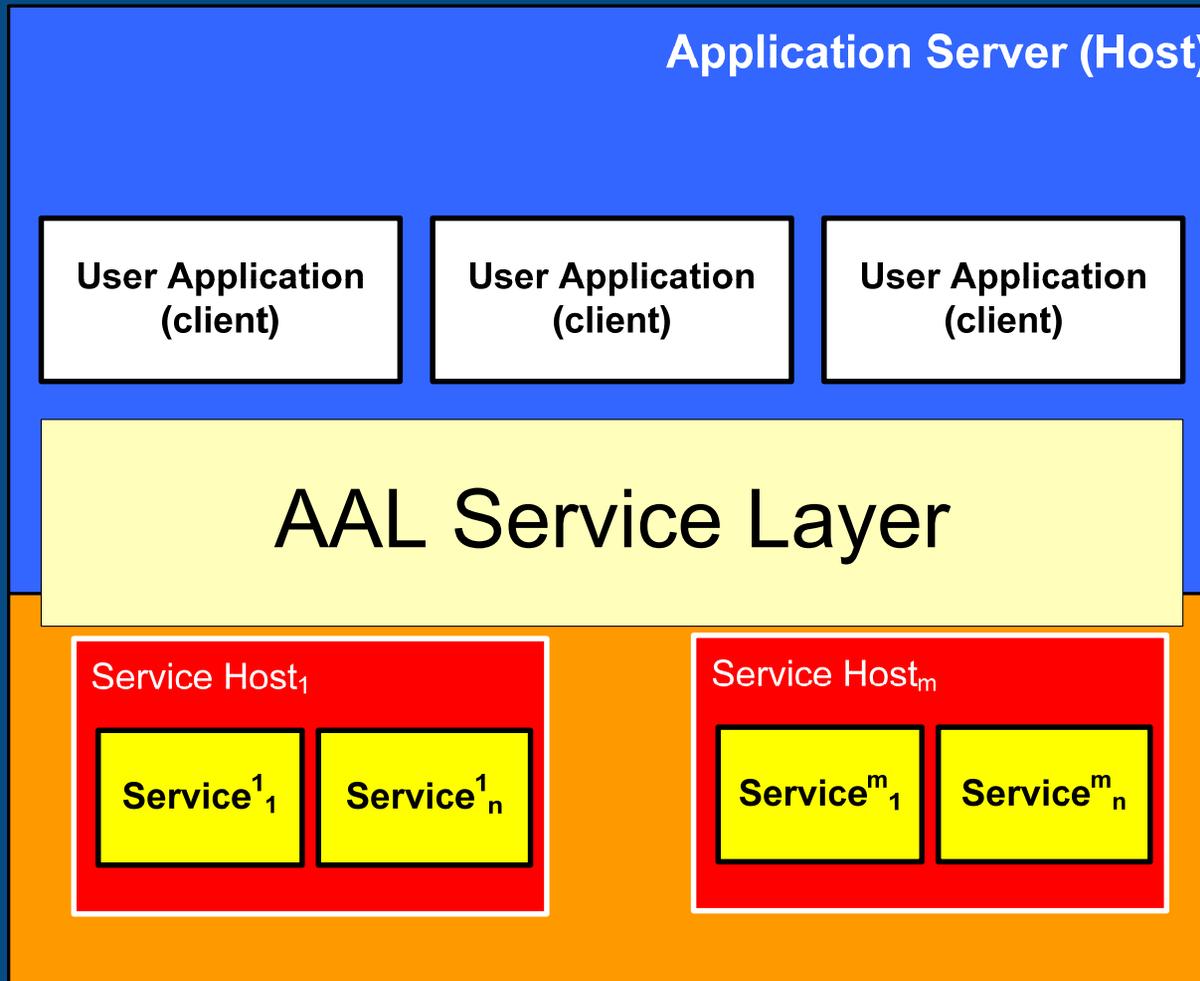
- Asynchronous Programming
 - Well suited for embarrassingly parallel workloads
 - Well suited for location abstraction
 - High latency call stacks do not block application
 - General model
 - Does not preclude more restrictive models
- AAL provides constructs to simplify use
 - Synchronization, Message delivery
 - Several message delivery modes
- Exception Handling



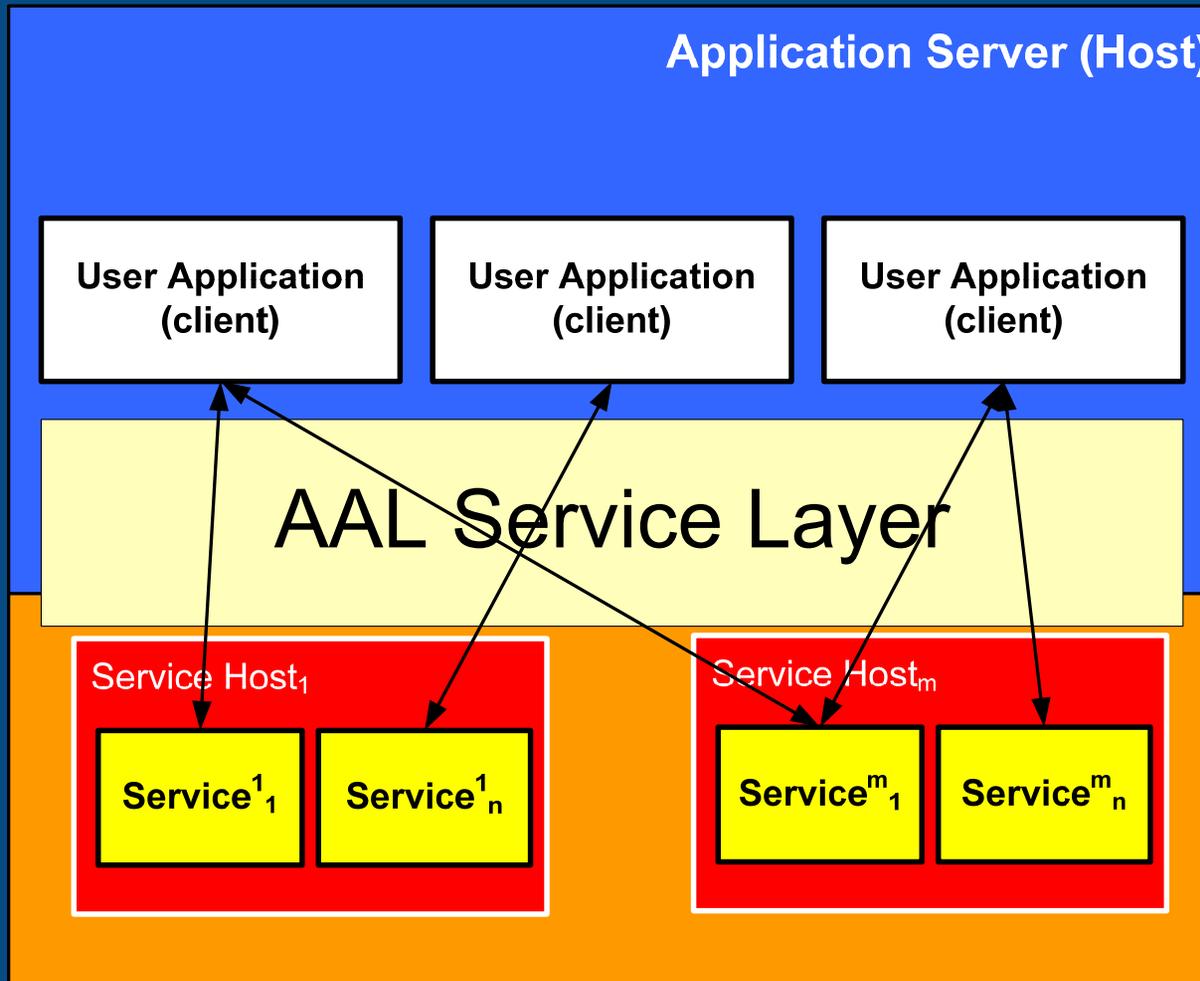
SOA Model



Client/Service Model



Client/Service Model



A large, 3D graphic of the words "THANK YOU" in a light blue, sans-serif font. The letters are rendered with a perspective effect, appearing to float above a dark blue rectangular shadow. The "Y" and "O" are particularly prominent due to their size and depth.

More Information:

- www.intel.com/go/QuickAssist
- Joe.Grecco@Intel.com

Legal Disclaimer

- 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.
- All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Intel, processors, chipsets, and desktop boards 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.
- Code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user
- Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.
- Intel, Intel Inside, and the Intel logo are trademarks of Intel Corporation in the United States and other countries.
- *Other names and brands may be claimed as the property of others.
- Copyright © 2008 Intel Corporation.



Risk Factors

This presentation contains forward-looking statements that involve a number of risks and uncertainties. These statements do not reflect the potential impact of any mergers, acquisitions, divestitures, investments or other similar transactions that may be completed in the future, with the exception of the Numonyx transaction. Our forward-looking statements for 2008 reflect the expectation that the Numonyx transaction will close during the first quarter. The information presented is accurate only as of today's date and will not be updated. In addition to any factors discussed in the presentation, the important factors that could cause actual results to differ materially include the following: Factors that could cause demand to be different from Intel's expectations include changes in business and economic conditions, including conditions in the credit market that could affect consumer confidence; customer acceptance of Intel's and competitors' products; changes in customer order patterns, including order cancellations; and changes in the level of inventory at customers. Intel's results could be affected by the timing of closing of acquisitions and divestitures. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Additionally, Intel is in the process of transitioning to its next generation of products on 45 nm process technology, and there could be execution issues associated with these changes, including product defects and errata along with lower than anticipated manufacturing yields. Revenue and the gross margin percentage are affected by the timing of new Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel's response to such actions; Intel's ability to respond quickly to technological developments and to incorporate new features into its products; and the availability of sufficient components from suppliers to meet demand. The gross margin percentage could vary significantly from expectations based on changes in revenue levels; product mix and pricing; capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; excess or obsolete inventory; manufacturing yields; changes in unit costs; impairments of long-lived assets, including manufacturing, assembly/test and intangible assets; and the timing and execution of the manufacturing ramp and associated costs, including start-up costs. Expenses, particularly certain marketing and compensation expenses, vary depending on the level of demand for Intel's products, the level of revenue and profits, and impairments of long-lived assets. Intel is in the midst of a structure and efficiency program that is resulting in several actions that could have an impact on expected expense levels and gross margin. We expect to complete the divestiture of our NOR flash memory assets to Numonyx. A delay or failure of the transaction to close, or a change in the financial performance of the contributed businesses could have a negative impact on our financial statements. Intel's equity proportion of the new company's results will be reflected on its financial statements below operating income and with a one quarter lag. Intel's results could be affected by the amount, type, and valuation of share-based awards granted as well as the amount of awards cancelled due to employee turnover and the timing of award exercises by employees. Intel's results could be impacted by adverse economic, social, political and physical/infrastructure conditions in the countries in which Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Intel's results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust and other issues, such as the litigation and regulatory matters described in Intel's SEC reports. A detailed discussion of these and other factors that could affect Intel's results is included in Intel's SEC filings, including the report on Form 10-K for the fiscal year ended December 29, 2007.

