# Distributed I/O Processing for High Speed Networks

## PCI Spring Developers Conference

**Byron Gillespie**

**Strategic Development Manager**

**Intel Corporation**

intel®

# Byron Gillespie

**Intel Corporation**
**5000 West Chandler Blvd. MS CH6-319**
**Chandler, Arizona 85226**
**(602) 554-2653, Fax (602) 554-7347**

## Current Activities

**Byron Gillespie works as a strategic development manager for Intel's Computing Division. He is responsible for working with customers to define the requirements for future i960® microprocessor products. Currently, Byron is providing technical support for the 80960 RP Intelligent I/O Processor and leading the system validation effort for the i960 RP processor.**

## Author Background

**Since joining Intel in 1991, he has defined three i960 microprocessors, including the i960 RP processor for intelligent PCI I/O applications. Previously, he had eight years experience writing software for embedded avionics applications. The embedded applications used a variety of Intel processors including the superscalar i960 CA microprocessor Gillespie received a B.S. in computer science from Northern Arizona University in 1983.**

intel®

# Distributed I/O Agenda

- Current Intelligent I/O Application Model
- System Performance Issues
- Next Generation Intelligent I/O Application Model
- Hardware PCI enhancements
- Focus on reducing host processor interrupt servicing and processing

# Intelligent I/O Application Model

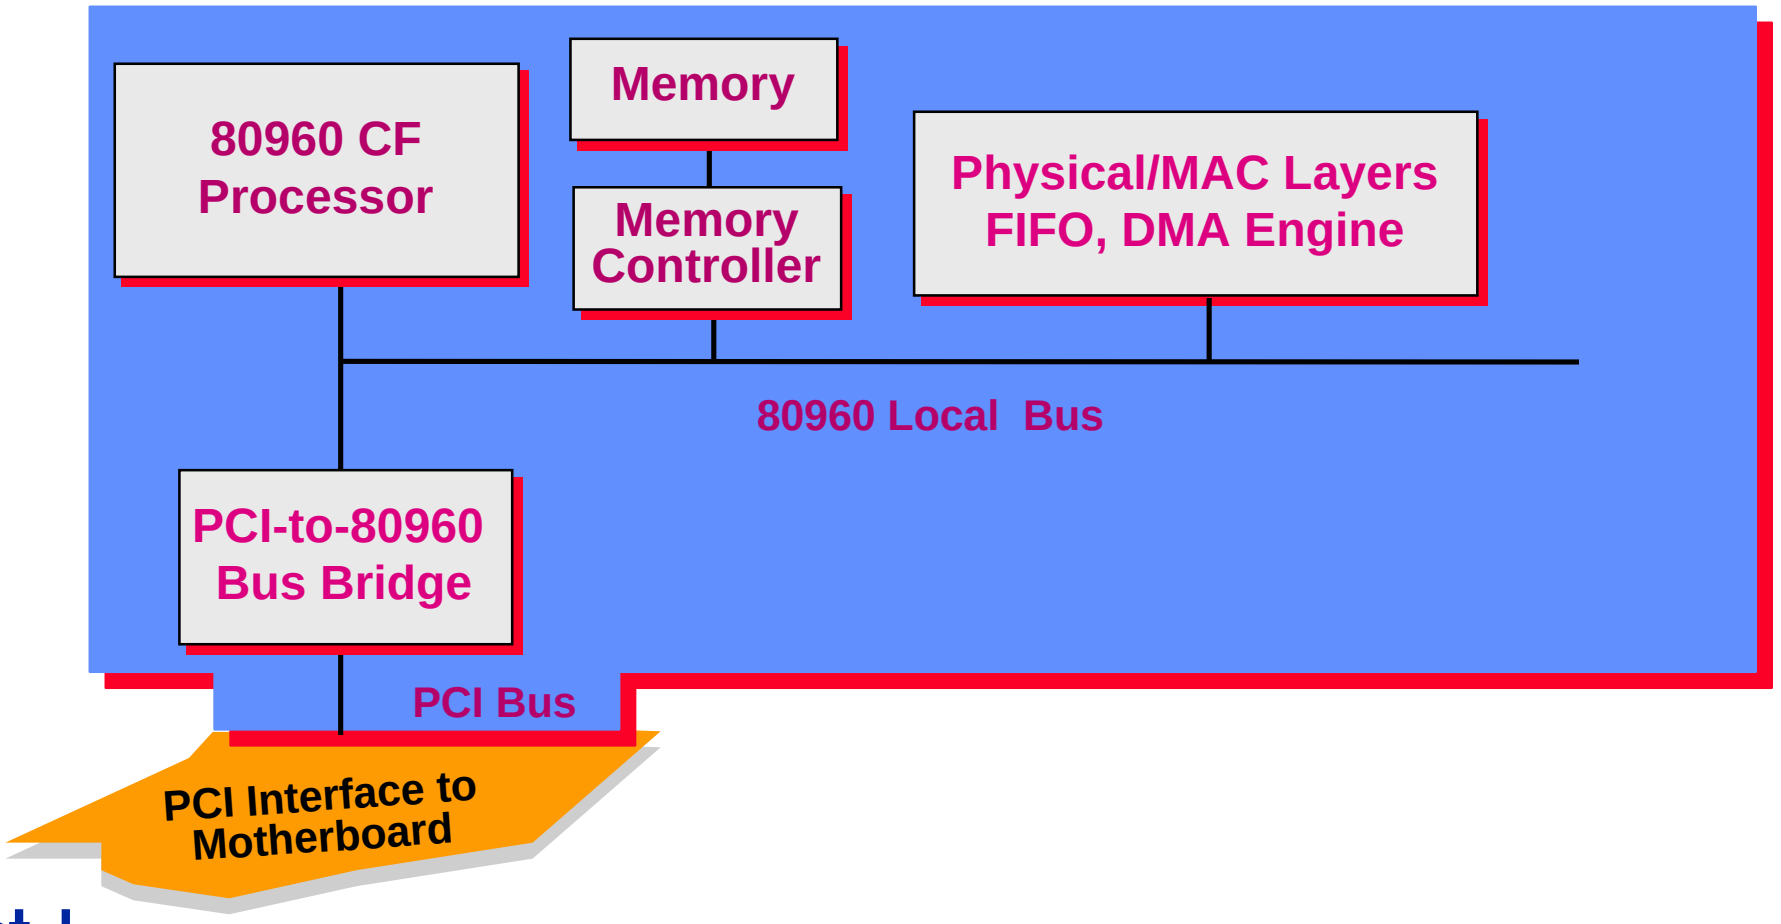


intel®

# Role of an I/O Processor

- Reduced host CPU utilization for I/O tasks
  - Handling interrupts
- Improve the performance of the client/server environment
  - Server I/O:
    - Simultaneous "conversation" with many clients
    - I/O scalability required

# Cost of Interrupt Analysis

- 100 MB/sec Ethernet Example using current generation of intelligent adapter I/O subsystem
- Dumb Ethernet card example
  - Host CPU utilization = 50%
  - Interrupt latency impacts network utilization
    - Lower inter-frame spacing between packets
    - ~ 8us versus ~12us

# System Analysis

- Re-directing host interrupts can significantly improve system performance
- Less contention for host memory resulting in higher PCI performance
- Supports greater number of chain links
  - Keeps the transmit pipe full all the time
  - Receive data buffering and packing

# Fast Ethernet Card Demonstrates Benefits of Intelligent I/O For Networks

Server performance with four PRO/100 adapters

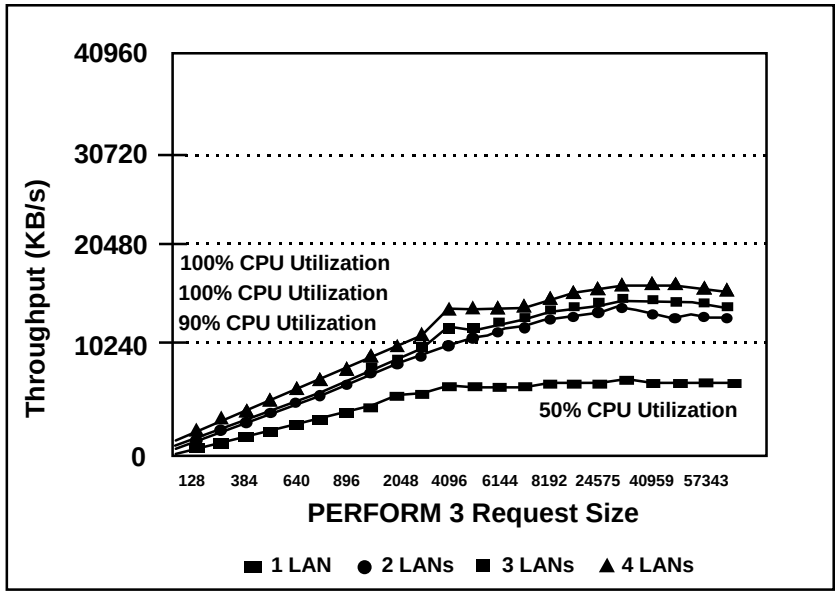


Server performance with four intelligent PRO/100 Server Adapters

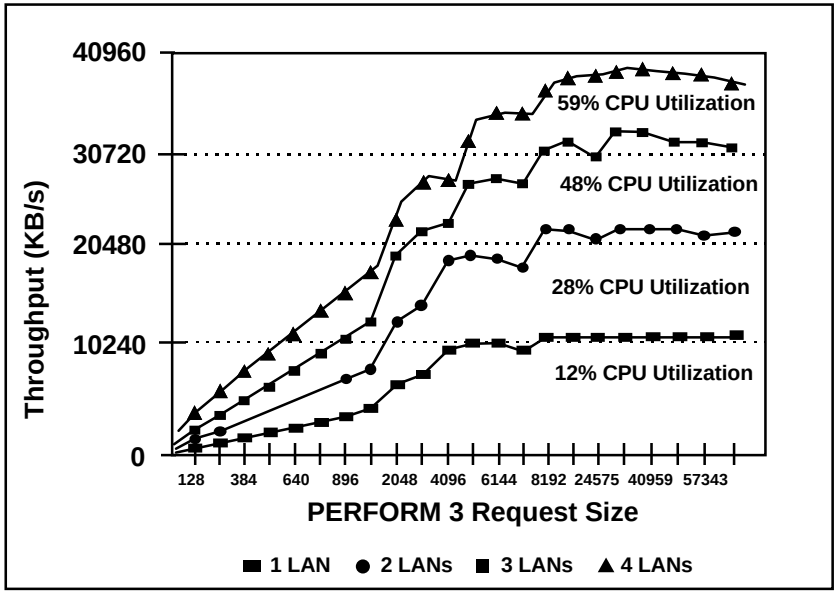


Source: Intel

intel®

# Focus On System Performance

- System architecture affects I/O performance
- Interrupt sharing
  - Affects both receive and transmit frame data
  - Overall network bandwidth is affected
- Watch surrounding PCI components and systems
  - Poorly designed PCI interfaces can reduce PCI performance
  - Chipset performance affects host memory access latency and PCI bandwidth
  - Disk performance affects network response time

# Next-Generation Intelligent I/O Application Model

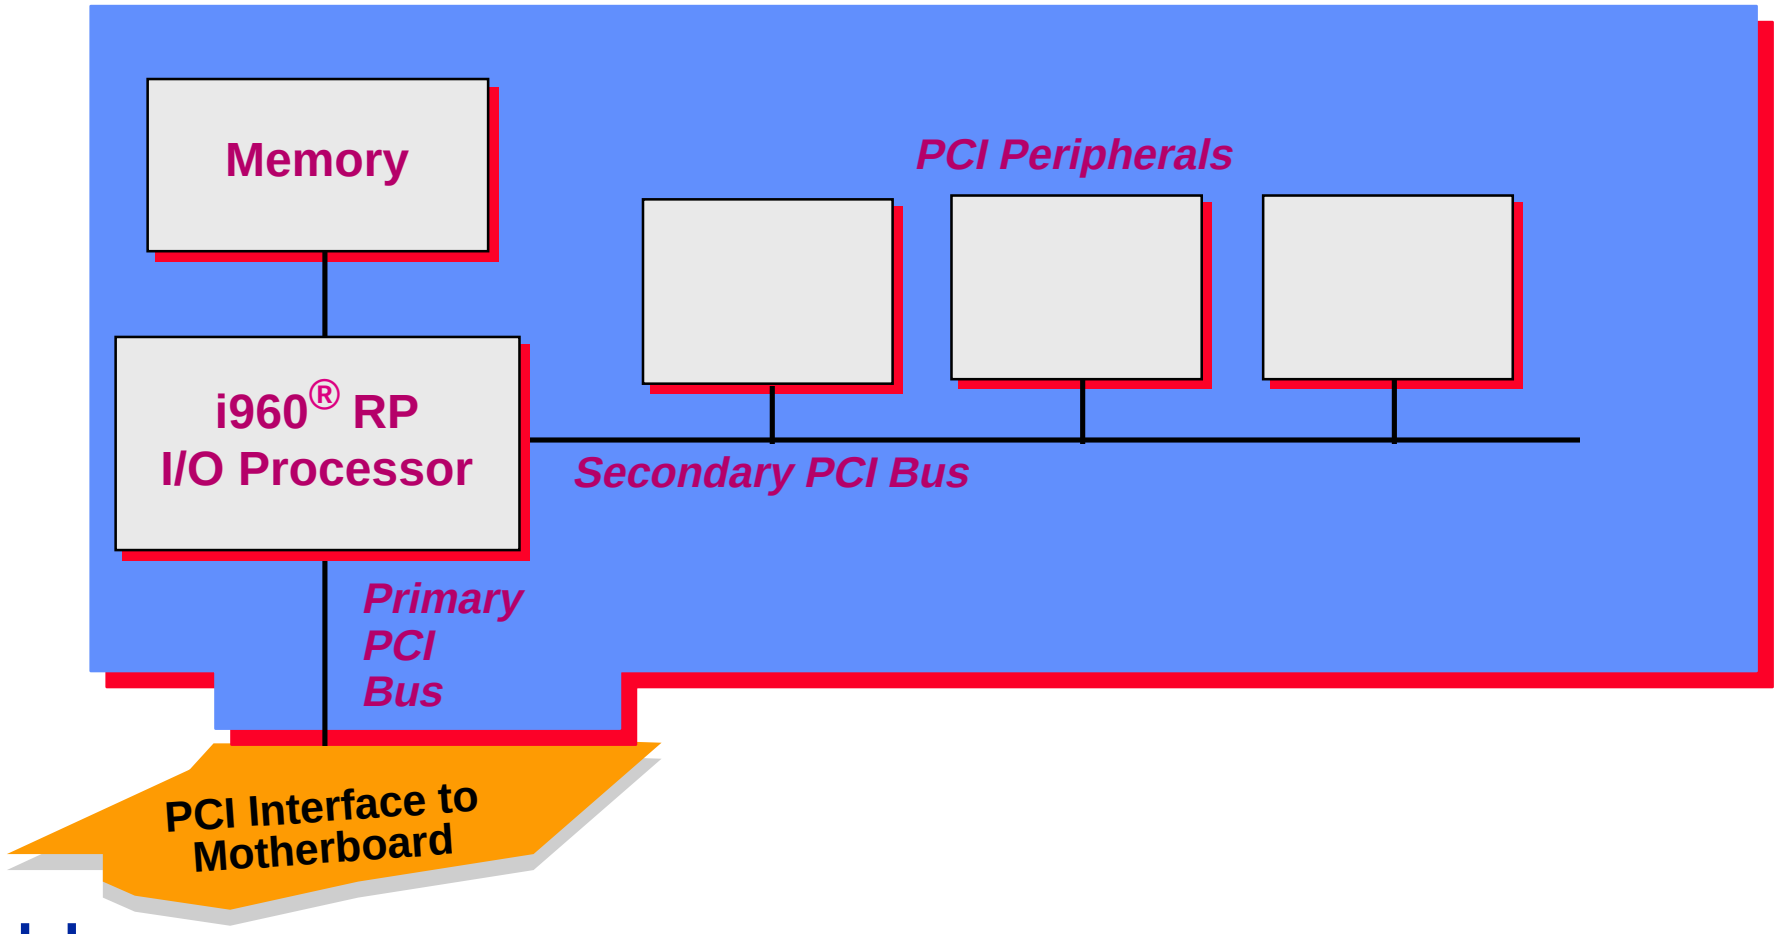


intel®

# i960® RP Processor Block Diagram

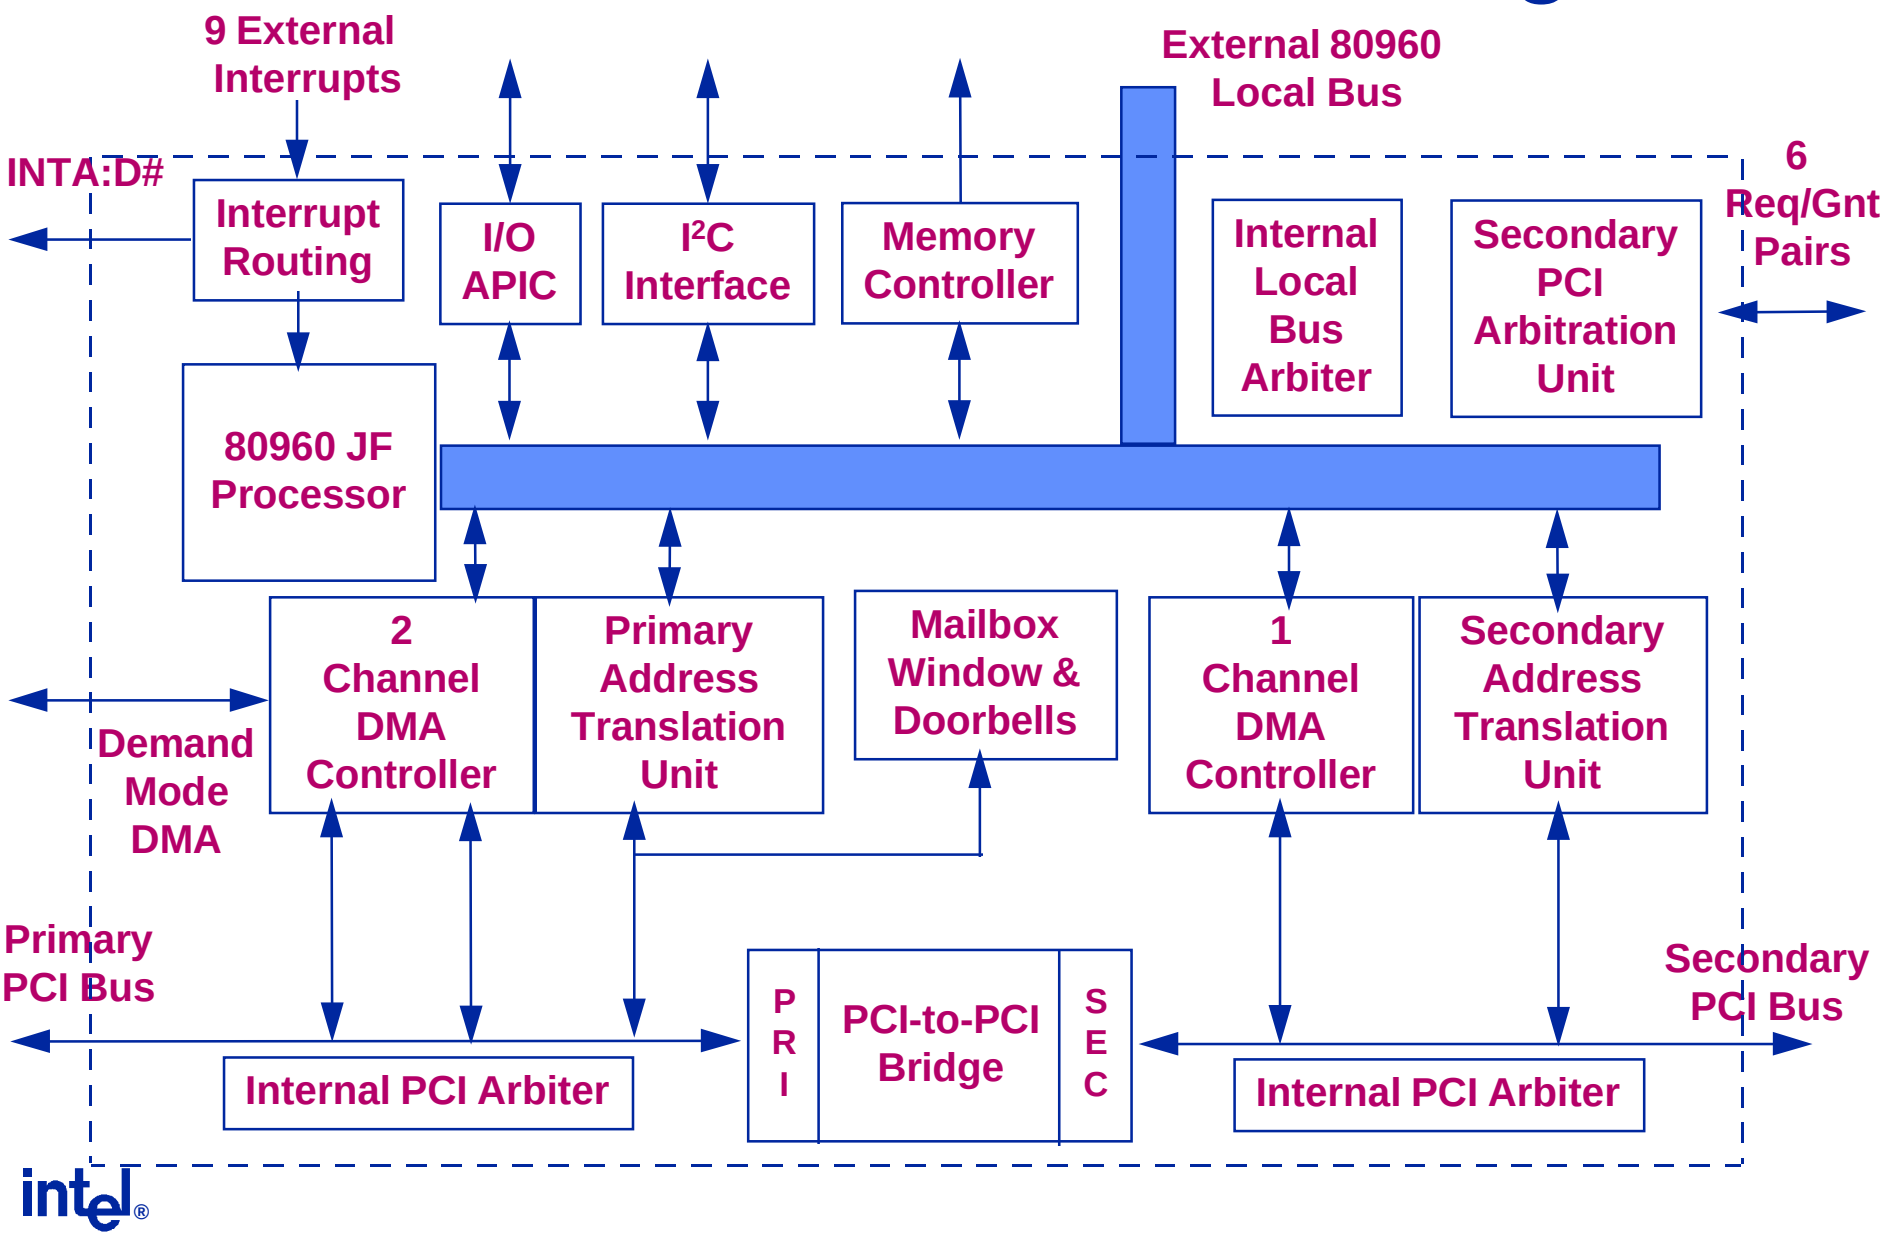

# Hardware Enhancements to Supplement System Performance

- PCI bus local to the card
  - Utilize PCI components designed with PCI performance in mind
- Use PCI commands to fullest
  - Intelligent prefetching strategy for MRMultiple and MRLine commands
  - MWInvalidate for large data transfers
- Hardware Message Queuing Mechanism
  - Simplifies and reduces driver software to reduce latency

# PCI Read Performance

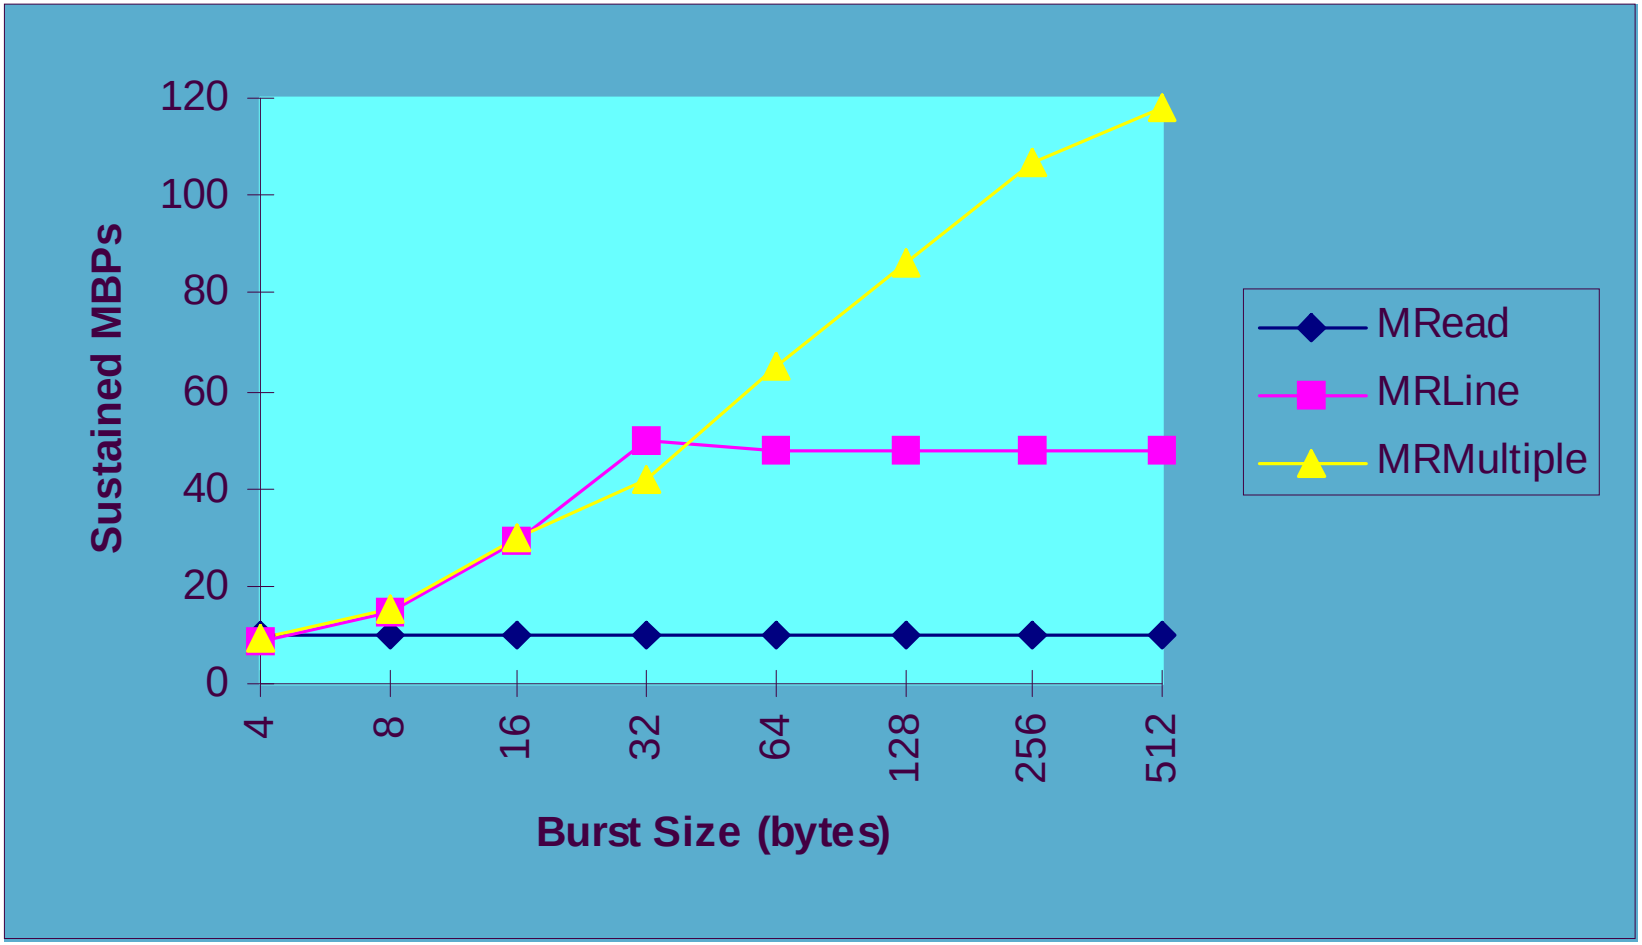


Source: Intel IAL Lab Measurement

intel®

# PCI Write Performance

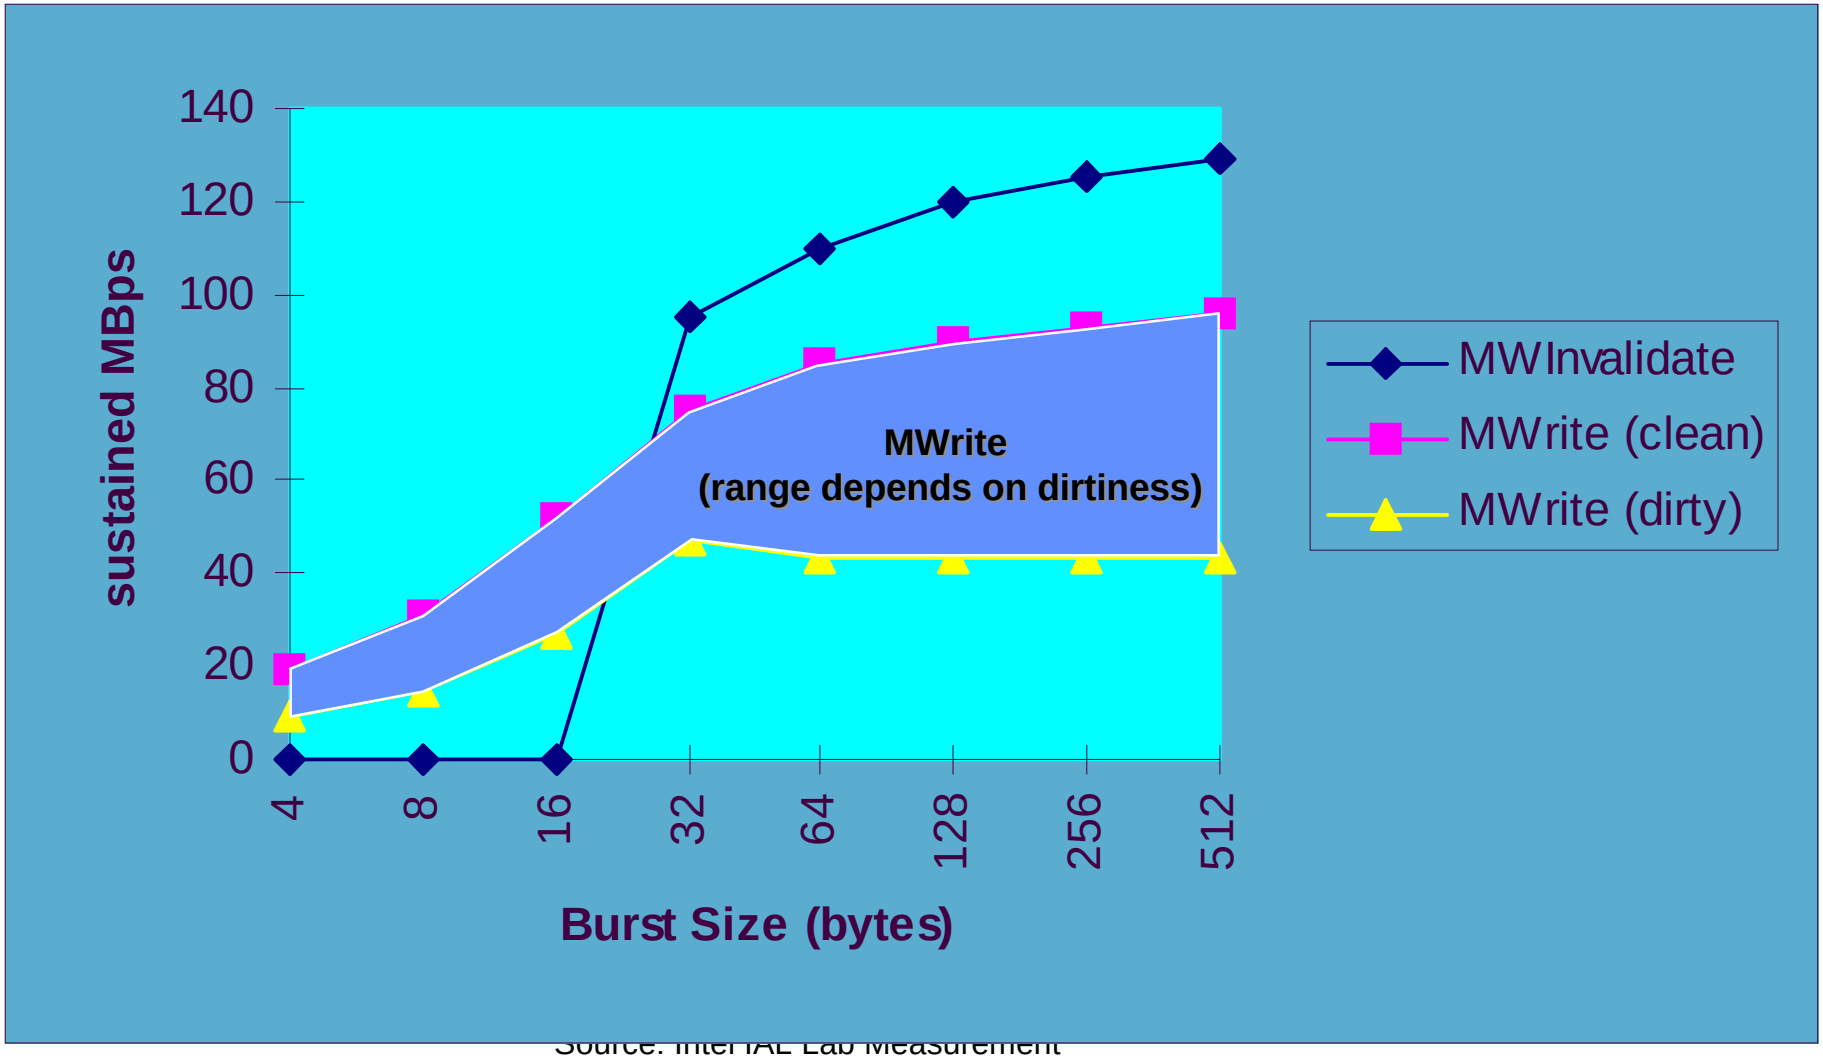


Source: Intel IAL Lab Measurement

intel®

# Reducing Interrupts with Intelligent I/O Subsystems

- Request and Response Queue Model
  - Reduces Host Interrupts
- New Problem: Spin locks as processors contend for queues in SMP systems
  - Utilize the "atomic" nature of PCI to avoid spin locks

# Request & Response Queue Model

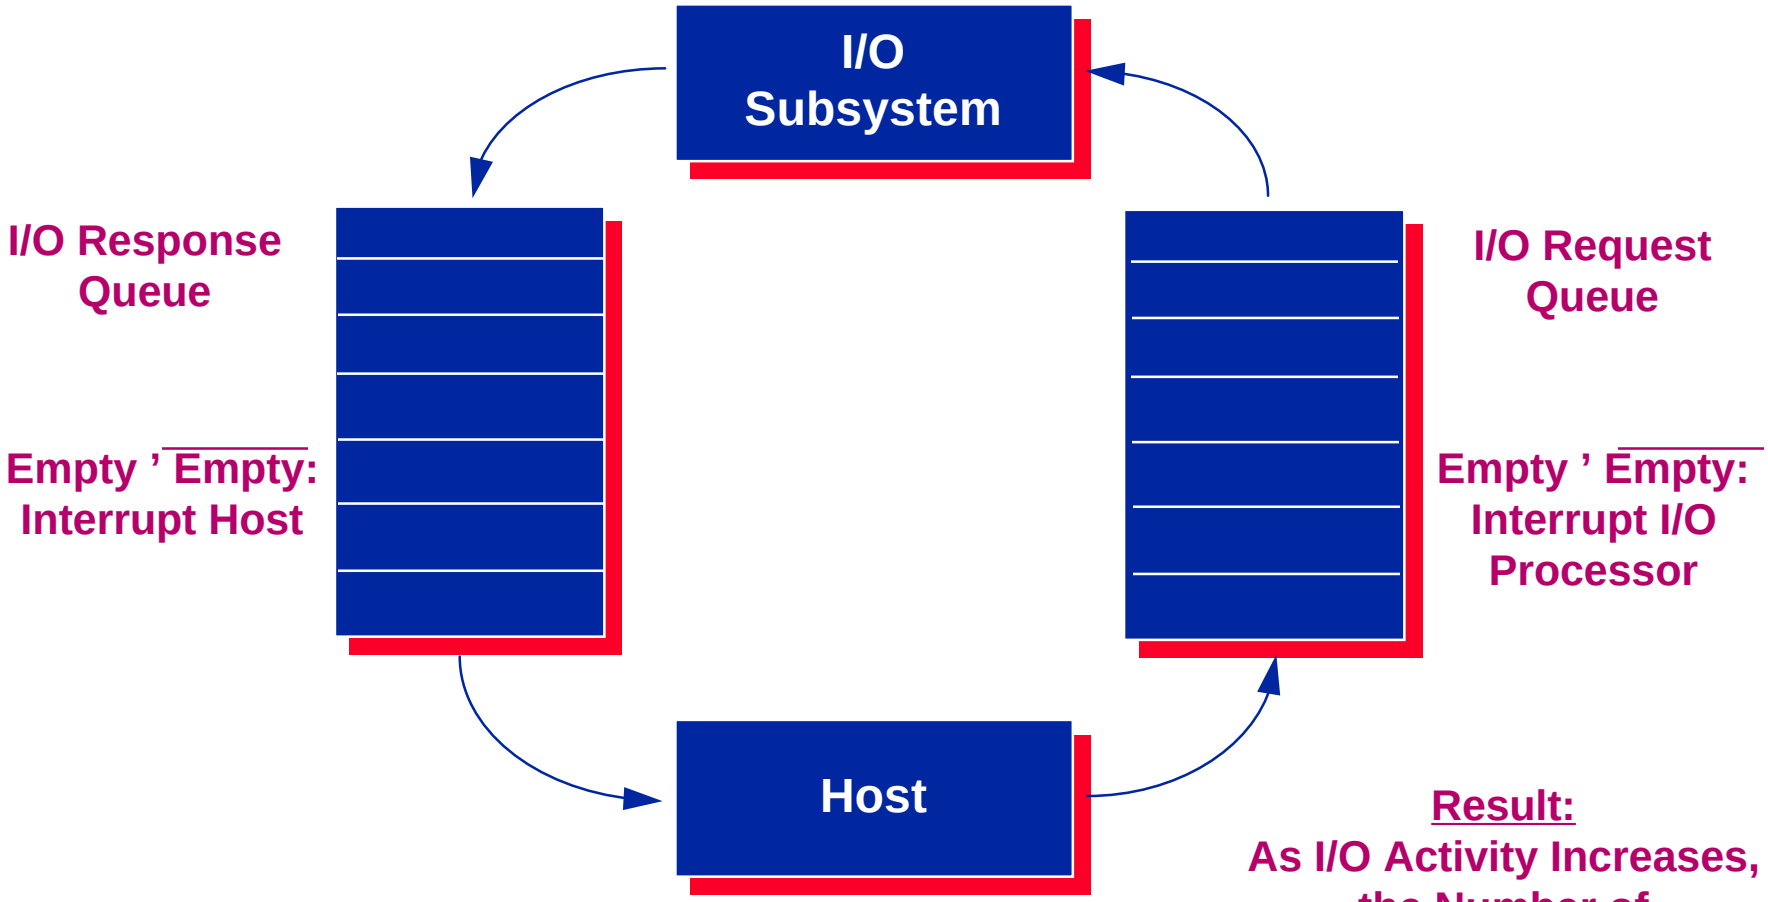


**Result:**
As I/O Activity Increases, the Number of Interrupts Decreases

intel®

# Messaging Unit: Circular Queues

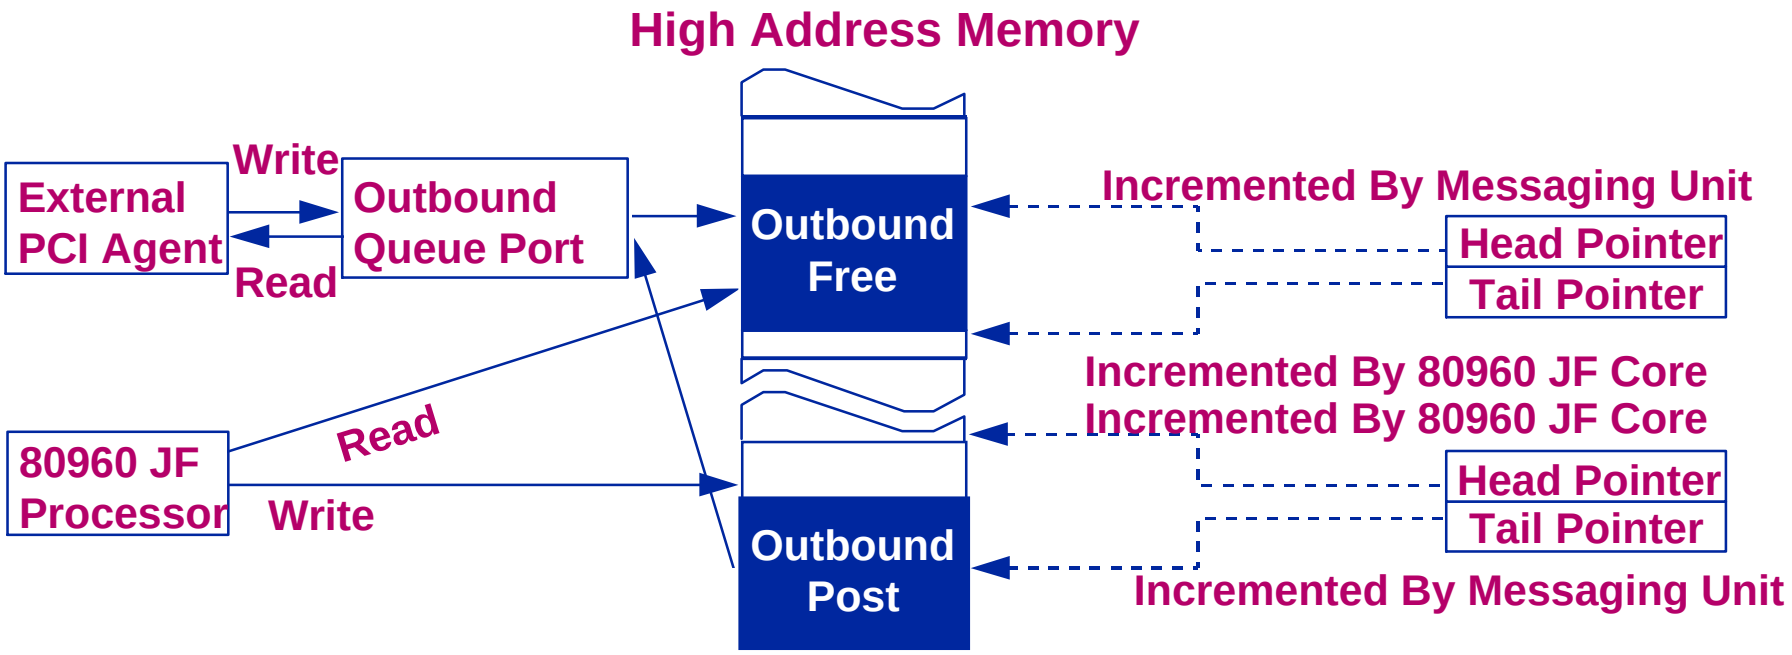


intel®

# Hardware-Assisted Queue Structure

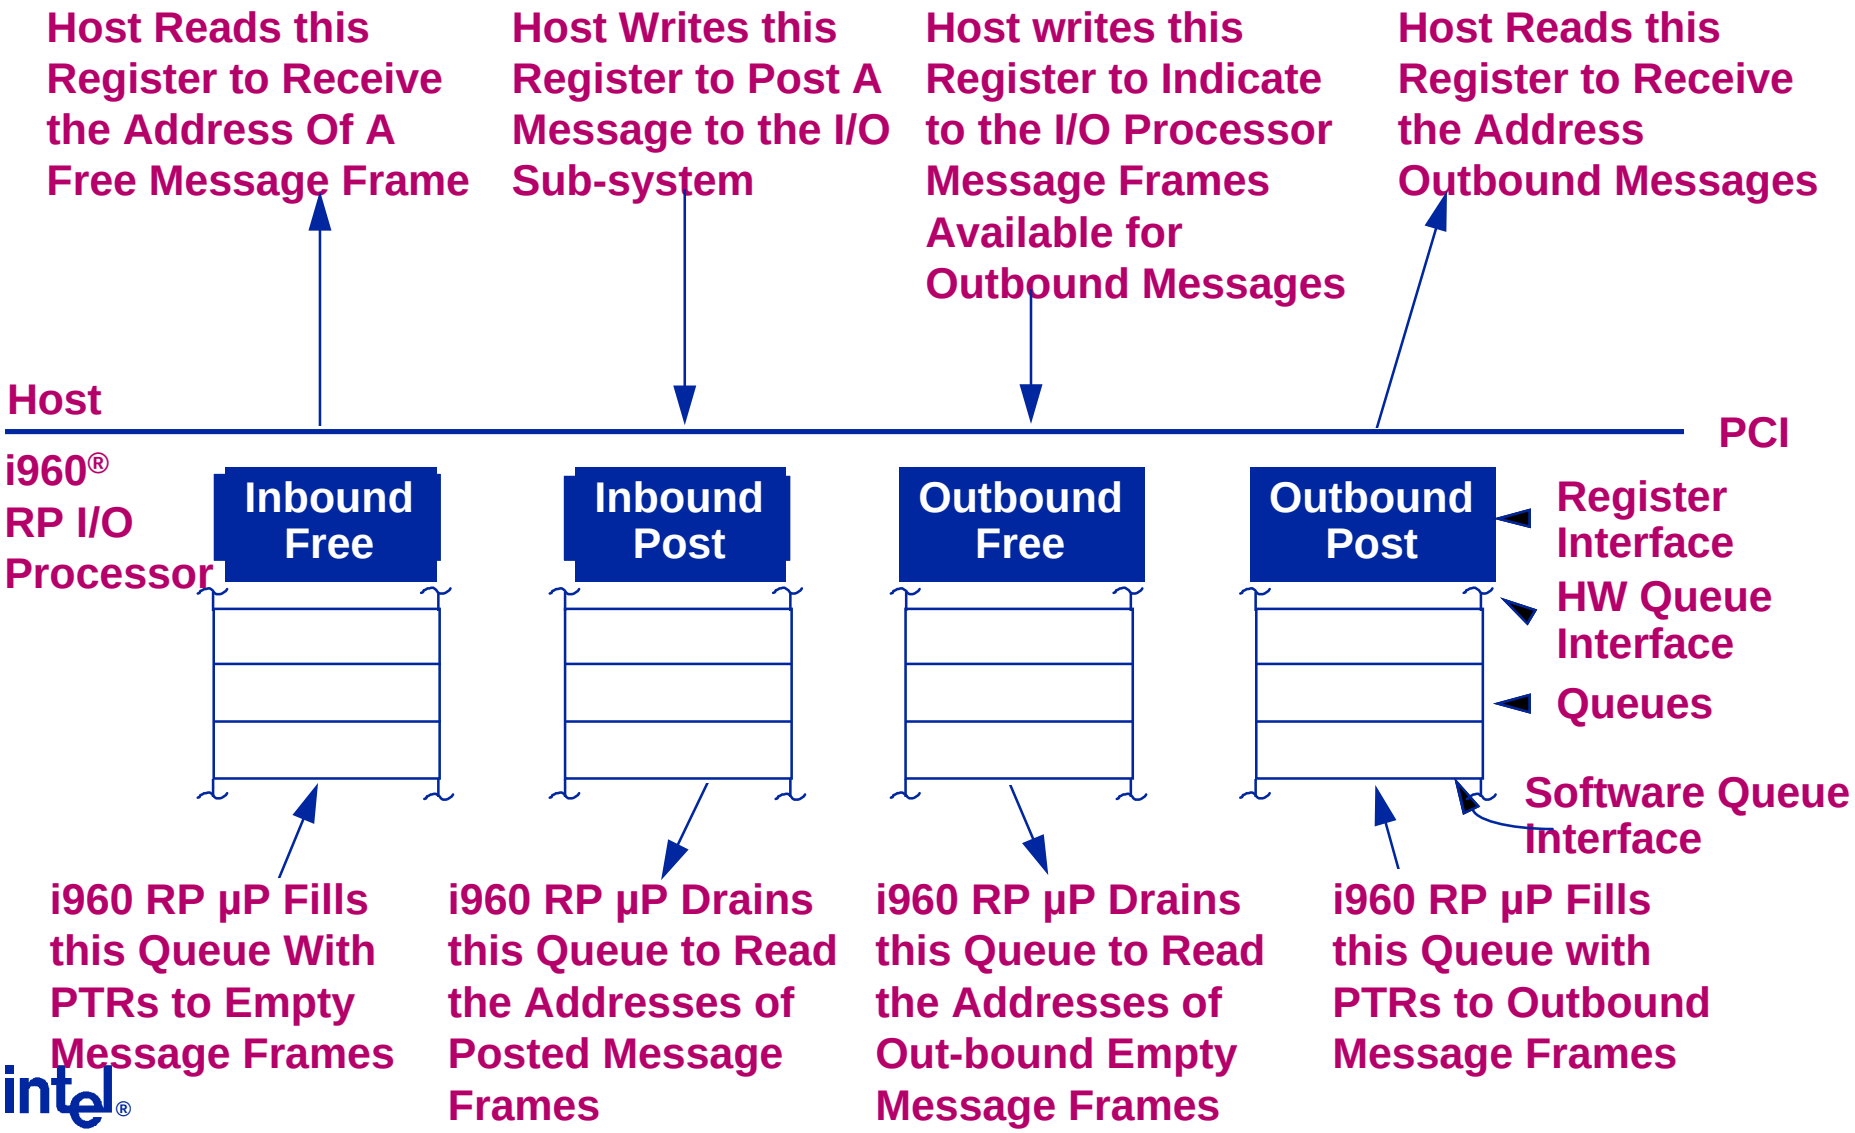


intel®

# Initialization

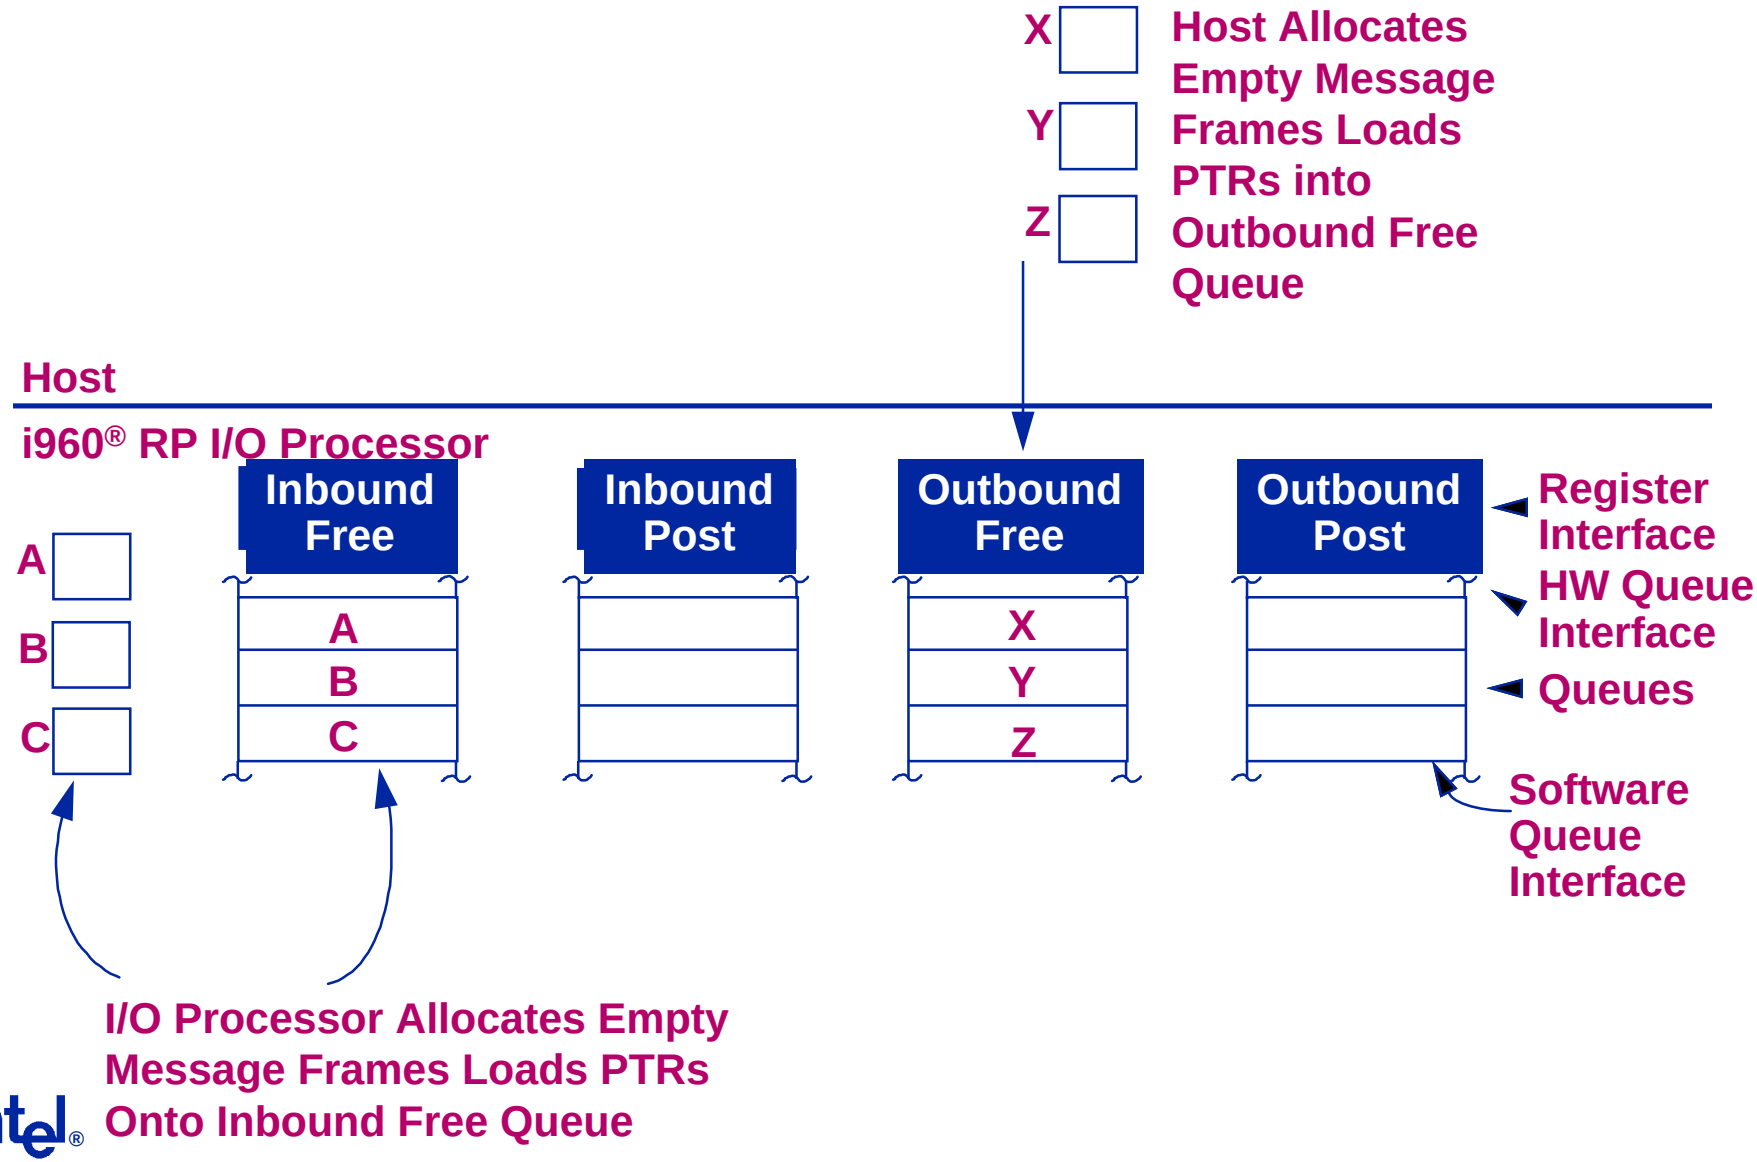

# Posting an I/O Request

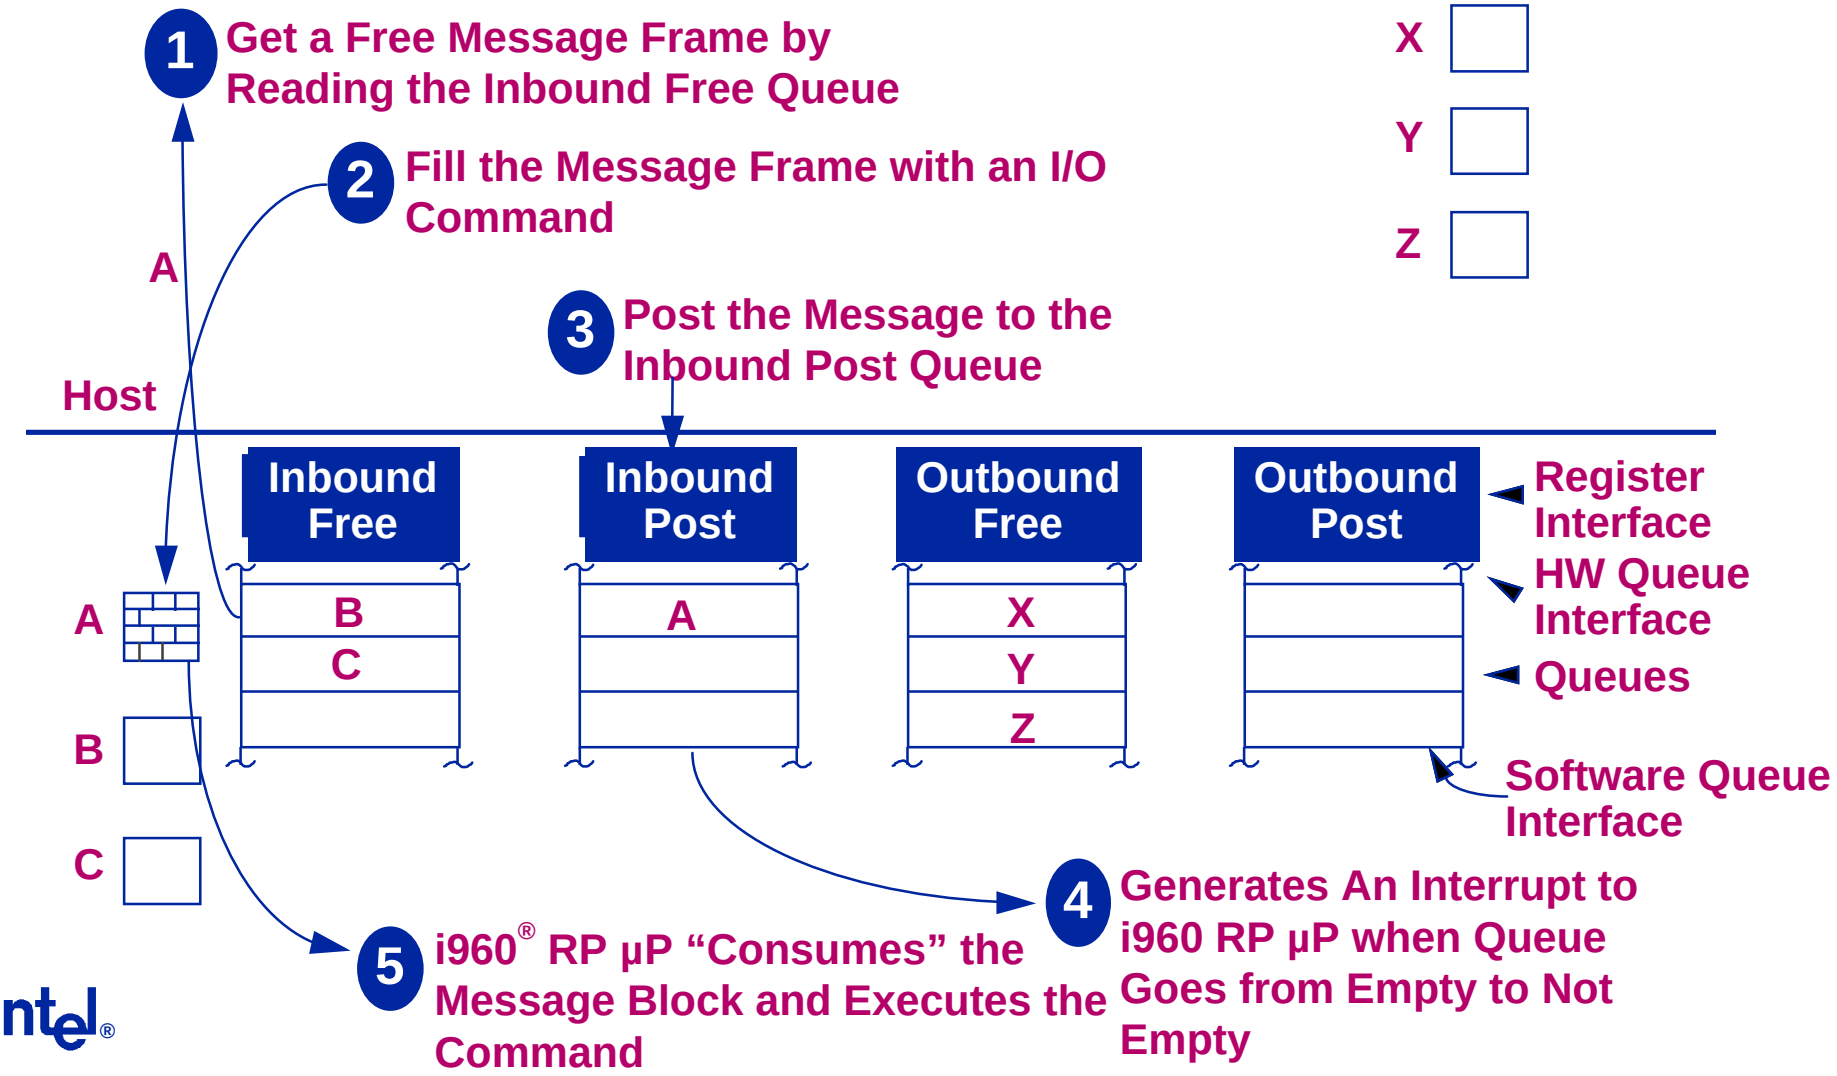

# Reply to an I/O Command

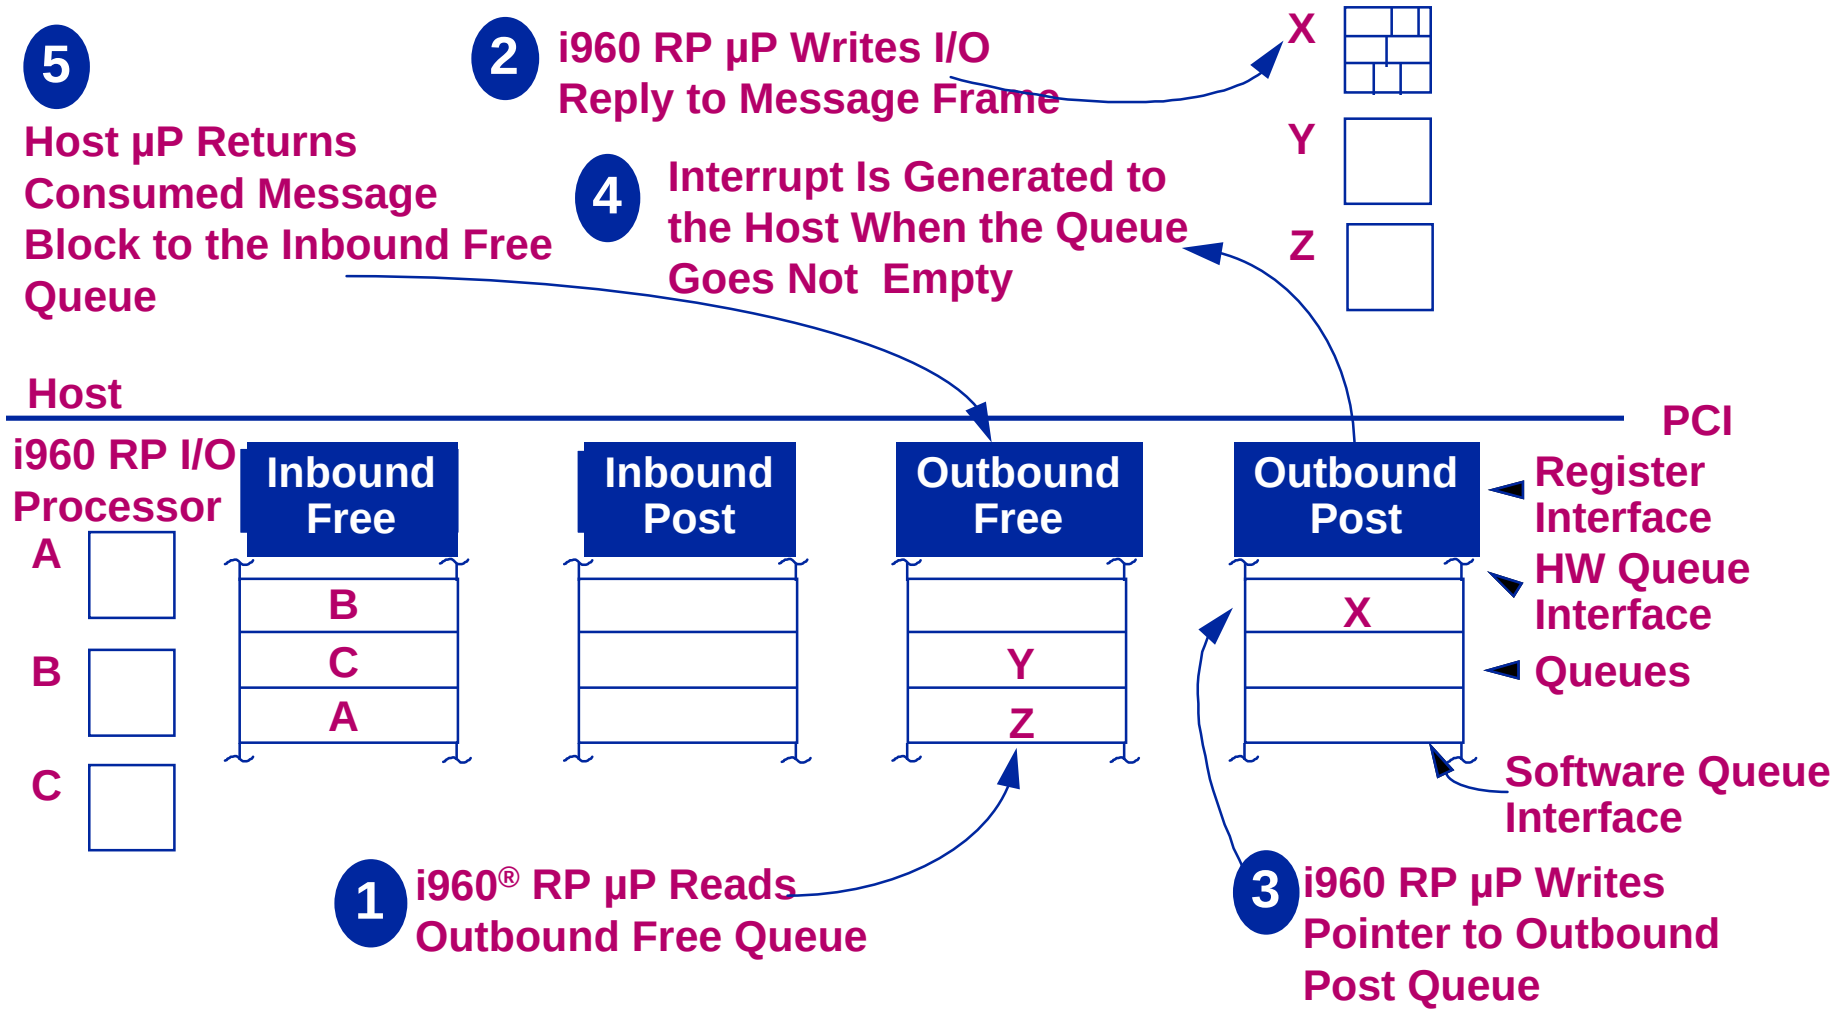


intel®

# Summary

- Distributed I/O processing brings enterprise class I/O to PC servers by
  - Improving server throughput
  - Allowing more host CPU cycles for applications
  - Enables server I/O to scale with the frequency of the host CPU and in SMP systems
  - Improves the client/server performance
- The i960® RP processor enables high performance I/O subsystems
  - Removes the spin-locks incurred by SMP systems
  - Removes I/O bottlenecks for servers

intel®