Intel Architecture: Features & Futures

For Servers & Workstations

Stephen L. Smith

Corporate Vice President, Microprocessor Products Group
General Manager, Santa Clara Processor Division
Intel Corporation

intel®

Agenda

- High End IA Roadmap overview
- Foster Processor preview
- Merced™ Processor features & status
- McKinley Processor preview
- Summary

intel®

High End IA Roadmap

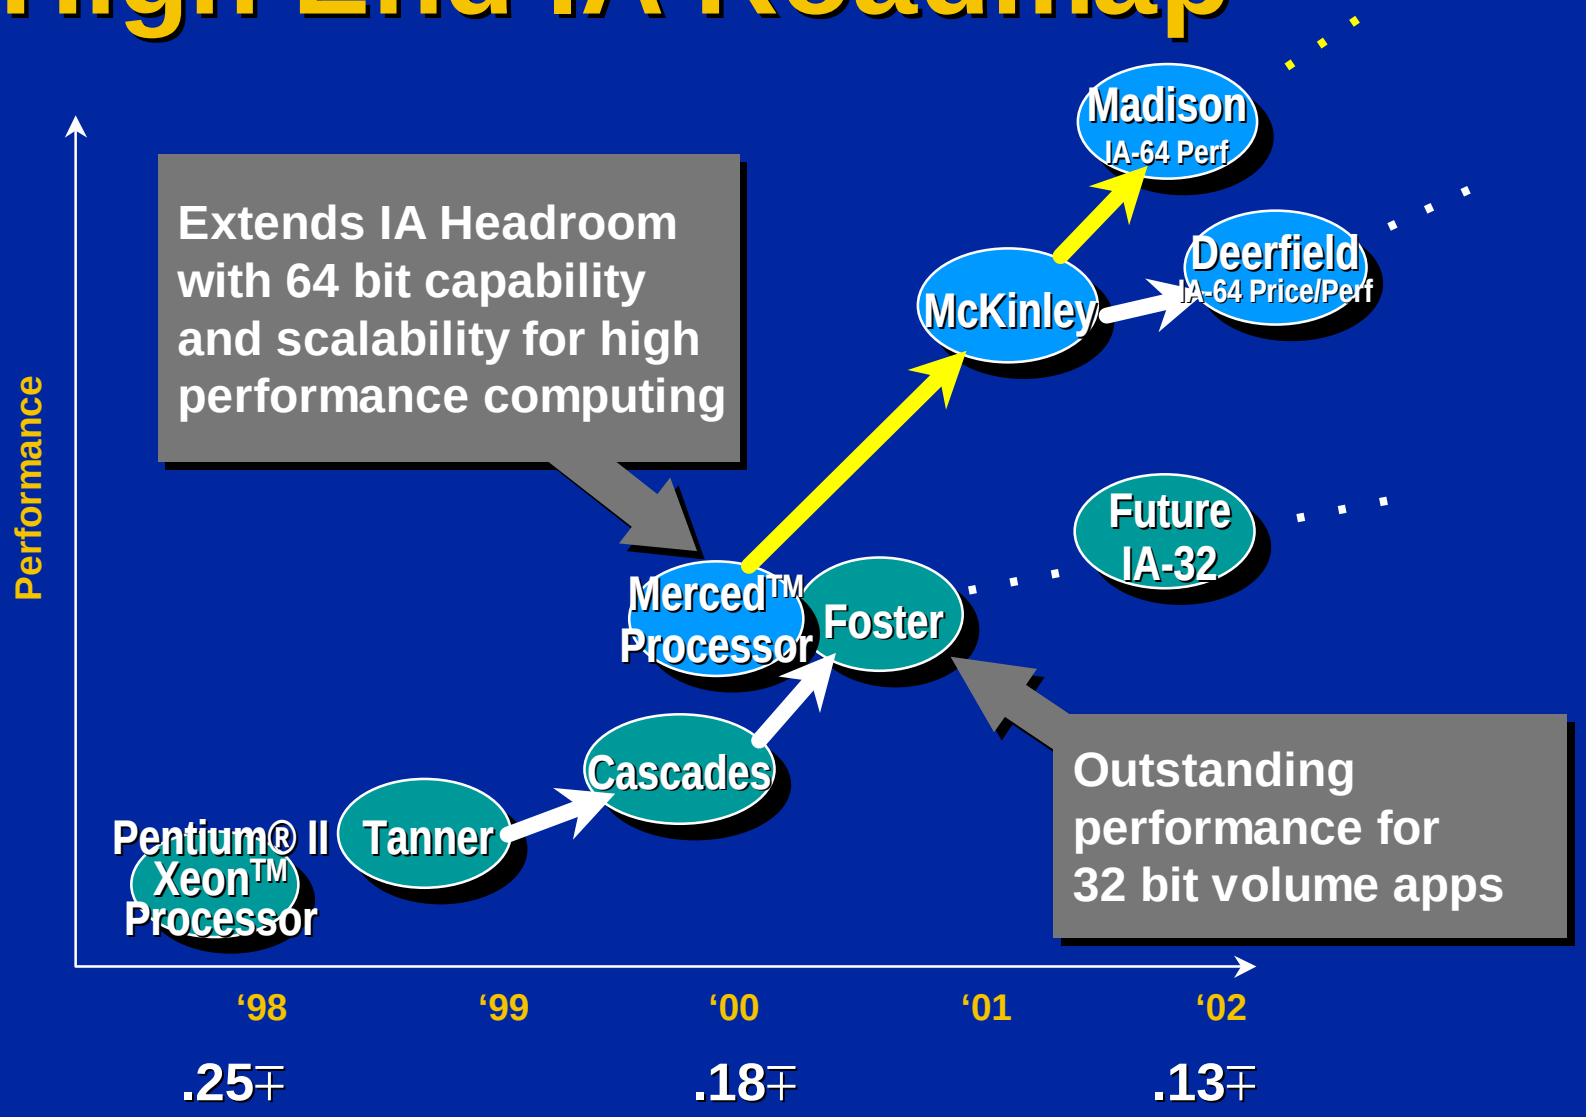

World class Server and Workstation Roadmap

intel®

*Intel code name

Solutions Focus

Server Apps

Data Warehouse/DSS
High Capacity OLTP
LOB/ERP
Security
Directory Services
Message Transaction
Collaborative
Publication
File & Print

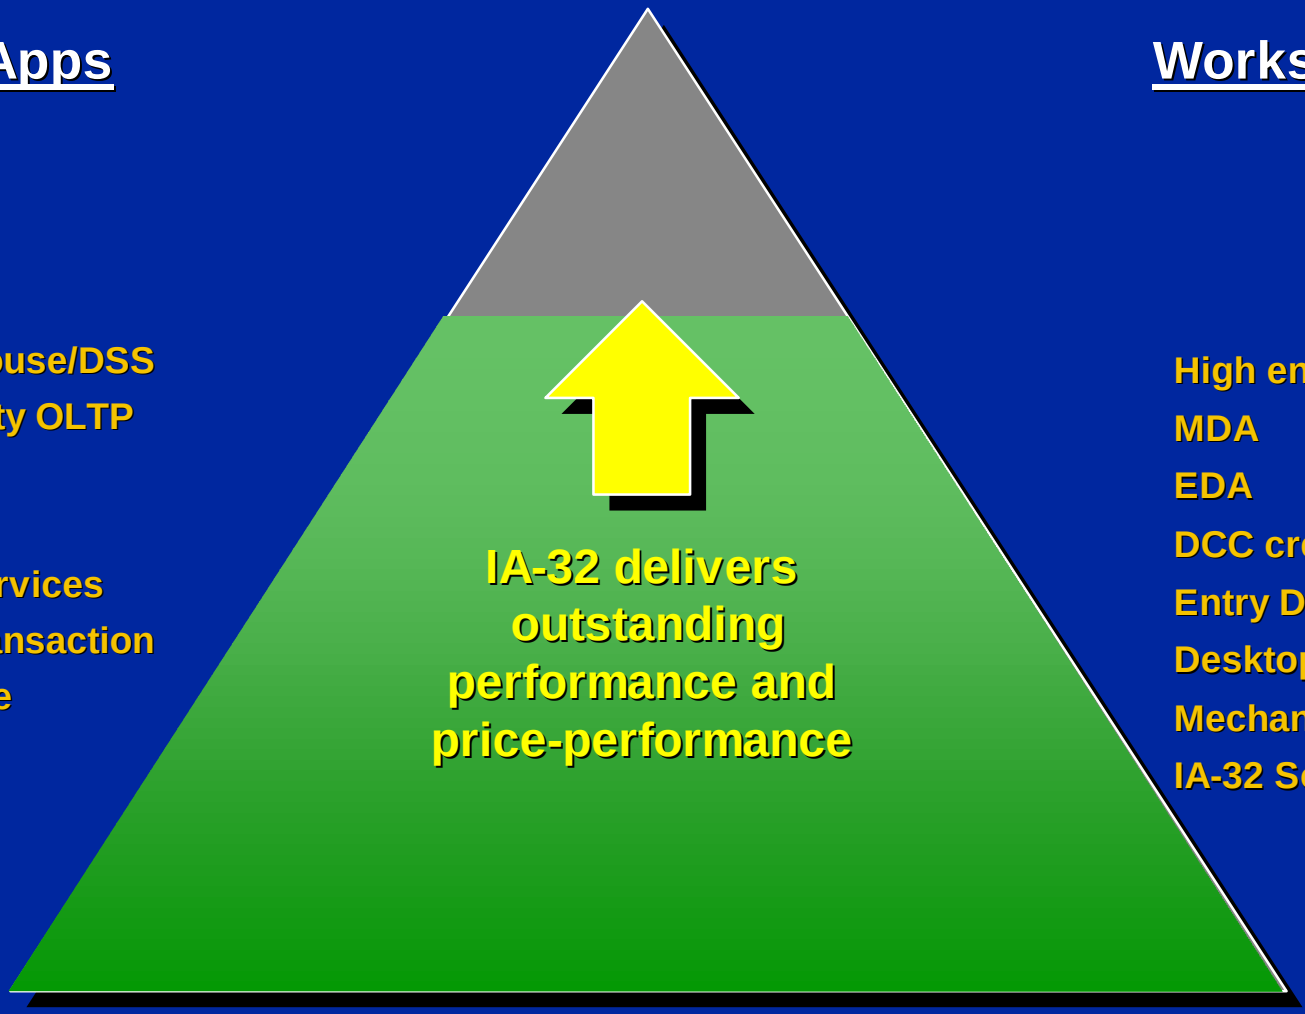

Workstation Apps

High end DCC
MDA
EDA
DCC creation/design
Entry DCC
Desktop Publishing
Mechanical Design
IA-32 Software Eng

intel®

*Intel code name

Solutions Focus

Server Apps

- High performance technical computing
- Very large memory DB
- Highest-capacity OLTP
- Highest end DSS Solutions
- Data Warehouse/DSS
- High Capacity OLTP
- LOB/ERP
- Security
- Directory Services
- Message Transaction
- Collaborative
- Publication
- File & Print

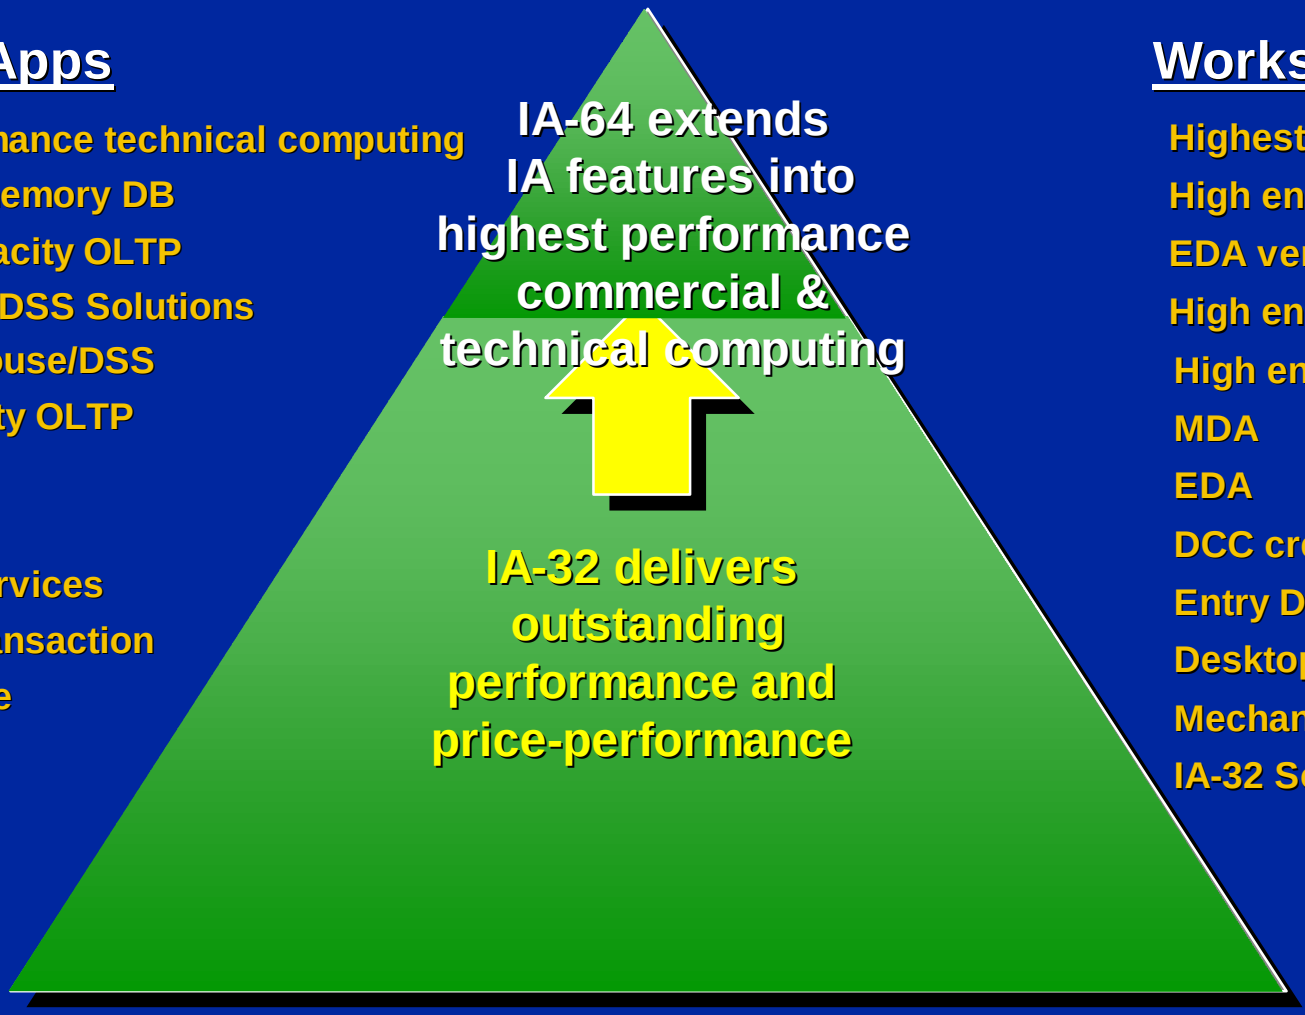

Workstation Apps

- Highest end CAE Analysis
- High end Software Eng
- EDA verification/synthesis
- High end technical analysis
- High end DCC
- MDA
- EDA
- DCC creation/design
- Entry DCC
- Desktop Publishing
- Mechanical Design
- IA-32 Software Eng

Complementing IA-32 and IA-64 products enable full range of Server/Workstation solutions

intel®

*Intel code name

Continuous IA-32 Innovations

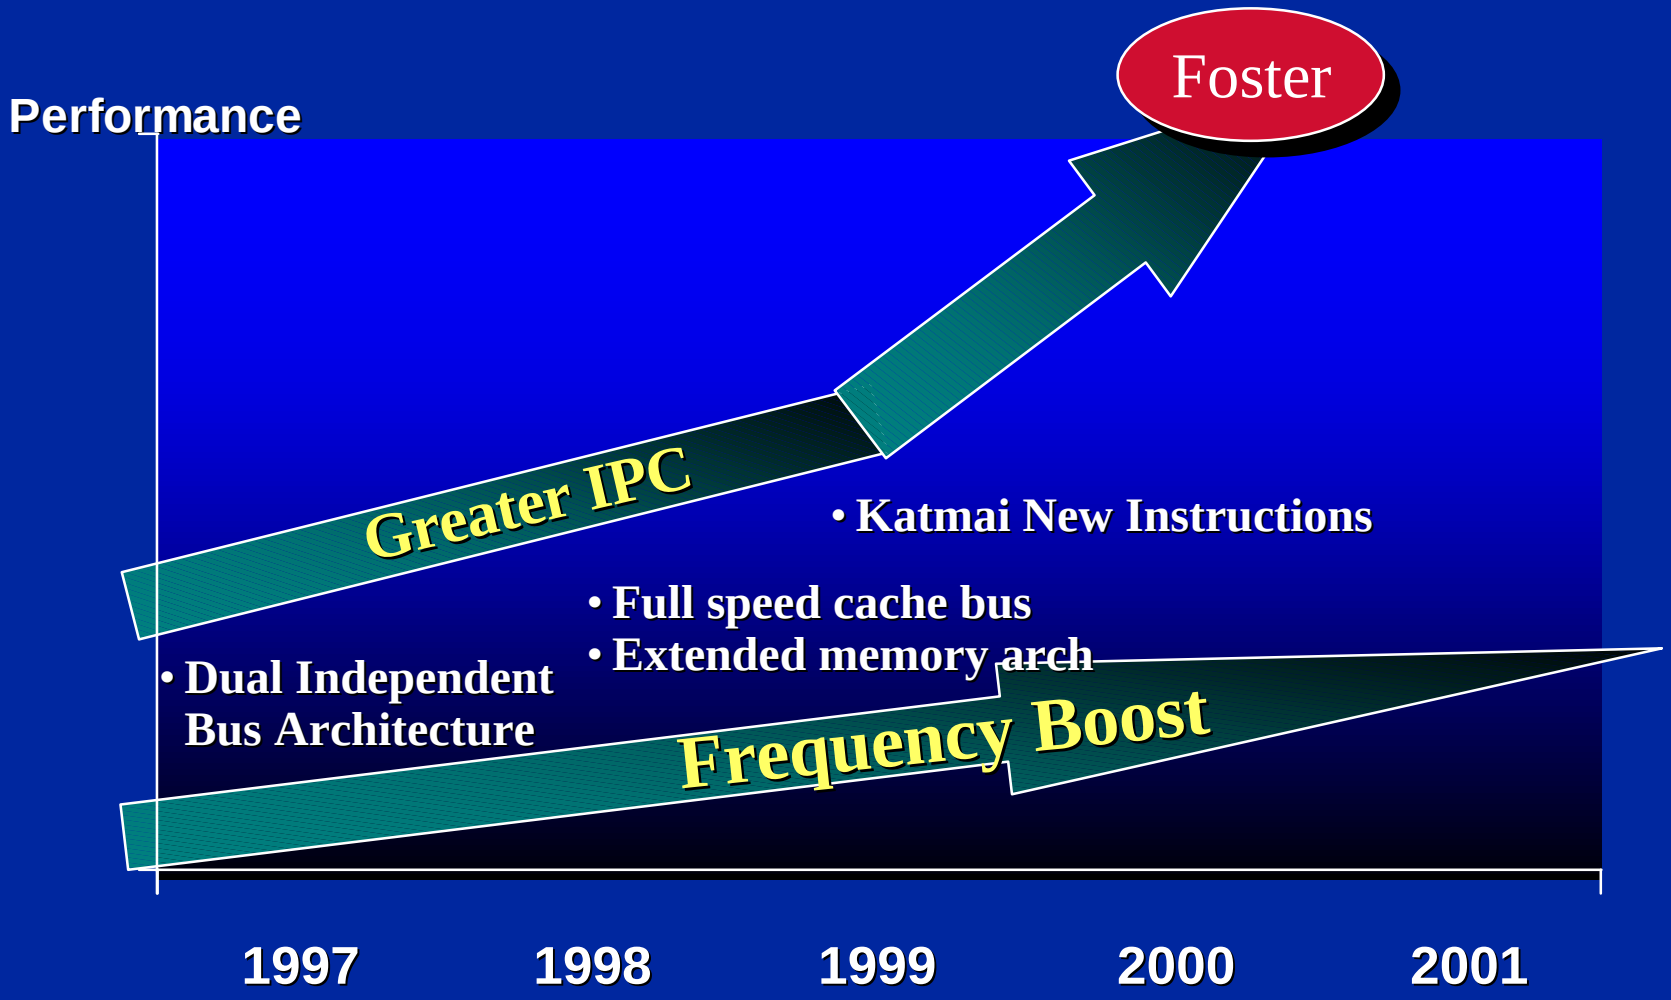

intel®

Continuous IA Innovations

New 32-bit microarchitecture
- –Implements trace cache for instruction decode
- –Enhances branch prediction

Faster frequency targeting 1GHz and beyond

Foster

Large on-chip L1 and L2 cache

Improved system throughput
- –Bus bandwidth of 3.2 GB/sec
- –Cache bandwidth increases
  - –L1 at 32 GB/sec, L2 at 8 GB/sec

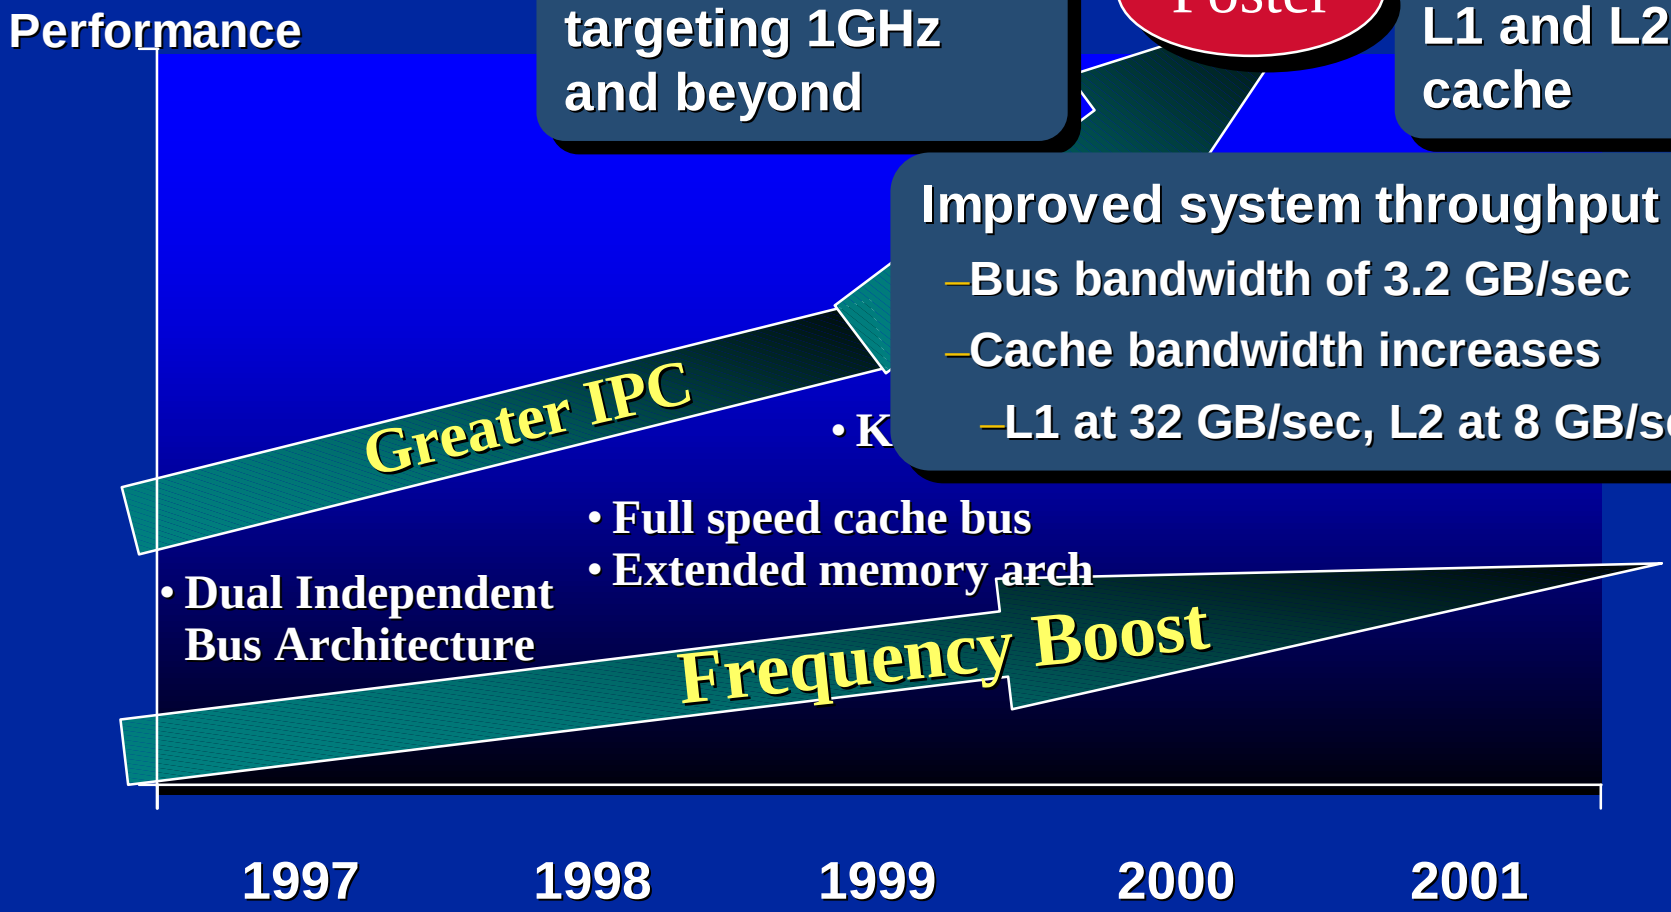

intel

Merced™ Processor FPU

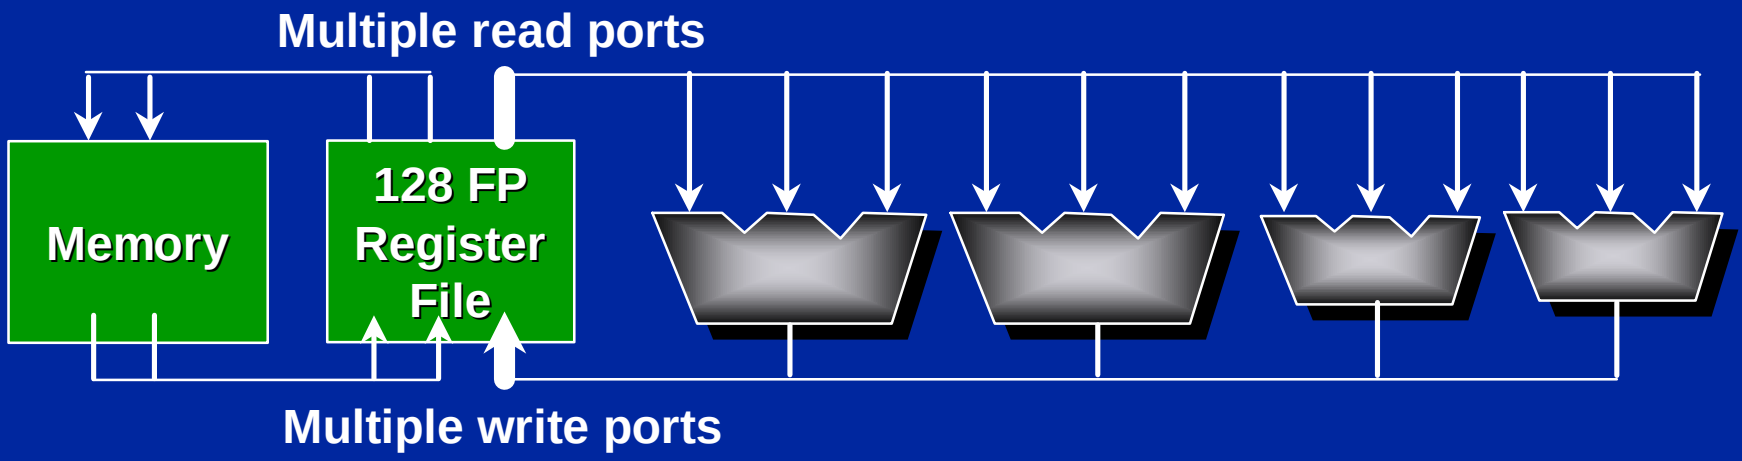

- 2 Extended Precision (EP) FMACs, 2 SP FMACs
  - Execution of up to 8 SP FLOPs / cycle
  - 4 EP FLOPs / cycle
- > 20x Pentium® Pro processor and ~3x Tanner performance on 3D graphics

intel®

IA-32 Hardware Execution

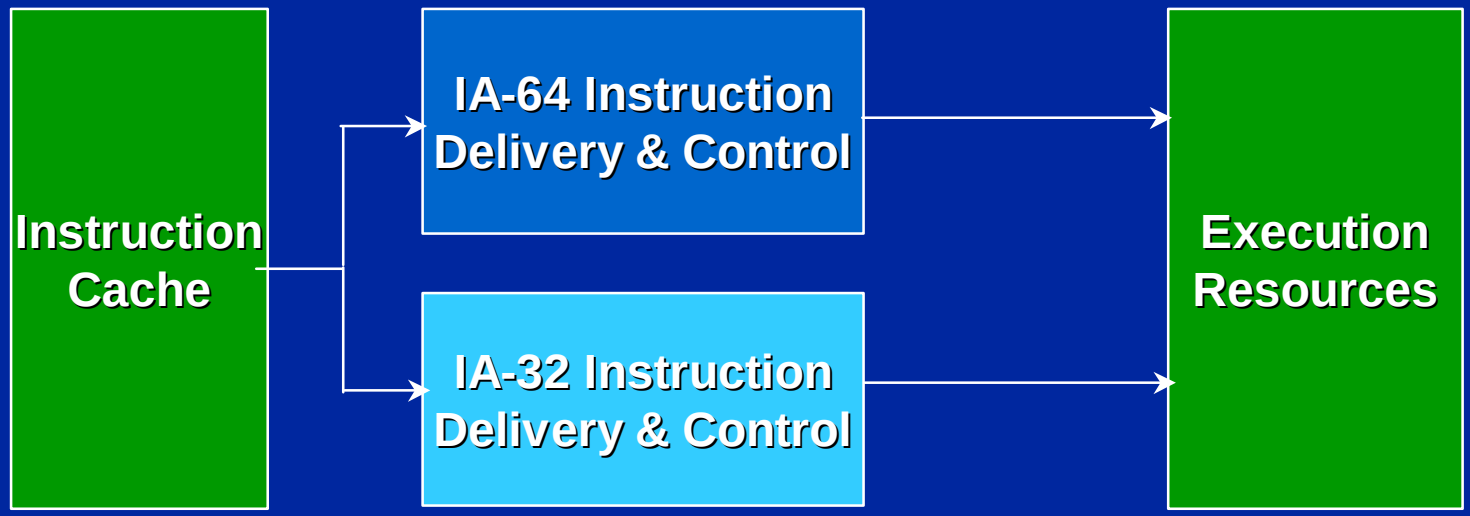

IA-32 Engine:
- IA-32 Instruction set decoder
- Dynamic execution

Shared resources:
- ALUs
- Registers
- Data cache

The only 64-bit processor with complete IA-32 binary compatibility

intel®

Merced™ Processor Manages Memory Latency

- Innovative three level cache hierarchy
  - Separate instruction & data L0 caches
  - Larger, unified L1 cache on die
  - L2 off die provides large overall capacity
- Highly efficient bus and memory utilization
  - Enhanced deferred transaction support
  - Cache line size optimized to conserve bandwidth
  - Dedicated, full speed L2 bus frees system bus for MP
  - Increased page size up to 256MB

intel®

Merced™ Processor Error Handling

- Extensive ECC coverage on processor and bus
  - L1 cache, L2 cache, L2 bus, system bus data
  - Full hardware support for correcting single bit ECC errors
- Enhanced machine check architecture
  - Processor and platform error correction via HW/ FW handshake and OS
  - Data poisoning provides greater system availability through process level error containment
- Comprehensive error logging
  - Error type, cache level, cache tag/data, corrected errors, transaction type, etc.

Supplements numerous processor, platform and OS features for enterprise-class RAS

intel®

Merced™ Processor Progress

- Microarchitecture definition complete
- Final stages of functional RTL validation-booting OS kernel
- Timing convergence exceeding goals
- Validation of circuit design making good progress
- Comprehensive pre-silicon MP validation using thorough RTL co-simulation environment

On-track for samples targeted for mid-1999

intel®

Merced™ Cartridge Preview

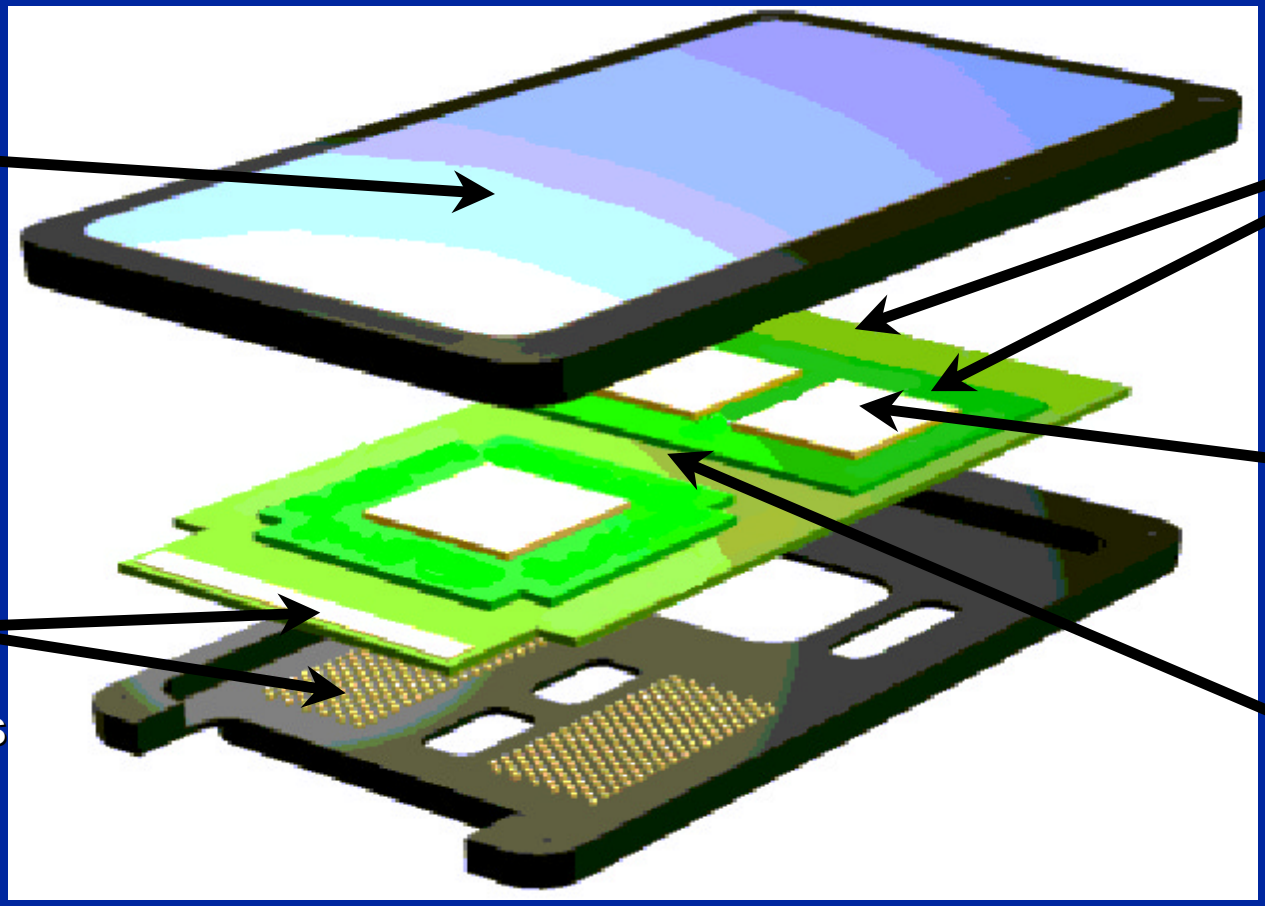

Efficient heat dissipation technology

Cost effective, performance substrate

Intel designed static cache RAM

Separate signal & power connections for signal integrity

Full speed cache bus

intel®

Merced™ Platform Program

- Key server & workstation vendors with multiple designs
  - Fault tolerant, massively parallel and technical computing designs
  - 4 to 512 MP servers and 2/4 MP workstations
- Multiple OSes making good progress: UNIX and NT
  - HP-UX, Solaris, SCO, SGI IRIX, Digital Unix, Novell Modesto, Win64
- Intel and industry shipping 64 bit SDKs and pre-silicon software development tools
- Top ISVs porting server and workstation applications
- Executing on Plan
  - Compiler optimization meeting key milestones
  - Multiple IA-64 OSes and apps booting on Merced simulator
  - Chipsets and systems designs on track for first samples

Industry converging on IA to reap common hardware foundation benefits

intel®

McKinley Processor

- McKinley extends Merced™ processor technology
  - Enhanced microarchitecture doubles IA-64 software performance
    - Frequency : Target > 1 GHz
    - IPC : Increased number of execution units
    - Very large, high speed on chip caches
  - Bus is superset of Merced bus: ~3X bus bandwidth
- Full Merced & IA-32 software compatibility
- Target production : Late '01

McKinley extends Merced processor benefits

intel®

Summary

- High performance IA server and workstation roadmap
  - Foster delivers outstanding performance for 32- bit apps
  - Merced™ processor adds 64 bits, increased headroom/ scalability
  - McKinley extends Merced processor benefits with 2X performance
  - Future IA-64 proliferations planned for .13∓ technology
- Merced processor is on track for mid 2000 production
  - Systems, OS, Applications, Tools aligned
- Common IA foundation brings greater choice to high performance segments
  - Variety of hardware, software and channel choices

IA is the Unifying Architecture

intel®