

HIGH-END PACKET SWITCH FABRIC FOR METRO AND ENTERPRISE CORE NETWORKING SYSTEMS

KEY ARCHITECTURAL ADVANTAGES

- **Global Arbitration**
 - Patented global arbiter provides the best scheduling performance; the scheduler has global state information for every queue and port in the system.
- **Virtual Output Queuing (VOQ)**
 - Supports up to 16K Virtual Output Queues for efficient scheduling.
- **Single-Stage Buffering**
 - Single-stage buffering on ingress only.
 - No packet reassembly or reordering required.
- **Up to 16 Classes of Service (COS)**
 - Support for up to 16 AF classes of service.
 - Three levels of drop precedence per COS.
- **Flow-Based Queuing on the Egress**
 - Flow-based queuing on the egress with support of up to 16K virtual ports.
 - Up to eight levels of hierarchy.

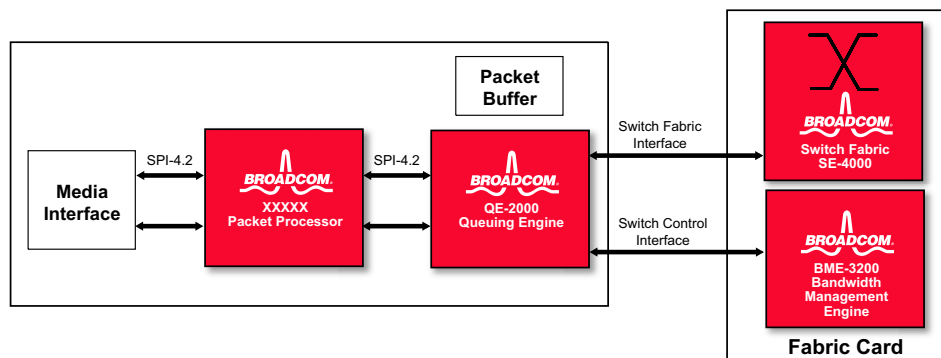
SUMMARY OF BENEFITS

- **Scalable Architecture**
 - Fabric scales from 20 Gbps to 1.28 Tbps.
- **Non-Blocking Performance**
 - Line-rate operation under all packet sizes and traffic patterns for unicast and multicast traffic.
 - No head of line blocking.
- **Advanced Quality of Service (QoS)**
 - Supports Service Level Agreements (SLAs) and bandwidth guarantees.
 - Maintains priority and classes of service through the fabric.
 - Egress queuing per virtual port.
- **Reliability**
 - Fault tolerance and redundancy features.
 - Error detection and correction mechanisms.
 - ECC memory protection in all devices in silicon, and CRC protection at the system level.

APPLICATIONS

- Metro-Ethernet Core and Edge Switch-Routers
- Enterprise Core and Wiring Closet Switch-Routers
- IP Routers
- Multi-Service Provisioning Platforms

Sandburst Fabric Block Diagram



Ultra-Scalability

The Sandburst[®] switch fabric is a highly scalable, non-blocking packet fabric that scales from 20 Gbps to 1.28 Tbps with no performance degradation, regardless of QoS profiles or traffic mix. Designed to meet the rigorous requirements of both the Enterprise Core and Metro-Ethernet markets, the switch fabric delivers bandwidth guarantees, advanced traffic engineering, and per-flow queuing across its extremely low-latency and highly-resilient fabric.

Optimal Traffic Management

The switch fabric delivers these services using a patented Virtual Output Queued (VoQ) architecture. The architecture supports up to 16K queues to enable over 16 classes of service per physical port or physical queuing per customer for service provider applications. The switch fabric optimally schedules traffic across its fabric using a global view of bandwidth requirements. Its patented global scheduling algorithms enable support of highly-efficient and granular bandwidth guarantees, and accurate prioritization of flows through the fabric.

Failover Schemes

The Sandburst switch fabric offers a rich selection of hardware failover schemes for its system OEM customers. The fabric supports 1+1, and N:N graceful redundancy schemes with sub-50 ms failover for control and data traffic types.

Advanced QoS

All of the switch fabric devices (see column to the right) provide sophisticated mechanisms for QoS management through the fabric and on the egress, a highly desired feature in the Metro-Ethernet Networking market. The devices fully support the IETF DiffServ Per Hop Behaviour (PHB) traffic management scheme supporting all the IETF-defined PHBs; Assured Forwarding (RFC 2597), Expedited Forwarding (RFC 2598), and Best Effort. In addition, packets can be queued per customer on the egress and hierarchically managed. The QE-2000 supports up to eight levels of hierarchy.

Additional Egress Queuing Functionality

Systems designers have the flexibility of using the egress queuing functionality in the QE-2000 along with Virtual Output Queuing. This is achieved by using the QE-2000 in mixed mode and dynamically partitioning memory buffers into VoQs and egress queues. In egress mode, the entire 16K queues can be used for hierarchical egress queuing.

Broadcom[®], the pulse logo, **Connecting everything**[®], Sandburst[®], and the Connecting everything logo are among the trademarks of Broadcom Corporation and/or its affiliates in the United States, certain other countries and/or the EU. Any other trademarks or trade names mentioned are the property of their respective owners.

Connecting
everything[®]

BROADCOM CORPORATION

5300 California Avenue
Irvine, California 92617

© 2008 by BROADCOM CORPORATION. All rights reserved.

Sandburst_Switch_Fabric-PB01-R 02/20/08

Silicon Devices

Queuing Engine (QE-2000)

- 20 Gbps full-duplex traffic manager
- Manages up to 16K queues
- Supports up to 16 Classes of Service (CoS)
- Deep ingress buffers for Round Trip Time (RTT) and QoS—using up to 512 MB of external commodity DRAMs
- Strict Priority (SP) and Weighted Fair Queuing (WFQ) scheduling
- Fully supports the IETF DiffServ model
- BE, EF, AF IETF PHB
- Per-flow egress queuing with up to eight levels of heirarchy
- 2 x SPI-4.2 (OIF) 48/64 channels per SPI-4.2 interface
- PCI—32-bit, 66 MHz

Bandwidth Management Engine (BME-3200)

The BME-3200 is the second generation global bandwidth managerelement of the Scalable Ethernet Switching product line. The BME-3200 connects to each of the QE-2000s in a system, offering highly reliable N+1 operation. The BME-3200 enables a highly scalable, non-blocking packet fabric thaqt scales from 20 Gbps to 1.28 Tbps with no performance degradation, regardless of QoS profiles or traffic mix.

The BME-3200 is designed to operate and meet the vigorous requirements of both the Enterprise Core and Metro-Ethernet markets. In tandem with the QE-2000 and SE-4000 it forms part of the switch fabric which delivers bandwidth guarantees, advanced traffic engineering, and per-flow queuing across its extremely low-latency and highly-resilient fabric.

- Patented VOQ architecture
- Administers IP SLAs within the packet fabric
- Arbitrates up to 32 QE-2000s using a dedicated out-of-band Bandwidth Management Interface to communicate with each QE-2000
- Dynamically allocates bandwidth to queues to provide bandwidth guarantees
- Globally tracks VOQ lengths across all QE-2000s
- Integrated WRED engine that calculates multiple-drop probabilities for each CoS in the system

Switching Engine (SE-4000)

- 40 x 40 cross-point switch
- Over 40 in. repeaterless operation of FR4
- Self-routing cross-connect
- Spatial multicasting
- Time tolerance architecture with Plesio-synchronous interfaces, reducing clock complexity and cost
- Hardware-based link and system monitoring



Phone: 949-926-5000

Fax: 949-926-7761

E-mail: info@broadcom.com

Web: www.broadcom.com