

Carrier-Class Edge Routing Required for Next-Generation Cable Services

By deploying high-performance, carrier-class, intelligent routers at the edge of the network, operators can efficiently manage bandwidth while delivering next-generation services over cable infrastructure. They can develop effective partnerships with revenue-sharing partners and can scale new services as they become increasingly popular.

Multiservice cable networks require carrier-class edge routers with high levels of intelligence so that all traffic flows can be efficiently classified and treated according to network policies. Operators and their partners will increasingly deliver diverse services over cable networks, and they can prepare for growing demand by ensuring that the network is architected for scalability and flexibility.

The term “edge router” can be interpreted differently depending on perspective. The network edge perceived by a terabit router or optical switch is not the same edge seen by a Digital Subscriber Line Access Multiplexer (DSLAM) or a cable router.

This white paper considers the edge router to be the transition point from the Hybrid Fiber Coax (HFC) access network to the regional backbone. The key feature of this position in the network is that the edge router is the first trusted device in the network and must therefore have the intelligence to implement traffic classification, management and policing. Network infrastructure



equipment in the core – whether optical or routed – need not provide these functions and can concentrate on processing large numbers of packets with simplified per-packet logic.

The edge router is the first trusted device in the network and must therefore have the intelligence to implement traffic classification, management and policing.

With carrier-class routing, cable operators can deploy increased intelligence at the edge of the HFC network to make the edge network core-agnostic. It does not have to matter if the core is based on hierarchical routing or on a flat, optical network, and operators therefore gain the flexibility to evolve the core network as newer technologies are introduced.



Carrier-Class Edge Routing

Traffic flows must be classified at extremely rapid rates at the edge of the network to streamline bandwidth efficiency, enable value-added applications and services, and prevent future bottlenecks. Enhanced services such as Voice over IP, interactive gaming, and multimedia applications are in their infancy but are expected to grow rapidly. Operators need to prepare for the onslaught of new services over the cable network so they can capitalize on these revenue opportunities.

Edge Router

► The Role of the Edge Router

The routers at the edge of the network identify and classify traffic flows and need to provide per-flow treatments according to network policies. After treating the flows, the router must efficiently forward the traffic to the appropriate destination. Traffic treatments include applying the appropriate Quality of Service (QoS) controls as well as implementing Admission Control and other traditional router services.

Routers are highly-intelligent devices that are protocol sensitive and operate at the three lower layers of the OSI model, using the Physical, Link and Network Layers to provide addressing and switching. For edge routers to be effective, they also need to offer Layer 4 routing to ensure application-aware, end-to-end transport and support advanced-load balancing to ensure the optimization of network infrastructure assets. Routers forward traffic based on pre-programmed routing considerations such as the destination address, packet priority level or route congestion level. They can contain traffic within a subnet, forward it to the core network or forward traffic to the backbone network of third-party partners.

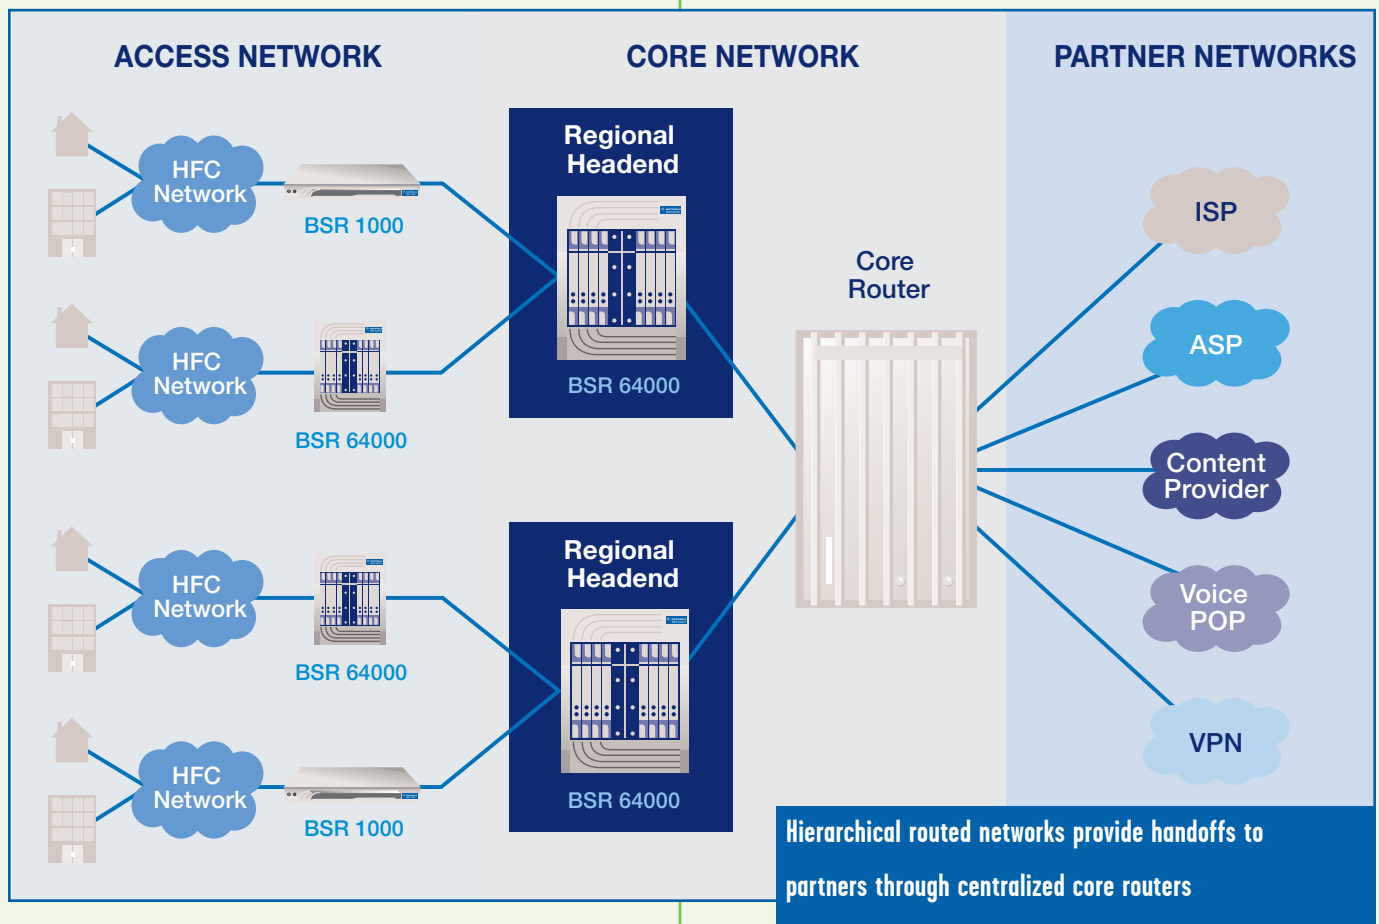
Routing to the Core

The granularity with which edge routers can identify, classify and treat packets varies. Multiple System Operators (MSOs) increasingly need to map traffic flows across both

the cable network and the core network of one or more providers. This requires high-performance packet classification at the edge of the network with packets for the core network marked according to MultiProtocol Label Switching (MPLS) and/or Diff-Serv standards. This allows operators to provide end-to-end QoS treatments across their core networks and across the backbone networks of their revenue-sharing partners. Whether the traffic is being routed over the HFC network, the operator's backbone, or the core networks of multiple revenue-sharing partners, the key to successful implementation is to add intelligence to the edge of the network. Operators also need advanced traffic engineering features using protocols such as the Resource Reservation Protocol (RSVP) to enable efficient end-to-end routing.

Eliminating the “Trombone Effect”

By implementing increased performance and intelligence at the edge of the network, carriers can develop new services and better manage demands on backbone bandwidth. They can segment traffic at the edge of the network and contain traffic within local domains. Some cable networks are architected with limited intelligence at the edge of the network, with all traffic being routed to the network core for treatment. These hierarchical topologies can create massive bottlenecks as the network scales to accommodate new services and subscribers. Under this centralized topology, the intelligence resides at the core of the



network and traffic is unnecessarily carried back to the core for classification when it could have been more efficiently routed at the network edge.

For example, a customer conducting a large file transfer to a business down the street could access the cable network through the distribution hub and have the traffic flow sent to the core router for processing – so it could be then shipped back over the cable network to the originating router for delivery to the nearby company. This process of shipping traffic back to a central site is often referred to as the “trombone effect”, since even local traffic traverses the backbone before being returned to the edge of the network. With intelligent edge routers, carriers can contain

local traffic and establish the optimum routing path according to network policies.

Bringing Content and Applications to the Customer

Carrier-class routers at the edge of the network present greater opportunity to reduce backbone congestion by distributing content and application servers closer to subscribers. They can be deployed in the distribution hub to improve the user experience and reduce the traffic traversing the core. Operators can even charge premium pricing for readily-available, high-value content, and they can cost-effectively offer high-bandwidth services such as MPEG video or interactive multimedia gaming applications.

Carrier-Class Edge Routing

Enabling Multiprovider Applications and Services

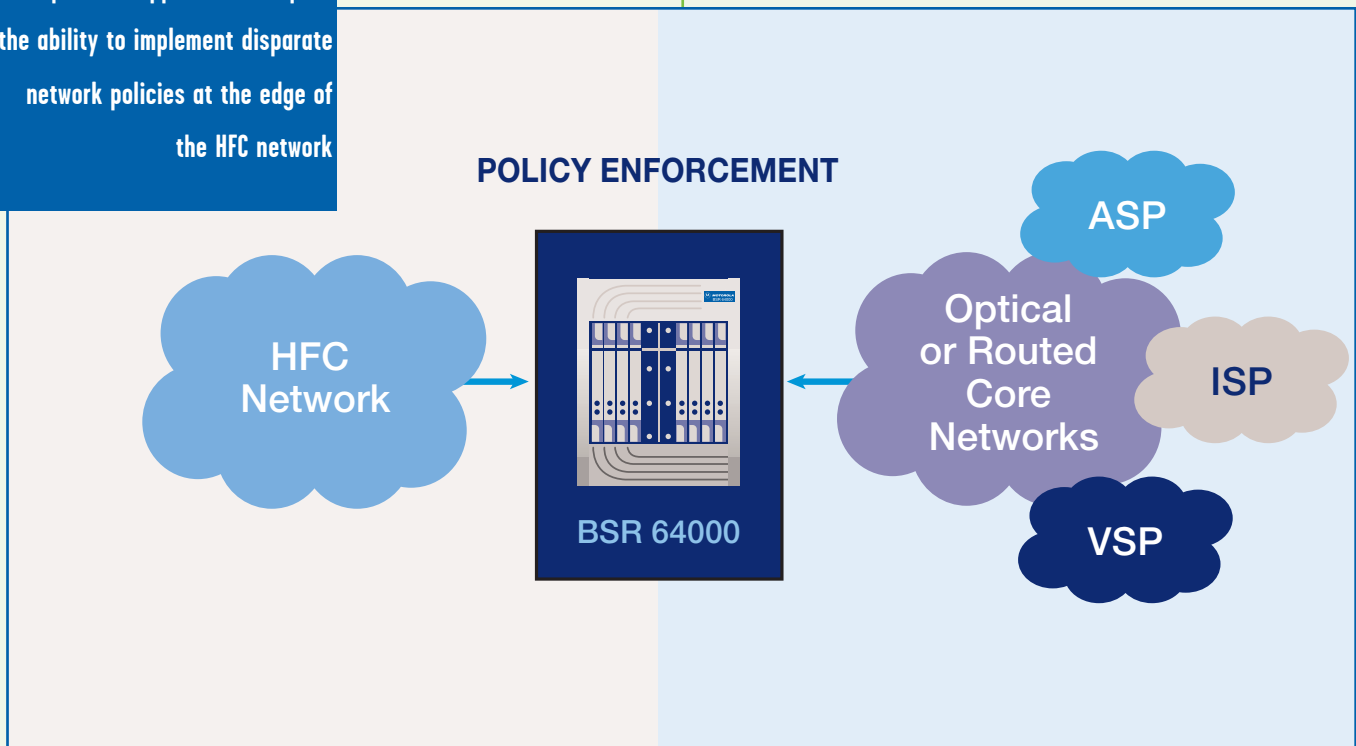
The edge router becomes even more critical in multiservice applications, since the operator has the opportunity to efficiently treat each flow at the edge of the network to eliminate costly and unnecessary routing of all traffic back to one or more central locations. Operators need to be able to efficiently route traffic within their own HFC network and also route traffic to the core networks of multiple providers in revenue-sharing applications. This requires high-levels of reliability as well as sophisticated routing intelligence so that traffic can be routed inter-domain between separate networks so operators can establish peer-to-peer relationships with partners.

Multiprovider routing requires robust, standards-based support, since operators have to be prepared for diverse network topologies and protocols used on third-party

networks. No operator knows what the next “killer app” will be, but they have to make sure they have a standards-based infrastructure that can ensure compatibility with different providers of content, services, and applications. By implementing standards-based routing at the edge of the network, carriers can protect investments in network equipment while preparing for Open Access and creating new revenue opportunities through wholesale partnerships.

Revenue-sharing partners will want varying levels of control over routing services. For example, voice carriers will want carrier-class reliability and control at the edge of the network to ensure the success of their services while a specialized Application Service Provider (ASP) may prefer to have the operator perform these functions. The deployment of carrier-class edge routers provides the maximum investment protection to support evolving network requirements.

Multiprovider applications require the ability to implement disparate network policies at the edge of the HFC network



Edge Router Requirements

Management and Control

Operators can allow multiple providers to deliver services independently over the shared HFC infrastructure. This requires advanced QoS, isolation and policing capabilities at the network edge that enable MSOs to deliver flexible, measurable, and enforceable SLAs to both providers and subscribers, and allow the delivery of real-time services over the HFC network from multiple sources. For example, a subscriber could obtain voice services from an Inter-Exchange carrier, VPN services from a competitive local exchange carrier, business applications from an ASP and Internet access services from the cable operator.

Sophisticated service creation and advanced partitioning features are required to support multiple providers delivering services over the shared cable access network. Advanced network management control is essential so that revenue-sharing partners can create and manage new services while gaining real-time insights into network performance and the impact of services on their allotted shares of the HFC access network resources. The ability to allow partitioned management views is critical so that third-party providers can optimize their use of the shared cable access network. They also need management control over service creation so they can create service profiles and apply them to subscribers to bulk provision new services.

Partitioned management allows operators to effectively manage bandwidth and deploy services while delivering the segmentation and isolation necessary to ensure that all providers operate within the MSO-established bounds and prevent any provider from usurping additional bandwidth and resources. Operators can offer residential, business-class and wholesale services to maximize revenue. They can establish and enforce network policies at the edge of the network for admission control and rate limitations to ensure that all SLAs can be met while optimizing revenue from available network capacity.

Carrier-Class Edge Router Requirements

Once a carrier establishes relationships with revenue-sharing partners, the edge router becomes a shared device that must support multiple network environments. In broadband networks, the sheer volume of traffic requires carrier-class levels of reliability. The failure of a router at the edge of the network results in lost revenue and customers and can even result in the operator being held liable for violating SLAs with revenue-sharing partners – such as a VoIP carrier.

Redundancy

Operators – and their subscribers and revenue sharing partners – cannot tolerate failures at the edge of the network, so carrier-class routers are required to ensure non-stop operations. The term “carrier-class” is often quantified as delivering Five Nines of Reliability – 99.999 percent reliability, or less than six minutes of unscheduled downtime per year. This requires redundancy for all critical system elements and compliance with NEBS standards.

Performance

Edge routers need the ability to instantly recognize source traffic flows and match the traffic to the appropriate provider in real-time, and they need advanced capabilities to police traffic flows and ensure that no provider consumes network resources that exceed agreed-upon specifications. Open Access support requires that routers make complex forwarding decisions based on multiple fields in the IP packet header rather than simply using the destination address. Hardware-based packet-processing is required to rapidly classify traffic at the edge of the network for transmission across HFC infrastructure as well as the core backbone of multiple providers.

Carrier-Class Edge Routing

QoS

Best-effort service is adequate for standard Internet access offered by ISPs but is insufficient for most enhanced services because operators need the ability to offer different QoS levels for specific subscribers – and for specific services. Third-party providers must be able to establish SLAs with cable operators that document the resources required for each of many potential services. This requires the ability to create and enforce a hierarchy of nested QoS domains within the HFC infrastructure – which in turn requires sophisticated, high-performance packet filtering and forwarding. It also requires the ability to support end-to-end QoS guarantees across both HFC and third-party networks using industry standards such as MPLS and Diff-Serv.

Operators need to support advanced SLA parameters such as maximum bandwidth allocation, minimum bandwidth guarantees, bounded delays and bounded jitter. They will need the ability to define QoS parameters both statically (e.g. Gold/Silver/Bronze services) and dynamically (e.g. for services such as voice call set-up). At a minimum, operators need the QoS capabilities of Data Over Cable Service Interface Specification (DOCSIS) 1.1-based equipment, but they also need features beyond these standards to enable enhanced services over both HFC and service provider networks.

Isolation

To provide subscribers and third-party providers with predictable levels of service, it is essential that traffic flows be contained at each level of the QoS hierarchies. Overload or misbehavior within the HFC network by any given provider must be contained within the network resources committed to that service provider – and not be allowed to impact other providers sharing the network.

Isolation is needed to prevent unscrupulous or naive providers from massively overselling their service to the detriment of all other providers on the HFC network. The edge router must allow each service provider to isolate each of its subscribers so that none can impact other subscribers sharing a common domain. In addition, any overload or misbehavior within a subscriber service should be isolated to that particular service. For example, a CLEC offering Internet access and voice services must be able to prevent a subscriber's Web traffic from impacting his or her voice calls.

Policing

Policing of traffic flows is required to ensure the necessary isolation and enable SLA enforcement. The edge router needs to police traffic flows to make sure that each service provider is compliant with documented SLA parameters. The operators need the flexibility to ensure that knowledgeable users do not take advantage of the network QoS mechanisms to obtain services for which they have not paid. Traffic that exceeds SLAs should be handled according to SLA policies that determine whether excess flows should be dropped, assigned lower priority levels or routed at incremental costs.

Accounting and Metering

Allowing multiple service providers to operate over a shared access network requires robust features for reconciliation and billing, and the per-flow classification and routing of traffic at the edge of the network. Granular observability with detailed metering information needs to be captured by the edge router at the per-flow level to ensure that SLAs are enforced, and the sophistication and complexity of accounting can vary dramatically. In the simplest case, a provider could define an SLA and the MSO could implement a policing mechanism to ensure that it is

not exceeded. However, in most applications both the provider and operator will want to meter the SLA to ensure conformance.

Management and Control

In a multiple service provider scenario multiple partners will require management access to the next-generation CMTS/edge router platform. The ability to provide partitioned management via SNMP and a Command Line Interface (CLI) is important to most revenue-sharing partners so that they can gain insights into their services across the network. Providing CMTS and router functionality in a single system further simplifies management operations by eliminating the need to manage additional hardware.

Scalability

The edge router must be scalable to support new demands while maintaining the carrier-class redundancy required for critical services. Centralized routing calculations with distributed forwarding in hardware provide the ideal solution for simple configuration and maximum scalability. The routing software implementations must be similarly scalable. It is not enough for a vendor to proclaim support for leading protocols – such as Open Shortest Path First (OSPF) for intra-domain routing or Border Gateway Protocol 4 (BGP4) for inter-domain routing – the operator must carefully evaluate the software to be sure it is a carrier-grade implementation that can scale to meet future requirements as the network scales in response to service requirements not dreamed of today.

Security

As operators scale their networks to offer new services delivered by third-party partners, security increasingly becomes a major concern. The edge router must be able to

establish secure peering relationships that protect the HFC network and also protect misbehaving partners from harming traffic flows destined for other providers. This requires not only secure and scalable routing implementations, but also the hardware-based processing needed to support large amounts of filters that can be configured to ensure network security.

Service Creation

Swift and easy service creation is essential for accelerating the introduction of revenue-generating services. Tools are required that allow operators – and possibly their partners – to instantiate commands into edge routers to enable the rapid creation of services according to business policies – without the need for router-by-router reconfigurations. This also requires compatibility with back-end Operations Support System (OSS) infrastructure and intelligence at the edge of the network to enable automated service creation.

Policy-Based Routing at the Edge with the BSR 64000

The award-winning Broadband Services Router 64000 (BSR 64000) from Motorola provides broadband carriers with a competitive advantage in defining, deploying and managing broadband services. This next-generation CMTS/edge router allows cable operators to rapidly introduce differentiated data, voice and multimedia services for both corporate and residential subscribers and deliver QoS levels end-to-end from the network edge.

The SmartFlow™ features of the BSR 64000 system perform content-aware packet classification through Layer 4 to provide unprecedented QoS flexibility. Since all the processing-intensive filtering, forwarding, accounting, and QoS/SLA functions are performed in hardware at wire-speed, the BSR 64000 reduces latency to a fraction of that commonly found in mainstream, software-based routers.

Carrier-Class Edge Routing

The BSR 64000 is DOCSIS 1.0-qualified and compatible with DOCSIS 1.1 and PacketCable 1.0 standards. It can support up to 26 downstream transmitters and up to 104 upstream receivers in a single, space-saving chassis.

The BSR 64000 can also be deployed to integrate and extend QoS and routing functionality to legacy, first-generation DOCSIS or proprietary CMTS equipment.

Policy-Based Routing and MPLS

The BSR 64000 can be deployed as a carrier-class MPLS Label Edge Router (LER) to provide support for dynamic Label Switched Path (LSP) creation. Because of the BSR 64000's hardware-based distributed forwarding architecture, MPLS traffic flows can be routed at wire-speed. Treatments for QoS are based on the policies defined by-and-for each service provider to enable end-to-end traffic treatment across access, metropolitan and core networks.

The BSR 64000 offers a carrier-class implementation of Policy-Based Routing with MPLS to allow broadband providers to enable third-party, revenue-sharing partners to deliver content, applications and services over the access network. The BSR 64000 looks at multiple fields within packets to determine the appropriate routing and QoS. Because the BSR 64000 can look at the individual application flows, it can extend the capabilities of DOCSIS 1.1. This intelligent edge router can inspect multiple fields within packets to determine the appropriate routing and packet classification requirements. Packet routing is partially determined by looking at the source IP address, understanding to which service provider partner the IP address belongs, applying the appropriate MPLS label and then routing the traffic on the appropriate LSP to that partner for their handling. This examination allows the BSR 64000 to implement more sophisticated QoS policies than are possible by simply looking at the data's destination address. The BSR 64000 can assign QoS and routing

policies based on parameters such as service provider, subscriber and application.

Carrier-Class Platform

The cost-effective BSR 64000 includes flexible interfaces for SONET and Ethernet connectivity, and it eliminates the need for discrete CMTS equipment, up-converters, aggregation switches and routers. It offers unified management of routing, QoS and CMTS functions and scales economically to meet ever-increasing subscriber demands and the introduction of new services.

The BSR 64000 design is based on centralized routing and distributed forwarding and provides the benefits of simple configuration (single router appearance), scalable performance (each additional line card brings an associated forwarding engine) and low cost-of-entry (operators only purchase the forwarding power required). It can be deployed in the distribution hub to provide an interchange point between the regional fiber network and the cable plant and in a regional headend to interconnect the regional network with a backbone network and allow connectivity to local content servers and management systems.

The system offers high levels of resiliency and redundancy, and the 16-slot, NEBS-compliant chassis has a midplane architecture that enhances serviceability by de-coupling functional hardware modules from the physical I/O and connectivity. The platform has redundant power supplies and fans in an ultra-dense and modular chassis that provides industry-leading density to make efficient use of scarce headend real estate. The BSR 64000 offers a rich feature set for carrier-class routing. Operators gain the ability to optimize the network traffic for a given topology, and they benefit from the advantages of flow-based routing over simpler destination-based routing. Operators can leverage these carrier-class routing features to better control

and manage the network traffic to optimize revenue, and they can leverage the high-performance architecture to lower the cost of packet processing.

Carrier-Class Software Implementation

The BSR 64000 offers carrier-class routing software implementations that can scale as operators expand their networks and develop new partnerships with revenue-sharing partners. The robust routing software ensures high-availability for the network and superior network reliability.

Each routing protocol implemented on the BSR 64000 offers maximum expansion to support new services, subscribers and providers. Carrier-class implementations of intra-domain, inter-domain and multicast routing are supported, including: OSPF v2, RIP v1 and v2, BGP4, IS-IS, VRRP, IGMP, DVMRP, MPLS, Diff-Serv and PIM-SM/DM. Motorola has designed a set of routing protocol and routing policy-dependent features to facilitate the task of system management and network control.

With the BSR 64000, operators benefit from carrier-class routing and can scale their networks in terms of numbers of routes, interfaces and peering relationships. This allows operators to provide high-growth, business-class services over cable networks and add more subscribers with additional pools of IP addresses – without modifying the equipment or adding hardware. The BSR 64000 provides maximum flexibility in the number of network nodes that can be connected, which provide increased traffic management control and topology independence.

Cable operators can implement efficient and scalable peering arrangements with third-party providers using standard routing protocols. They can “future-proof” their infrastructure by deploying a platform that supports all the major routing protocols. This will allow operators to

effectively route traffic end-to-end across the cable access network and the core networks of multiple providers – each potentially using different routing protocols.

Operators can traffic engineer end-to-end connections using RSVP, and paths between MPLS Label Edge Routers (LERs) can also be evaluated using the MPLS Label Distribution Protocol (LDP) to determine whether the path has the appropriate network characteristics to support the desired traffic type. Operators can therefore assure the performance required to support the routing requirements of diverse traffic flows.

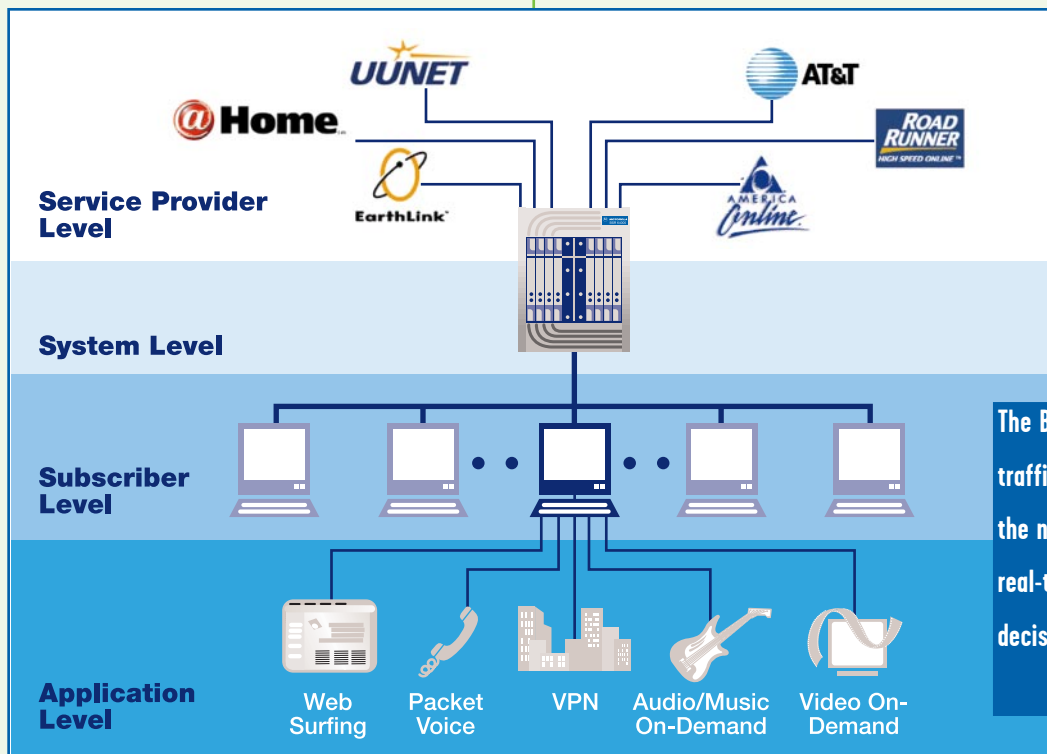
Carrier-Class Reliability

The BSR 64000 is architected for 99.999 percent reliability and provides the fault detection and switchover required for high-availability services. It includes Automatic Protection Switching (APS) and redundant controller and routing modules. The carrier-class BSR 64000 supports all of the traditional Central Office (CO) operational requirements – such as minimal disruptive software upgrades in redundant pairs, live insertion, version roll-back, integration into the alarming scheme and full NEBS compliance. Extensive fault detection and reporting features let operators optimize system performance and deliver primary-line IP telephony reliability.

Robust and Secure Multiprovider Support

Powerful partitioned management features allow third-party providers granular views of services, equipment and network resources. Robust and broad routing protocol implementations allow operators to effectively develop peering relationships so that the shared HFC access network can be used to deliver dedicated services to subscribers. Within the constraints imposed by the MSO, each service provider can deliver services based on individual policies and business practices. Cable operators

Carrier-Class Edge Routing



can isolate traffic flows so that any overflows or misbehaviors are contained within each service provider's allotted resources to ensure that no service provider can negatively impact network resources committed to any other provider.

Advanced routing security features support multiprovider security. Scalable routing software implementations allow hundreds of secure BGP4 peering sessions so that operators can successfully route traffic onto the core networks of third-party partners – without allowing any misbehaving partner to impact traffic destined for other providers. The BSR 64000 offers stable code even under malicious attacks and provides robust filtering mechanisms such as address screening and rate limiting to protect routers. Unlike conventional edge routers that may offer dozens of filters, the BSR 64000 offers thousands of filters and/or classifiers that can be implemented concurrently without degrading performance. It also offers security and

authentication extensions for the routing protocols implemented and the BSR 64000 enables secure, encrypted tunneling using the IPsec protocol for establishing Virtual Private Network (VPN) connections.

Flexible and Powerful Isolation and Policing

The BSR 64000 with SmartFlow allows per-flow policing and traffic shaping at wire-speed, thus enabling MSOs to provide SLAs to subscribers on either a short-term or long-term basis. High-performance, content-aware packet classification and forwarding at wire-speed provides unmatched abilities to offer customized QoS levels and guaranteed SLAs. Among the QoS parameters that can be included in SLAs are Constant Bit Rate (CBR); maximum bandwidth allocation for sustained traffic flows or bursts; minimum bandwidth guarantees with measurable Committed Information Rates (CIRs); bounded delays; and minimal packet loss guarantees.

Traceability

The BSR 64000 software provides tracing options to facilitate software as well as network-level debugging. Operators can debug the network while it is operational, and they can quickly identify and resolve any configuration problems, even if they are caused by peering arrangements with revenue-sharing partners. The BSR 64000 offers the ability to quickly identify and fix inconsistent network configurations without interrupting network operation. The tracing options can be configurable at many levels, and all logins and configuration changes are logged. All the network interface and flow level statistics are saved periodically in non-volatile memory, and any security violations are automatically logged to support network auditing.

High Operational Reliability

Each task domain in the BSR 64000 runs in its own protected memory space using the hardware Memory Management Unit with minimal impact on its performance. This ensures that the failure of one of the tasks does not negatively impact the operation of other tasks. Between these independent tasks, communication takes place via well-defined interfaces that provide inter-process communication, thus resulting in a very reliable software system. Similar type of protection is provided for the shared data to prevent accidental corruption and prevent partners from sending invalid routes.

Security

The number of security threats that providers and operators face is growing constantly, and any security violation for the control traffic has serious repercussions for operators, providers, and subscribers. To reduce exposure to these threats, network equipment must have inherent

security precautions that safeguard the operation, service and functionality of all supported services. The BSR 64000 provides multiple levels of access for different users, an inactivity log out timer, encrypted password protection and the ability to restrict Telnet management to specified IP addresses only. The BSR 64000 supports MD5 authentication for BGP sessions, and it also supports MD5-based routing message exchange for OSPF and RIP. This protects against violators stealing service from providers by getting the provider to advertise false routes, injecting a routing backdoor in the routing table and destabilizing the network by damaging the routing message exchange.

The BSR 64000 provides management and control plane security. This architecture immunizes the network from service attacks, reduces the risk of network instability and precludes the possibility of intruders crashing the network by injecting bogus routes into the router. All configuration changes and security violations are logged to provide a complete audit trail and help operators identify any intrusion attempts.

Interoperability

The BSR 64000 is most likely to be deployed in existing networks, and Motorola has designed the system for compatibility with established vendor equipment, including routers from Cisco Systems and Juniper Networks. The interoperability of the BSR 64000 has been independently verified by DOCSIS 1.0 qualification and successful completion of the University of New Hampshire interoperability test suites for routing protocols. The BSR 64000 offers seamless integration with the existing network as well as automatic switchover to new equipment with minimal downtime.

Network Management and Service Creation

The Advanced Provisioning Manager from Motorola provides powerful element management features so that operators and their partners can optimize the use of network assets and efficiently manage the BSR 64000. It abstracts the creation of QoS-enabled services so that they can be defined at the business policy level of the operator rather than the command line interface of the CMTS/router. The Advanced Provisioning Manager allows MSOs to offer Web-based customer self-provisioning and integrate service creation with the existing OSS infrastructure.

Partitioned Management Views

The BSR 64000 offers several options for management, control and administration. In headend locations with limited availability of trained staff, troubleshooting on the BSR 64000 is simple – with easy-to-read diagnostic LEDs as well as remote management capability to support provisioning, configuration and problem identification. The BSR 64000 supports Simple Network Management Protocol

(SNMP) v1 and v3 and offers a Cisco-compatible Command Line Interface (CLI) for ease-of-use and interoperability with legacy infrastructure. Partitioned management allows each provider to view its own network management environment to control its committed resources on the HFC network.

➤ A Next-Generation Solution for the Network Edge

The BSR 64000 offers operators a full-featured, next-generation edge-routing platform with an integrated CMTS and worldclass routing functionality that can scale as the network grows. Operators can deploy a single, integrated solution at the edge of the network with robust routing capabilities and one of the highest-density, lowest price-per-port CMTS in the industry. The BSR 64000 offers carrier-class, 99.999 percent reliability so MSOs can efficiently route multiservice traffic flows at the network edge and develop efficient peering relationships with revenue-sharing partners that lead to increased revenues and profits.



MOTOROLA