MOTOROLA
intelligence everywhere™

Routing Protocol Support in the BSR 64000

February 23, 2004

White Paper

Introduction

Cable operators need carrier-class implementations of major routing protocols so they can efficiently deploy new services and integrate diverse technology across their networks. The Broadband Services Router 64000 (BSR 64000) is a fully redundant, carrier-class, intelligent edge router with an integrated, high-density Cable Modem Termination System (CMTS). It supports major IP routing protocols and this proven platform allows broadband operators to integrate new services onto existing networks using industry-standard protocols.

The BSR 64000 is a high-performance, intelligent edge router that performs sophisticated traffic grooming and forwarding in a distributed fashion at the periphery of the broadband access network to enable end-to-end service delivery across access, metropolitan, and core networks. It is Data Over Cable Service Interface Specification (DOCSIS) and EuroDOCSIS 1.1-qualified, PacketCable 1.0-qualified, and compatible with DOCSIS 2.0 and Euro-DOCSIS 2.0.

This whitepaper presents an overview of the routing architecture of the BSR 64000 platform and provides a synopsis of major routing protocols used by cable operators and how they are treated by the BSR 64000. It emphasizes the importance of reliable and scalable implementations of routing protocols and demonstrates why the BSR 64000 is an ideal platform for supporting the evolving routing requirements of broadband cable operators.

Carrier-Class, Hardware-Based Routing

The BSR 64000 offers unified management of routing, QoS, and CMTS functions, and it scales economically to meet ever-increasing subscriber demands and the introduction of new services. The BSR 64000 architecture optimizes performance and reliability by centralizing routing table calculations while distributing packet-forwarding functions. The BSR 64000 is architected for 99.999 percent availability and provides the redundancy, fault detection, and switchover required for high-availability services. Because the architecture is based on centralized routing and distributed forwarding, the BSR 64000 provides the benefits of:

- Simple configuration (single router appearance)
- Scalable performance (each additional line card brings an associated forwarding engine
- Low cost-of-entry (operators only purchase the forwarding power required)

Since all the processing-intense filtering, forwarding, accounting, and QoS 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. Route calculations are performed in Supervisory Routing Modules, (SRMs) which operators can deploy in a redundant configuration. All interface modules connect to the primary and secondary SRMs via redundant control buses that allow the flow of control and management information from the SRM to the resource modules. Routing table calculations are centralized in the SRM to optimize system performance and simplify configuration and management. The BSR 64000 offers 1:1 SRM redundancy, and the primary SRM is always synchronized with the backup.

High-speed interfaces allow connectivity to local servers so operators can efficiently distribute content, services, and applications over the access network, and the BSR 64000 eliminates the need for discrete CMTS equipment, up converters, aggregation switches, and routers. The BSR 64000 allows operators to benefit from carrier-class routing and allows them to scale their networks in terms of numbers of subscribers, routes, peering arrangements, and services.

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.

Rich Protocol Support

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. A router chooses a path by looking in a routing table. Routers like the BSR 64000 are responsible for determining the next hop for each packet. Originally, these lookup tables were manually configured in a process known as static routing but most operators today implement dynamic routing because network topologies constantly evolve.

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 statistics, flow level statistics are saved periodically in non-volatile memory, and any security violations are automatically logged to support network auditing.

Flexibility to Support Diverse Routing Requirements

Each operator's network can be viewed in routing terms as an autonomous system or routing domain. Operators implement interior routing protocols within a domain, while exterior routing protocols are used for routing among autonomous systems. The BSR 64000 supports unicast protocols for transmitting router updates to neighboring routers, and it can transmit status updates to multiple routers throughout the network using multicast protocols. The BSR 64000 supports interior, exterior, and multicast routing protocols and allows operators to classify traffic at the edge of the network for routing within the operator's domain and/or to third-party providers. With the BSR 64000, operators benefit from high-availability routing and can scale their networks in terms of number of routes, interfaces, and peering relationships.

It provides maximum flexibility in the number of network nodes that can be connected and advanced traffic management control and topology independence. Robust routing functionality at Layer 3 and above enables efficient load sharing and balancing of network resources to optimize system availability. The BSR 64000 offers carrier-class implementations of the following protocols including hitless failover and upgrades so operators can scale their networks to support ever-increasing demands while ensuring uptime:

- Routing Information Protocol (RIP) v1 and v2
- Intermediate System-Intermediate System (IS-IS)
- Open Shortest Path First (OSPF) v2
- Border Gateway Protocol 4 (BGP4)
- Multiprotocol Label Switching (MPLS)
- Internet Group Management Protocol (IGMP)
- Virtual Router Redundancy Protocol (VRRP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Protocol-Independent Multicast-Sparse Mode (PIM-SM)

The following are highlights of the BSR 64000's support for these critical protocols. (For more information on BSR 64000's carrier-class architecture for routing data, voice, and multimedia traffic, please download the whitepaper at www.motorolabroadband.com.)

Routing Information Protocol

RIP is a dynamic routing protocol primary used in interior routing environments. It is deployed primarily in earlier cable data networks and offers limited scalability potential. RIP uses a distance-vector algorithm to discover routes and exchange route information. When an operator deploys RIP on the BSR 64000, the system will advertise a route as a vector of direction and distance. With RIP, a router sends its entire routing table to the next router every 30 seconds, and each router passes along this information. RIP uses a hop count to determine network distance.

The BSR 64000s on the network using this vector information to build route lookup tables. RIP is used primarily for small, autonomous networks, but for larger broadband networks the overhead caused by the transmission of the entire routing traffic table every 30 seconds imposes an ever-increasing burden. Many operators deploy RIP between a router located on the enterprise network and a BSR 64000 in the distribution hub. RIP allows operators to easily deploy enhanced services to the enterprise. In this scenario, operators can use RIP between an enterprise router and a BSR 64000 in the distribution hub, and then use a more robust interior routing protocol like OSPF within the interior domain.

Intermediate System-Intermediate System

This link state hierarchical routing protocol allows intermediate systems (routers) to exchange information on a single network to determine network topology. It is an interior protocol for routing traffic within an autonomous domain. IS-IS is used by some cable operators for routing traffic within their networks. Like RIP, it uses a single parameter for calculating routes.

IS-IS routers flood the cable network with link-state information, and other routers receive this information and build a database that describes all the routes on the network. A routing table is then calculated from this information. Like RIP, IS-IS results in network overhead burdens from link-state transmissions. IS-IS is used primarily on cable networks comprised largely of equipment from Cisco Systems, and the BSR 64000 offers a robust implementation of IS-IS that enables standards-based compatibility with legacy routers.

Open Shortest Path First

OSPF is an interior routing protocol developed for use within a service provider's own network. As a link state routing protocol, OSPF collects information about links between routers to build a map of the network topology. A BSR 64000 on the network that obtains a change to a routing table or detects a change in the network immediately multicasts the information to all other OSPF routers in the network.

Unlike RIP, which sends the entire routing table, OSPF will send only the part that has changed. With RIP, the routing table is sent to a neighbor host every 30 seconds but OSPF multicasts the updated information only when a change has taken place. Rather than simply counting the number of hops, OSPF bases its path descriptions on link states that take into account additional network information.

OSPF also goes beyond RIP by letting the broadband operator assign cost metrics to a given router so that some paths are given preference. RIP is supported within OSPF for router-to-end station communication so smaller operators with RIP subnets can evolve to OSPF as their network scales. Because of the increased efficiency, lower overhead and the intelligence to calculate routes not only based on distance but also on customized cost metrics, operators with large or growing networks will likely prefer OSPF as the standard for interior routing. The BSR 64000 supports fast reroute capabilities to enable automatic failover between Ethernet interfaces in the rare event of a failure in the interface card.

A redundant Ethernet interface is available for each platform to enable load sharing and protection, and OSPF enables fast reroute of traffic flows between these Ethernet interfaces. OSPF is the most widely used interior routing protocol within cable operator networks, and this protocol is mature and widely supported by the vendor community. The BSR 64000 offers carrier-class implementations of OSPF v1 and v2 so that operators can create large and scalable routing domains to support their residential and enterprise subscribers.

Border Gateway Protocol

The BSR 64000 allows operators to interoperate their networks with any number of peering partners using BGP, which is an exterior routing protocol used for exchanging routing information between autonomous networks. The BSR 64000 supports BGP4 and allows operators to automatically create two routing tables—one using a protocol like RIP or OSPF for routing within the autonomous network, and one using an exterior protocol like BGP for routing traffic from the operator's network to a third-party service provider.

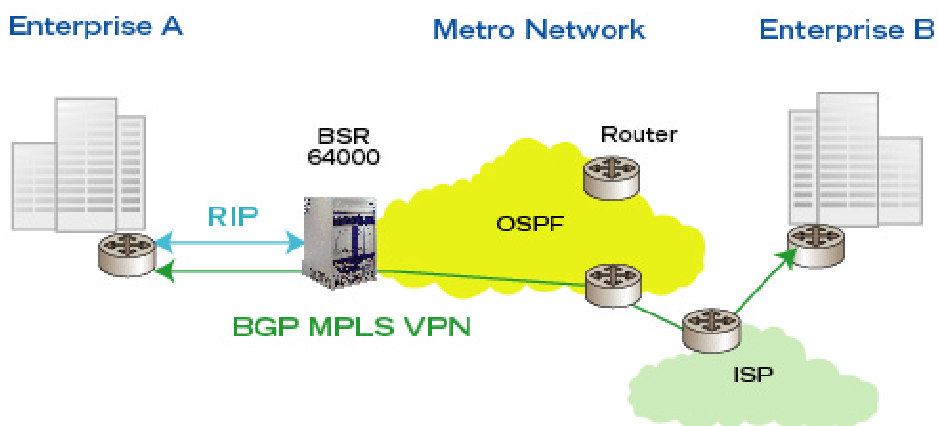

The BSR 64000 allows operators to concurrently implement multiple interior and exterior routing protocols to support diverse carrier-class routing applications.

Like OSPF, BGP allows routers to make route calculations based not only on distance but also on operator-defined cost-metrics. The BSR 64000 offers a scalable routing software implementation of BGP4 that allow operators to create hundreds of secure BGP peering sessions. They can successfully route traffic onto the core networks of third-party partners while protecting their networks from allowing any partner to impact traffic destined for other providers. The BSR 64000 also allows operators to use BGP4 in conjunction with MPLS to create flexible Virtual Private Networks (see below).

Multiprotocol Label Switching

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.

With the MPLS features of the BSR 64000, operators can also offer VPN services to commercial customers. Support for BGP/MPLS VPN (RFC 2547) allows cable operators to provide compelling VPN services for enterprise customers that encourage them to migrate from their incumbent carriers. Broadband operators can offer services with the same parameters for QoS, security, performance and at a fraction of the cost of traditional leased line or Frame Relay/ATM services offered by incumbent providers—and operators can provide enhanced value by allowing commercial customers to quickly and easily provision additional bandwidth.

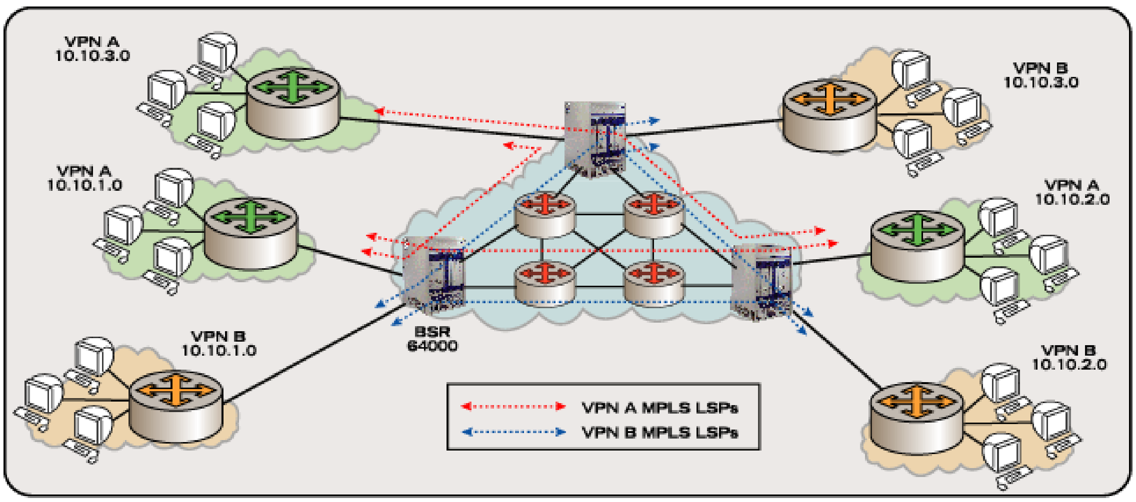

Cable operators maintain maximum flexibility for deploying VPNs. In this example, BGP/MPLS VPNs are deployed to create Label Switched Paths between VPNs while avoiding the limitations of non-routable and overlapping IP addresses.

Internet Group Management Protocol

Multicast routers must indicate their desire to be included in a multicast session. This is usually done with IGMP, a protocol that runs between immediately neighboring multicast routers. The BSR 64000 can use the protocol to inform a local multicast router whether it wants to receive transmissions from a particular multicast group.

Virtual Router Redundancy Protocol

VRRP is used for redundancy between routers. It eliminates single points of failure by supporting redundant router connections. Using VRRP, a virtual IP address can be specified manually or with Dynamic Host Configuration Protocol (DHCP). This virtual IP address is shared among the routers, and one of the routers is designated as the master router and the other router or routers are designated as backups. The master router sends a special VRRP advertisement packet to the backup routers, usually every second. If the advertisements stop, the backups assume the master is down the next-in-line router is elected to run live.

The BSR 64000 offers VRRP to provide fast, efficient recovery in the event of a failure. Dynamic routing protocols such as RIP and OSPF are time-consuming and inefficient for recovery purposes, since they must discover failed routes, alert other routers, run routing algorithms to calculate alternate routers, and then build new routing tables. VRRP can also be used for load balancing, and it is part of both IPv4 (the version of IP that most cable operators currently use) and IPv6.

Distance Vector Multicast Routing Protocol

DVMRP is a distance vector routing protocol for multicast, and it was developed based on RIP. It is used when a router receives a multicast packet and it wants to find out if other multicast routers to which it is connected would like to receive the packet. DVMRP sends the packet to all attached routers and waits for a reply. Routers not interested in the multicast transmission return a "prune" message, which essentially prevents further multicast messages for that group from reaching the router.

The BSR 64000 supports DVMRP, which is primarily used by cable networks in which multicast group members are densely distributed throughout the network. It is useful in applications in which many of the subnets contain at least one router and bandwidth is widely available, but since it transmits multicast messages to all routers within the domain it is less appealing for networks in which the routers are more sparsely distributed and bandwidth is less available.

Protocol-Independent Multicast-Sparse Mode

Routers running multicast protocols depend on their immediate neighbor for updates, which means that when the routers are sparsely distributed (as they are in most cable networks) the constant exchange of multicast data will result in severe degradation of performance. Operators therefore need improved abilities to maintain multicast trees across the access network. PIM-SM was designed to operate efficiently across wide area networks where routers are sparsely distributed across a wide region. It contrasts with PIM-Dense Mode (DM), which operates more like DVMRP and transmits multicast messages to all routers on the network without discrimination.

The BSR 64000 supports PIM-SM but does not support PIM-DM because of the excessive overhead burden. PIM-SM uses the traditional IP multicast model of receiver-initiated membership, supports both shared and shortest-path trees, is not dependent on a specific unicast routing protocol, and uses soft-state mechanisms to adapt to changing network conditions. PIM-SM routes multicast packets to multicast groups, and unlike DVMRP it is protocol independent and can be used in conjunction with unicast protocols such as RIP and OSPF. In contrast, dense-mode protocols such as PIM-DM or DVMRP are designed for networks in which multicast groups are widely represented and bandwidth is plentiful. DVMRP transmits additional overhead to all routers within the domain, while PIM-SM assumes that no router wants data unless it is explicitly requested. It therefore avoids flooding the network with multicast traffic.

The BSR 64000 supports PIM-SM so that operators can efficiently support multicast traffic and support the traditional IP multicast service model of receiver-initiated multicast group membership. Operators can implement PIM-SM with OSPF or RIP and gain the efficiency of multicasting without the overhead congestion across the regional network that can be potentially delivered using DVMRP or PIM-DM.

Secure Protocol Support with the BSR 64000

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, inactivity log out timer, encrypted password protection, and it offers the ability to restrict Telnet management to specified IP addresses only. It 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.

Enabling Next-Generation, Multi-Service Broadband Networks

The selection of the optimal routing protocols to be deployed over cable infrastructure should not be limited by the functionality of the CMTS/ router at the edge of the network. The BSR 64000 allows operators to swiftly deploy enhanced IP services to generate new revenue streams, capture market share and accelerate return-on-investment for network infrastructure. This carrier-class routing solution provides broadband carriers with a competitive edge in defining, deploying, and managing broadband services. Cable operators can deploy the carrier-class BSR 64000 in distribution hubs and select the right mix of protocols to meet their business requirements.

The BSR 64000 features carrier-class implementations of the leading IP protocols, and this fully redundant, high-density solution is architected for 99.999 percent availability so broadband providers can offer enhanced IP services, including IP telephony. The BSR 64000 provides the robust protocol support and the carrier-class architecture necessary to support voice, data, and multimedia services over cable networks. For more information about the BSR 64000—including whitepapers, brochures, and solutions briefs, please visit www.motorola.broadband.com.

MOTOROLA

MGBI

513292-001