 # 132
TECHNOLOGY WHITE PAPER

Riverstone's "Hitless Protection System"
Hot Failover for IP Routers

Tim Wu, Riverstone Networks

ABSTRACT

Riverstone's Hitless Protection System (HPS) increases router and network uptime on carrier networks, providing carrier-class hot failover characteristics similar to SONET and circuit-switched voice networks. It does so by eliminating the single leading source of switch and router downtime: reboots during software upgrades and control module restarts. Carriers who use HPS can save, depending on the size of their networks, hundreds of hours of router downtime on service provider networks.

HPS acts as a complement to frequently deployed Virtual Router Redundancy Protocol, which provides for inter-platform failover in cases of router crashes. Some carriers may choose to deploy a single HPS-capable router as a substitute for a pair of VRRP routers, halving equipment and fiber lease costs for that deployment.

The Hitless Protection System is one part of Riverstone's Carrier Class Metro initiative. It is one of a layered series of redundancy and resiliency technologies designed to maximize network availability and support high availability networking practices.

JUSTIFICATION: CUSTOMER DEMAND

Repeated surveys show that when choosing a service provider, customers value "availability" over every other criteria. "Network reliability" has ranked #1 in each of the last 5 years of Telechoice/Interactive Week's annual ISP surveys. In the 2001 survey, 95% ranked reliability as extremely or very important, making it by far the leading customer concern. The following graph shows the top five surveyed customer concerns when choosing a service provider:

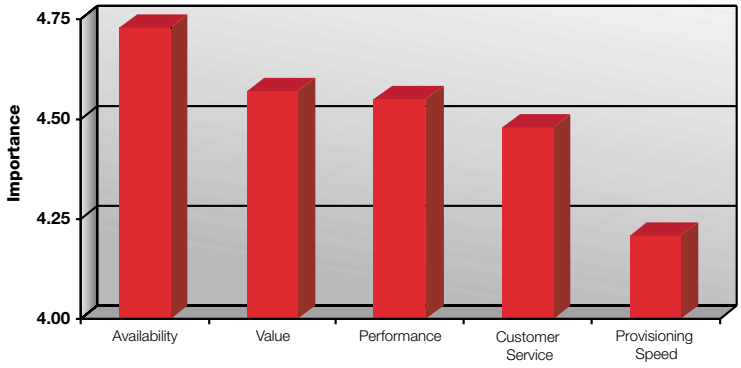

Source: Infonetics 2001 ISP Survey

River STONE NETWORKS

HOW IT WORKS

An average router encounters control module software crashes or requires software upgrades 2 to 4 times a year. Standard routers require a reboot at this juncture. Depending on the complexity of the configuration, a reboot can take from 5 to 40 minutes, during which the router is "down" (i.e., not passing traffic).

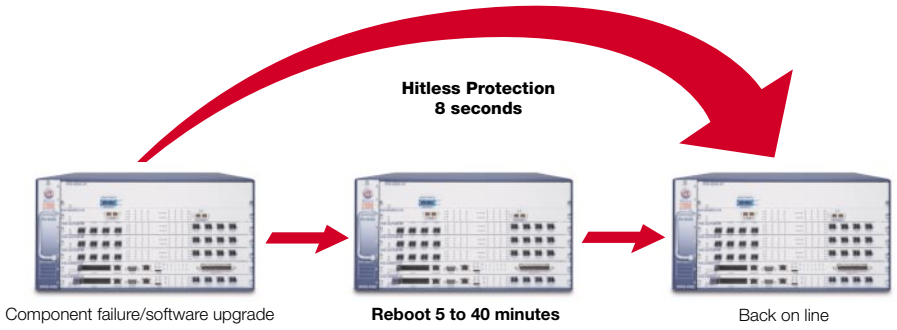

On the other hand, any Riverstone RS router equipped with Hitless Protection System and a redundant control module overcomes a failover or a software upgrade without a reboot. It does so by constantly keeping the redundant control module "ready for action" by continually updating state information with spare CPU cycles. As a result, in the case of control module failure or during software upgrades, it takes only 8 seconds for the new control module to come on line. Crucially, **no traffic is lost** during that time and all flows are maintained through Riverstone's patented ASICs technology.

ONE SMALL STEP FOR A ROUTER, ONE GIANT LEAP FOR NETWORK UPTIME

The Hitless Protection System eliminates reboots, the leading cause of router downtime on today's networks. The following calculations show, for a 500 chassis network typical of a medium-sized nationwide service provider, average expected savings of about **247 hours** of router downtime.

Router Downtime Calculations

Assume a router will encounter, on average, a total of 3 software upgrades or control module failures per year (actual range: 2-6)

Assume an average chassis router reboot takes 10 minutes (actual range: 5-40)

Competitor Router:	3 reboots x 10 min = 30 min/year router downtime
Riverstone with HPS:	3 hitless switchovers x 8 sec = 24 sec/year downtime*

Assume a service provider with 500 routers (average of 33.3/city in 15 cities)

Competitor Network:	500 routers x 30 min/year/router = 250 hours
Riverstone HPS Network:	500 routers x 24 sec/year/router = 3.34 hours*
Savings:	**246.66 hours router downtime/year**

** From failures or software upgrades. There may be other sources of downtime not surveyed here.*

RIVERSTONE HPS AS A VRRP REPLACEMENT: ONE ROUTER INSTEAD OF TWO

Today, many service providers deploy two routers and use VRRP to prevent router downtime from becoming network downtime. Riverstone's Hitless Protection System can either be used to complement VRRP, improving VRRP performance dramatically, or even, in some cases, used as a full substitute for VRRP.

Of course, in any mission-critical network, a carrier will want the double protection of VRRP and HPS. But situations may arise where the costs of maintaining two platforms are too high in terms of equipment, space, and power. In these cases, carriers can deploy a single HPS router in place of a dual VRRP system. The obvious advantage: a halving of both initial equipment and ongoing real-estate leasing, electrical, and fiber leasing costs.

The Hitless Protection System as a VRRP Replacement

Virtual Router Redundancy Protocol

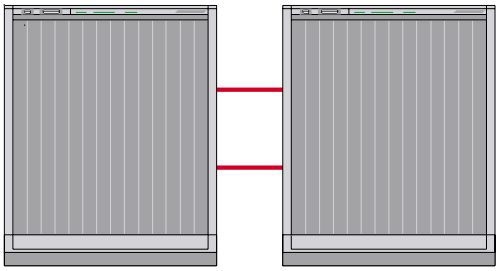

Standard Switches and Routers

- *~15-30 sec lost traffic/failure*
- *Double equipment costs*
- *Double fiber costs*

Hitless Protection System

Riverstone RS16000 Router

- *Zero lost traffic/failure*
- *Single equipment costs*
- *Reduced rack space & electricity usage*
- *Single fiber costs*
- *Can complement VRRP*

(Note: Not recommended for all network positions)

River STONE NETWORKS

HPS OVER VRRP

HPS can also be used as a powerful complement to VRRP. HPS adds another layer of protection, before VRRP kicks in. By reducing usage of VRRP, HPS can save considerable amounts of VRRP-related downtime.

In the calculations shown below, a service provider with 500 chassis routers using HSP as a complement to VRRP can save around 22 hours of lost traffic per year.

Network Traffic Loss Calculations

Assume a router will encounter, on average, a total of 3 events requiring VRRP or HSP protection per year (actual range: 2-6)

Assume that VRRP takes 60 seconds to pick up flows (actual range: 15 sec-1 min)

Competitor Router:	3 reboots x 1 min = 3 min/year lost traffic downtime
Riverstone with HPS:	3 hitless switchovers x 8 sec = 24 sec/year lost traffic*

Assume a service provider with 500 routers (average of 33.3/city in 15 cities)

Competitor Network:	500 routers x 3 min/year/router = 25 hours
Riverstone HPS Network:	500 routers x 24 sec/year/router = 3.34 hours*
Savings:	**21.66 hours lost traffic/year**

** From failures or software upgrades. There may be other sources of downtime not surveyed here.*

HPS SETS THE STANDARD IN ROUTER RESILIENCY: THE ONLY FEATURE OF ITS KIND IN THE INDUSTRY

The Riverstone RS router is the only router available that features a Layer 3 and Layer 4 redundant router failover system. While there do exist implementations at Layer 2, there are no other systems that provide protection for routed traffic, or when services are based on Layer 4 information such as dynamically provisioning bandwidth to guarantee a committed information rate for voice or video traffic.

In summary, Riverstone's new Hitless Protection System sets a new standard for platform layer redundancy in an IP router. HPS is a standard feature of RapidOS version 8.0, and is available on every RS router that has a redundant control module.

River STONE NETWORKS

Acronyms

10GbE:	*10-Gigabit Ethernet*
ABR:	*Available Bit Rate*
ACL:	*Access Control List*
ANSI:	*American National Standards Institute*
API:	*Application Program Interface*
ASIC:	*Application-Specific Integrated Circuit*
ASP:	*Application Service Provider*
ATM:	*Asynchronous Transfer Mode*
BGP:	*Border Gateway Protocol*
BPDU:	*Bridge Protocol Data Units*
CAR:	*Committed Access Rate*
CBR:	*Constant Bit Rate*
CLEC:	*Competitive Local Exchange Carrier (ie: MCI, Sprint, etc.)*
CLI:	*Command Line Interface*
CMTS:	*Cable Modem Termination System*
CoS:	*Class of Service*
CR-LDP:	*Constraint-Based LDP*
CSMA/CD:	*Carrier Sense Multiple Access/Collision Detection*
CSP:	*Content Service Provider*
DiffServe:	*Differential Services IETF Standard*
DLL:	*Data Link Layer*
DOCSIS:	*Data Over Cable System Interface Specification*
DS1/DS3:	*Digital Signal, Level 1 (1.54 Mbps) or 3 (44.7 Mbps)*
DSL:	*Digital Subscriber Line*
DSLAM:	*DSL Access Multiplexer*
E1/E2:	*European Trunk 1/2 (2 Mbps/34.3 Mbps)*
EMS:	*Element Management System*
ERP:	*Enterprise Resource Planning*
EXP:	*Experimental bits*
FDDI:	*Fiber Distributed Data Interface*
FEC:	*Forwarding Equivalence Class*
GARP:	*Generic Attribute Register Protocol*
GVRP:	*GARP VLAN*
HPS:	*Hitless Protection System*
HSSI:	*High Speed Serial Interface*
IANA:	*Internet Assigned Numbers Authority*
IBGP:	*Internal Border Gateway Protocol*
IETF:	*Internet Engineering Task Force*
IGP:	*Interior Gateway Protocol*
IP:	*Internet Protocol*
IS-IS:	*Intermediate Systems to Intermediate Systems*
ISP:	*Internet Service Provider*
ITU:	*International Telecommunications Union*
LAN:	*Local Area Network*
LDP:	*Label Distribution Protocol*
LEC:	*Local Exchange Carrier*
LER:	*Label Edge Router*
LFAP:	*Lightweight Flow Accounting Protocol*
LSP:	*Label Switched Path*
LSR:	*Label Switched Router*
MAC:	*Media Access Control*
MAN:	*Metropolitan Area Network*
MCDN:	*Microcellular Data Network*
MDI:	*Media Dependent Interface*
MDU:	*Multiple Dwelling Unit*
MIB:	*Management Information Base*
MLPPP:	*Multi Layer Point to Point Protocol*
MMDS:	*Multi-point Multi-channel Distribution System*
MMF:	*Multi-mode fiber*
MPLS:	*Multi-Protocol Label Switching*
MSN:	*Management Network Server*
MSP:	*Metropolitan Service Provider*
MSTP:	*Multiple Spanning Tree Protocol*
MTU:	*Multiple Tenant Unit*
MVST:	*Multiple-VLAN Spanning Tree*
OC-3/OC-12:	*Optical Carrier 3/12 (155 Mbps/622 Mbps)*
OSPF:	*Open Shortest Path First*
PCS:	*Physical Coding Sublayer*
PDU:	*Protocol Data Units*
PE:	*Provider Edge*
PEF:	*Packet over Ethernet over Fiber*
PES:	*Packet over Ethernet over SONET-based optical*
PEW:	*Packet over Ethernet over WDM-based optical*
PHY:	*Physical Layer*
PMA:	*Physical Media Attachment*
PMD:	*Physical Media Dependent*
PON:	*Passive Optical Networking*
POP:	*Point of Presence*
POS:	*Packet over SONET*
PPP:	*Point to Point Protocol*
PSTN:	*Public Switch Telephone Network*
PVC:	*Private Virtual Circuit*
PVST:	*Per-VLAN Spanning Tree*
QoS:	*Quality of Service*
RBOC:	*Regional Bell Operating Company (ie: PacBell, etc.)*
RED:	*Random Early Discard*
RSTP:	*Rapid Spanning Tree Protocol*
RSVP:	*Resource Reservation Protocol*
RSVP-TE:	*RSVP-Traffic Engineering*
rt-VBR:	*real-time Variable Bit Rate*
SAN:	*Storage Area Network*
SAP:	*Session Announcement Protocol*
SLA:	*Service Level Agreement*
SMF:	*Single-mode fiber*
SNMP:	*Simple Network Management Protocol*
SONET:	*Synchronous Optical Network*
STP:	*Spanning Tree Protocol*
T1:	*Trunk 1 (1.544 Mbps)*
TCP/IP:	*Transmission Control Protocol/Internet Protocol*
TDM:	*Time Division Multiplexing*
TOS:	*Terms of Service*
UBR:	*Undefined Bit Rate*
VBR:	*Variable Bit Rate*
VLAN:	*Virtual LAN*
VoD:	*Video on Demand*
VPN:	*Virtual private network*
VRRP:	*Virtual Router Redundancy Protocol*
WAN:	*Wide Area Network*
WDM:	*Wave Division Multiplexing*
WFQ:	*Weighted Fair Queuing*
WRED:	*Weighted Random Early Discard*
XGMII:	*10G Media Independent Interface*

Riverstone Networks, Inc.
5200 Great America Parkway, Santa Clara, CA 95054 USA

877 / 778-9595 *or* ***408 / 878-6500*** *or* ***www.riverstonenet.com***