REPRINT

NetworkWorld GERMANY IDG

www.networkworld.de
Knowledge projection/
lead for IT professionals

22/01

NetworkWorld GERMANY 17 2001 tested

EANTC
NetworkWorld Partner Lab

Test Series: MPLS Router with RSVP-TE

Strong Pair for Carrier Networks

www.networkworld.de
Table with the measured values of the Data-Plane-Tests

AVICI TSR

NetworkWorld GERMANY 17 2001 tested ★★★★☆ = very good

Manufacturer
Avici
www.avici.com
Price: upon request

Technical Data
The core router functions with a passive backplane without central switching fabric architecture. It has a switching capacity of 400 Gbps per bay. Up to 40 modules can be installed per bay.

Test Results
- High performance for all traffic streams through three-dimensional processor architecture
- Robust RSVP implementation
- Easily operated through clear configuration menus

MPLS support for routers is no longer a rarity. The technology market is ready as the two tested devices, both with RSVP-TE support, Avici's "TSR" and Riverstone's "RS 8000", prove in impressive fashion.

GABRIELE SCHRENK, ANDREAS WURM

More and more carriers are choosing MPLS (Multiprotocol Label Switching) to offer their customers value-added services like virtual private networks (VPN). Reason enough for NetworkWorld Germany, together with Partner Lab EANTC, to test two additional MPLS routers (to see previous MPLS tests, see volumes 06/01 and 14/01). Once again, the focus is the traffic engineering expansion of the resource reservation protocol (RSVP-TE).

The test subjects: Riverstone RS 8000 and Avici TSR

Two unlike test subjects were put through extensive performance and functionality tests this time: Avici's TSR core router, and Riverstone's RS 8000 edge router. The devices target different markets, and therefore are so different that they can be viewed as partners rather than competitors. The goal of Avici and Riverstone is to offer a collaborative end-to-end solution for both carrier and service providers. The performance of the TSR at the core supplements the manifold possibilities and performance of the RS 8000 in the MAN (metropolitan area network). Both manufacturers put great value on the interoperability of the devices, which is why we took the compatibility of the two routers into consideration while subjecting them to common test scenarios during our evaluation.

Based on their own statements, Riverstone developed their router to "transform pure bandwidth into profitable services for MANs". According to the manufacturer, the RS family delivers intelligent bandwidth with versatile service possibilities, offering routing as well as switching. It supports the most important interfaces like Gigabit Ethernet, Packet over Sonet (POS), and ATM (asynchronous transfer mode). The RS 8000 uses Asics' fourth generation to offer dynamic

RIVERSTONE RS 8000

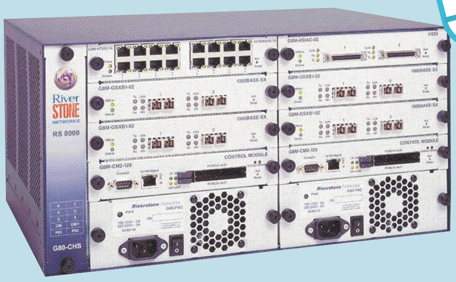

Foto: Riverstone

Manufacturer
Riverstone Networks
www.riverstonenetworks.com
Price: upon request

Technical Data

The modular edge router provides interfaces for Gigabit Ethernet, Packet over Sonet (POS) and asynchronous transfer mode (ATM), in addition to supporting VPNs. The backplane has a capacity of 32 Gbps.

Test Result

- Excellent throughput of data in the performance tests
- Stable RSVP implementation and an already well-performing beta version
- Easy to operate configuration

NetworkWorld Germany 17 2001 tested ★★★★☆ = very good

TEST RESULTS

	Signaling-Functionality RSVP-TE	Call Capacity RSVP-TE	Data Plane IP	Data Plane LER	Overall Operation and Features
Riverstone	★★★★★ (4,5)	- Not tested	★★★★★ (5)	★★★★★ (5)	★★★★☆ (4)
Avici	★★★★★ (5)	★★★★★ (5)	★★★★★ (4.5)	- Not tested	★★★★☆ (4)

★= unsatisfactory ★★= acceptable ★★★= good ★★★★= very good ★★★★★= excellent

bandwidth regulation and VPN networks that are MPLS-ready. Reliability and availability provided by Riverstone's hitless protection system are intended to make the devices suitable as carriers.

Test preparation and planning

The manufacturer supplied software version 8.020 for the RS 8000, and beta version 9.0.A.18 for the tests dealing with prioritizing traffic streams. The router was equipped with eight Gigabit Ethernet modules; the backplane had a capacity of 32 Gbps.

Avici designed the TSR and SSR router families with the goal of satisfying the demands for bandwidth and service of both carriers and service providers. The result is an unusual system architecture. It allows high scalability and enables cost-effective capacity expansion from five Gbps to five Tbps.

Avici loaded their product with software version 4.1.2.0. For test purposes, the router was equipped with two Gigabit Ethernet modules and two POS-OC-48 modules. In addition, four POS-OC-12 modules were installed. The passive backplane has no central switch fabric architecture; the switching capacity to the backplane is 60 Gbps per interface. The TSR features a distributed architecture. In the main-frame world, it would be called "massive parallel": there is no single backplane, as each neighboring network processor is connected in three dimensions. Because of this, the backplane capacity varies depending on the cabinet configuration (bay). One bay can be equipped with up to 40 router module cards. In addition, network administrators can connect several bays. These connected bays logically also act as routers. Therefore, a capacity of up to five Tbps is possible.

We used the Spirent Communications AX4000 as our diagnostic tool to analyze protocols and to generate throughput loads and IP packets with and without MLPS labels. At the same time, the AX4000 emulates the RSVP-TE protocol at its interfaces and can, therefore, act like one or more MPLS router. It analyses and generates data with routing speed at its ATM, POS, or Gigabit Ethernet ports. During the MPLS testing of the two test subjects, we concentrated not only on the ability to do the job, but also on the ability to work together. In addition to the overall performance,

ASSESSMENT

	Signaling	Throughput	Equipment and Operation	Overall Assessment
Weight	1/3	1/3	1/3	
Riverstone	★★★★★ (4,5)	★★★★★ (5)	★★★★☆ (4)	★★★★★ 4,5
Avici	★★★★★ (5)	★★★★★ (4,5)	★★★★☆ (4)	★★★★★ 4,5

★= unsatisfactory ★★= acceptable ★★★= good ★★★★= very good ★★★★★= excellent

Testing capabilities reach their limits

The two MPLS switch routers, Avici's TSR and Riverstone's RS 8000, showed off their best side at the EANTC Lab. Once again, in the third phase of our MLPS test series (see NetworkWorld 14/01 and 06/01), we established that these manufacturers of carrier network devices impress with high quality standards. Besides other tested areas, both test subjects showed excellent results during the evaluation of the traffic engineering software.
The low error rate of the tested devices is gratifying. Nevertheless, the perfect router does not exist. The possibilities are limited by the available test methods and the limited duration of the testing. For example, measuring the IP throughput with backplane capacity in the terabit realm can be achieved only with an incredible array of measurement equipment. The same can be said for testing maximum MPLS data throughput. Avici, for example, can achieve up to 15,000 paths per port. If there were weaknesses in implementation at the beginning of this year, they are no longer an issue. The measuring stick for the next series of tests will have to be raised.

Gabriele Schrenk is a member of the executive board of the European Advanced Networking Test Center.

we tested the devices for traffic engineering support, and whether they, as MPLS routers, conform to signalling standards and sent the IP packets correctly. We did not test whether the two test subjects route at wire-speed at full backplane capacity. In our opinion, such tests are of questionable interest to the end-user, since no carrier works at a network capacity over 60%. The first part of the testing

PERFORMANCE CHARACTERISTICS: AVICI TSR AND RIVERSTONE RS 8000

	Riverstone RS 8000	Avici TSR
Signaling Protocols		
LDP	+	+
CR-LDP	-	-
RSVP-TE	+	+
OSPF-TE	+	+
BGP-MP (RFC 2547)	As provider	As provider
Routing Protocols		
OSPF	+	+
IS-IS	+	+
RIP	+	-
BGP-4	+	+
Multicast Routing (PIM, DVMRP)	+	+
IP Services		
IP-VPNs	+	+ (As provider)
IPsec	-	-
Diffserv-MPLS-Mapping, E-LSPs and L-LSPs	+	Planned 2nd half of 2002
Layer 4 switching	+	-
Stateless Packet filter	+	Planned 1st half of 2002
Security measures against denial of service attacks	+	+
MPLS Options		
ATM / FR / Ethernet 802.1q over MPLS (Martini Draft)	+ (Nor for ATM and FR)	-
Strict Priority Queueing	+	Planned 1st half of 2002
Supported Physical Interfaces		
Packet over Sonet (POS) (STM-1, STM-4, STM-16, STM-64)	+ (Not STM-64)	+
STM-1 / STM-4 channelized	-	-
Gigabit Ethernet	+	+
Fast Ethernet	+	-
ATM (serveral bandwidth)	E3, STM-1, STM-4	-
MPLS supportive interfaces	POS, GE	POS, GE

was focused on the signal protocol RSVP-TE.

RSVP-TE implementation

The two devices were tested to prove that they master what they are built for: building pathways through the MPLS network according to conventional norms. We assumed that the IP routing works properly. Our main interest lies in how the test subjects utilize traffic engineering with RSVP-TE in an MPLS network.

The testers connected the Spirent emulator to the router with two interfaces each, either Gigabit Ethernet or POS. In order to determine the path the RSVP-TE flow is to take, network administrators can use explicit route objects (ERO). They can be used for traffic engineering, increasing the capacity of existing network resources. It is possible to detach single packet streams from the IP routing mechanism, so that alternative, less used paths can be taken in order to fully utilize the capacity of the router. Both routers understood and used EROs correctly.

Record route objects (RRO) are important to avoid loops while building label switched paths (LSP). Just like the ERO, the RRO has a list of routers that constitute an LSP through the MSLP network. Unlike the EROs, this list is created while the path or RESV message finds its way to the target router. This characteristic is ideal according to the standard. It requests a confirmation from each node by means of adding its IP address to the end of the incoming RROs. The receiver of the list can follow the LSP based on the information contained in the RRO. These can then be used for network administration, or as data for ERO input. Both routers passed this test, as well, without any problems.

Next, we checked how both routers reacted to data overload throughput at the interface when the link was overbooked. Both devices should deny further throughput requests, and both devices reacted as they should. The RS 8000 made a small mistake by sending the incorrect error code with the path denial;

Deployment of MPLS routers

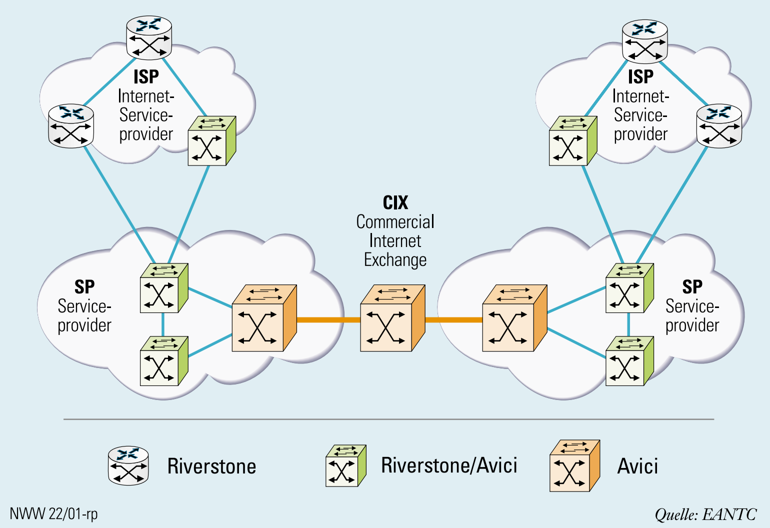

Collaboration: the RS 8000 is suited for use on the edge of the network. The TSR from Avici is designed as a core router; however, it could also do its job at the edge of the network.

Riverstone will fix this problem in the near future. As far as the interoperability with other manufacturers is concerned, this incorrect error code should not interfere with normal operation.

Should an error or overload of a given interface/link occur, the LSPs can be prioritized and redirected according to their priority. The setup and hold priorities are defined for each LSP. This setup priority regulates the importance during the setup of the path. In contrast, the hold priority regulates the claim for the already existing path. Setup and hold priorities can be set from zero to seven – the smaller the number, the higher the priority. The actual software that we used for the Riverstone router did not support this mechanism yet. The manufacturer gave us the newest beta, version 9.0.A.18, instead, which is capable of prioritizing data streams. Using the beta version, the router passed the test. Avici did not show any weaknesses either. The LSPs marked as low priority were successfully delayed in favor of those marked more important.

RSVP-TE path capacity

The path capacity test offers quantitative data on how well the router handles a large number of LSPs. We controlled the setup of paths using a restricted time window. The data load generator emulated two connected routers, which kept each tunnel in soft state mode by regularly sending path and RESV messages. The virtual routers renewed the paths by sending periodic messages that reset the timers. If one of those messages does not arrive, the communication partner presumes that the LSP no longer exists.

In doing this, a realistic data load for the network components was achieved. The design of the RSVP protocol requires a higher administrative effort, which rises proportionally with the number of paths used, as opposed to the label distribution protocol (LDP).

The LDP is another protocol used to build connections over LSPs in MPLS networks. Our data load generator built 10,000 LPSs between two ports of the TSR routers. This reached the maximum capabilities of our Adtech emulator. The Avici TSR mastered the demand without any problem; all paths were set up and successfully held in soft state.

MPLS data throughput

According to Avici, the current software version, 4.1.2.0, limits the LSPs to 15,000 per port. Riverstone states the capacity of the RS 8000 to

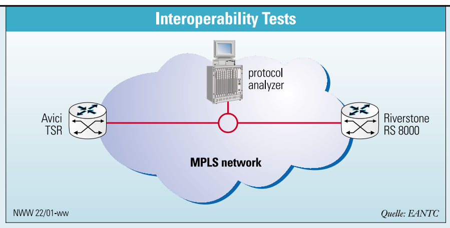

No weaknesses: both devices worked together flawlessly, and mastered all required tasks without problems

be 3,000 LSPs per port. The TSR, acting as an Internet core router, therefore, must handle many different LSPs — customer data streams — at the same time. Riverstone has a different field of application, demanding different capabilities; therefore, we did not evaluate its performance in this area.

With all large IP/MLSP routers, the current level of technology allows data to be routed at the maximum physical transmission rate possible by the connections without blockage, called wire-speed. Sometimes problems occur with a larger number of parallel IP or MPLS streams. For this test, we chose a realistic RSVP-TE range of 500 to a maximum of 1,000 RSVP paths.

An MPLS router can assume several different roles. It can be an ingress node at the incoming point of the network, analyzing, classifying, and sending incoming IP packets via LSP to the next node. During this process, the router assigns the packets to an LSP and identifies them with the corresponding label.

Used at the core of the network, the label switch router (LSR) guides incoming data with an LSP to the appropriate exit port and next LSP. A look-up in the forwarding table attaches a label, which provides the path to the next network node, until the egress router at the end of the MPLS network removes the label, and sends the IP packet according to its routing table to the next connected network. In order to check the packet loss rate, the following conditions are defined by the IETF:

- Constant maximum load and variable packet size, from 64 byte to maximum frame length
- Constant packet size of 64 bytes at a variable load of 20-100%

We connected the routers to two interfaces of the AX4000 load generator. The test examined the following roles and tasks of the MPLS devices:

PRIORITIZING OF DATA STREAMS

	1st LSP	2nd LSP	3rd LSP
Bandwidth (Mbps over Gigabit Ethernet/POS OC-12)	RS 400 / TSR 250	RS 400 / TSR 250	RS 400 / TSR 250
Test 1: Simultaneous build of three paths, Hold priority = Setup priority			
Setup / Hold priority	3 / 3	5 / 5	4 / 4
Expected result	up	down	up
Result Riverstone	up	down	up
Result Avici	up	down	up
Test 2: LSP 3 competes with established LSP 1 and LSP 2; different Hold and Setup priorities			
Setup / Hold priority	3 / 2	4 / 1	5 / 0
Expected result	up	up	down
Result Riverstone	up	up	down
Result Avici	up	up	down
Test 3: 3rd LSP is not of higher priority than already established LSP 1 and LSP 2			
Setup / Hold priority	3 / 2	4 / 1	2 / 2
Expected result	up	up	down
Result Riverstone	up	up	down
Result Avici	up	up	down
Test 4: 3rd LSP is of higher priority than the 1st LSP, but lower than the 2nd LSP			
Setup / Hold priority	3 / 2	4 / 1	1 / 1
Expected result	down	up	up
Result Riverstone	down	up	up
Result Avici	down	up	up
Test 5: 3rd LSP is of higher priority than the 1st LSP, but lower than the 2nd LSP; however, superseding of 1st LSP is not enough for the build of the 3rd			
Bandwidth (Mbps over Gigabit Ethernet/POS OC-12)	RS 400 / TSR 250	RS 400 / TSR 250	RS 800 / TSR 500
Setup / Hold priority	3 / 2	4 / 1	1 / 1
Expected result	up	up	down
Result Riverstone	up	up	down
Result Avici	up	up	down
Test 6: 3rd LSP is of higher priority than the existing ones; however, superseding of both LSPs in not enough for the build of the 3rd			
Bandwidth (Mbps over Gigabit Ethernet/POS OC-12)	RS 400 / TSR 250	RS 400 / TSR 250	RS 1200 / TSR 800
Setup / Hold priority	3 / 2	4 / 1	0 / 0
Expected result	up	up	down
Result Riverstone	up	up	down
Result Avici	up	up	down

Source: Spirent

The tests were done with the "Adtech AX 4000", a protocol analyzer and load generator manufactured by Spirent Communications. The device generates IP packets with or without MPLS labels. At the same time, the AX 4000 emulates the RSVP-TE protocol at its interfaces, and can behave like one or more MPLS routers. The ATM, POS, or Gigabit Ethernet interfaces generate and analyze data at wire-speed.

- Label edge router: classifies incoming data streams and forwards them to the MPLS network
- Label switch router: routes MPLS packets
- IP throughput without MPLS: as comparison data

Routers at the beginning of the network have the largest workload. The ingress label edge router (LER) classifies and analyses all the incoming IP packets. That means it assigns a forward equivalence class (FEC), attaches the appropriate label, and transmits them over the relevant exit port to the LSP.

The LER is generally at the limit of the provider network and connects one or more customers of the supplier with their services. Such a device must support different protocols and offer security measures. The RS 8000 by Riverstone mastered all tests with a varying number of LSPs and packet sizes without packet loss. The TSR, as a core component, is primarily used as an LSR. Its task is to route the incoming data traffic according to the LERs or LSRs. Therefore, we did not test the way it acts as an LER. In the core of an MPLS network, the LSRs are responsible for routing the received IP packets, possibly with a new label, to the appropriate exit port. These labels are retrieved from a label information base (LIB). The LIB contains all the existing LSP's, and describes the incoming and outgoing ports, as well as possible reservations and priorities.

After all paths in the network are built, a LSR has only one task: to switch labels based on the LIB. Routing duties are only necessary in the case of an error. Both devices passed these tests without losing packets.

For comparison purposes, we also tested pure IP performance. In this test, both test subjects routed packets without labels. They determined the next network node by calculating the target address of the packet. The RS 8000 by Riverstone showed excellent results. Even with smaller packet sizes, no loss was noted during the test. The Avici TSR achieved a throughput of 100% with a realistic mix of 500 and 1000 different IP flows having packets of all sizes.

A very conservative, but in effect theoretical, lab test illustrates that the R&D department at Avici did excellent work when developing the TSR: using a single, full bandwidth IP data stream with IP packets of 64 Bytes, the TSR routed using only 55% of available capacity. The manufacturer commented that the switching architecture of the TSR and SSR models is aimed at routing IP packets in the correct order in all possible configurations. A single IP flow cannot possi-bly demonstrate the strength of the Avici routers' massively distributed resources.

Configuration of the routers

The command line interface (CLI) by Riverstone is well documented and easy to use. Simple MPLS configurations can be done quickly and easily without help.

The Avici menu is nearly identical to Cisco's internetwork operating system (IOS). In addition, Cisco's "config text files" can be loaded and used in the TSR. The software will alert the system administrator if there are functions that are not supported. As mentioned earlier, we tested the interoperability of both routers. We were especially interested in the following:

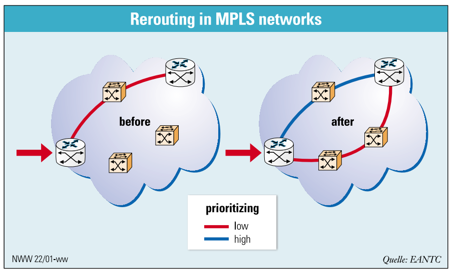

Prioritizing of data streams: RSVP-TE supports dynamic redirecting, which allows diversion of lower prioritized data streams to substitute routes.

- Setting up of LSPs with LDP
- Setting up of traffic engineering LSPs with RSVP-TE

We used open shortest path first (OSPF) as the routing protocol between the two routers. Neither one of the routers had any problems working together during the evaluation, nor were any problems noted during completion of all tests.

ABOUT THE AUTHOR

GABRIELE SCHRENK

As a member of the Executive Board at EANTC (www.eantc.de), Gabriele Schrenk is responsible for testing and consulting.