

# Accounting for Profitability: Riverstone's Lightweight Flow Accounting Protocol

**ABSTRACT** A leading feature of the Riverstone RS router is its flow and traffic accounting capabilities. The RS platform was designed for service providers and therefore engineered with the ability to account for every byte and packet of every flow without affecting the routing and switching performance. Based on Riverstone's Lightweight Flow Accounting Protocol (LFAP), accounting capabilities are integrated into the ASIC architecture of every Riverstone router, and are a part of what makes Riverstone's routers stand out in the service provider market.

This paper is intended to provide a full introduction to the central capabilities of Riverstone's accounting architecture. The paper will also explain how service providers can use Riverstone's capabilities to improve their network management and also, if desired, transition or improve a usage-based billing model.

## WHY DO ACCOUNTING FEATURES MATTER?

Whether a service provider uses a flat-rate or usage-based billing model, there is one fact that cannot be avoided — **bandwidth to service providers is not free**. Investing in the equipment necessary to provision bandwidth has real costs, which need to be recouped. The goal of service providers should be to deploy and to maintain the necessary amount of equipment, or bandwidth, that is needed to meet its business objectives. Small errors in this process can lead to large losses.

Accurate network usage information is central to any coherent model of network management. Service providers that have the full picture of network usage will always be at an advantage when it comes to both allocating existing bandwidth and planning the deployments of the future. A reliable and accurate network accounting system is central to this goal.

Usage-based billing can be understood as perhaps the best way to achieve the goal of matching customer revenue and service provider costs. It is true, however, that the North American market has been slow to move toward usage-based billing. Industry experts expected that value-added services would help service providers move away from flat-rate billing, but the full effects have yet to be realized. The situation, however, is much different overseas. In Europe, Asia, and Latin America, usage-based billing for even basic access services is the standard. We may begin to see a transition within North America back to usage-based billing as a means of improving service provider margins.



**THE RIVERSTONE SOLUTION:  
LIGHTWEIGHT FLOW  
ACCOUNTING PROTOCOL  
(IRFC 2124 LFAP V1)**

LFAP is an accounting protocol that was developed by Riverstone to meet the flow accounting needs of service providers. The protocol is particularly valuable for gathering information for Layers 3 and 4 flows. Accounting data for Layer 2 networks can also be collected using LFAP as long as the device is running in Layer 4 bridging mode.

The initial version was submitted to the IETF as an informational RFC 2124. The current implementation is LFAP version 4. The latest LFAP specification is available to any interested parties at the following link: <http://www.nmops.org/draft-rfcd-info-calato-01.txt>

The hardware implementation of LFAP is the outstanding aspect of Riverstone's accounting architecture. Across the RS platform, flow and traffic data is accounted for by the hardware, rather than software, so that turning accounting features does not impact system performance. Periodically, the router will gather the flow data and send the information to an accounting server using the LFAP protocol. The time interval at which the flow data is sent to the accounting server is configurable down to every five minutes.

In version 4 of LFAP, the following information is captured and transmitted, and can be used as the service provider wishes:

**LFAP v4 Fields Collected**

- Source/destination IP address
- Source/destination AS
- Source/destination port
- Ingress/egress port
- Type of service (DSCP)
- Protocol
- Total bytes
- Total packets
- Total time of flow

Version 5 of LFAP is expected to be available in Riverstone's RapidOS 9.0; LFAP v5 will also be submitted to the IETF as an informational RFC. In addition to the accounting fields in LFAP v4, the following additional fields will be added to the technical specifications (the implementation of LFAP v5 in RapidOS 9.0 will include a subset):

**LFAP v5 Fields Collected**

- Egress ATM Subinterface
- Egress Frame Relay Subinterface
- Next Hop IP
- TCP Control Bits
- Next Hop AS
- Source Address Netmask
- Destination Address Netmask
- Destination BGP Community
- Source BGP Community
- Traffic Index (Riverstone specific)
- VLAN ID
- VLAN priority
- Source Virtual Address (for NAT/LSNAT)
- Destination Virtual Address (for NAT/LSNAT)
- MPLS (TBD)



**CARRIER CLASS  
ACCOUNTING**

LFAP was designed for service providers, and as a result features that are designed for the demands of modern networks. The protocol works in conjunction with Riverstone hardware-based routers to capture traffic data reliably so that network managers can collect flow accounting data and transmit the data over a reliable protocol. A key benefit of Riverstone's hardware traffic accounting capabilities is that the accounting feature does not impact system performance.

**1. TCP Based vs. UDP Based**

LFAP is inherently more reliable than other accounting protocols since it is TCP based at the transport layer. The flow accounting data is sent from the router to the accounting server and the delivery of the data is confirmed using TCP's capabilities. An accounting protocol that uses UDP supplies no delivery guarantees. Without a reliable transport mechanism, issues of packet order and reliability could be introduced as the traffic accounting data travels from the network element to the data collector.

**2. Real Time Traffic Flow Data**

LFAP is designed to deliver real-time flow data, even in the middle of a flow. The LFAP protocol specifies that the router establish a flow traffic file once a flow begins. As the flow continues, the router will gather the incremental traffic flow data and transmit that data to the accounting server. Other protocols, on the other hand, wait until the flow has expired and then send all flow traffic data to the accounting server.

**3. Fail-Over Accounting Servers**

Riverstone's current implementation also allows for the configuration of up to two fail-over servers. When a router is configured to send LFAP data to an accounting server, a primary server is established. Two additional backup servers can also be specified. In the event that the primary server fails, the router will automatically recognize the failure and send that traffic data to the backup server.

**4. LFAP MIB**

RapidOS provides a MIB that can be used to monitor the LFAP accounting process in the Riverstone equipment. As a result, service providers can be assured that the accounting process is functioning and can be alerted when the accounting process fails.



## THE USES OF LFAP DATA

There are two principal uses for flow accounting data. The first is for network usage analysis: the study and understanding of how users are using the network. The second is for usage-based billing.

### **Network Usage Analysis – The Key to Service Provider Profitability**

Understanding exactly how the network is actually being used is the key to maximizing the efficiency and profitability of that network. Today, service providers are frequently faced with increasing bandwidth demands from individual customers without a corresponding increase in revenue. The situation calls for carefully considered expansion that only network usage analysis can facilitate, and perhaps a greater move to usage-based billing.

In an ideal world, users would be considerate and limit the amount of bandwidth they consume during peak hours. The reality is the opposite. Bandwidth hogging users and applications consistently consume a disproportionate share of the available bandwidth. For example, Netzero, a dialup ISP, recently reported that 12% of users accounted for 53% of their bandwidth usage. While bandwidth-intensive applications like file-sharing (Napster) and streaming video create greater bandwidth demands, flat-rate billing means only the same amount of revenue is collected. Bandwidth abuse by a small percentage of the customer base results in unwanted network congestion for everybody else. Oftentimes, service providers must expand their network capacity to maintain service levels.

However, when considering network expansion, service providers should be asking questions such as: Which applications are consuming the most bandwidth? When are the peak traffic times? Do I need to add network capacity or can I implement network policies to improve performance? Riverstone's ability to collect accurate flow and traffic accounting features allow service providers to gain an accurate understanding of how the network is being used.

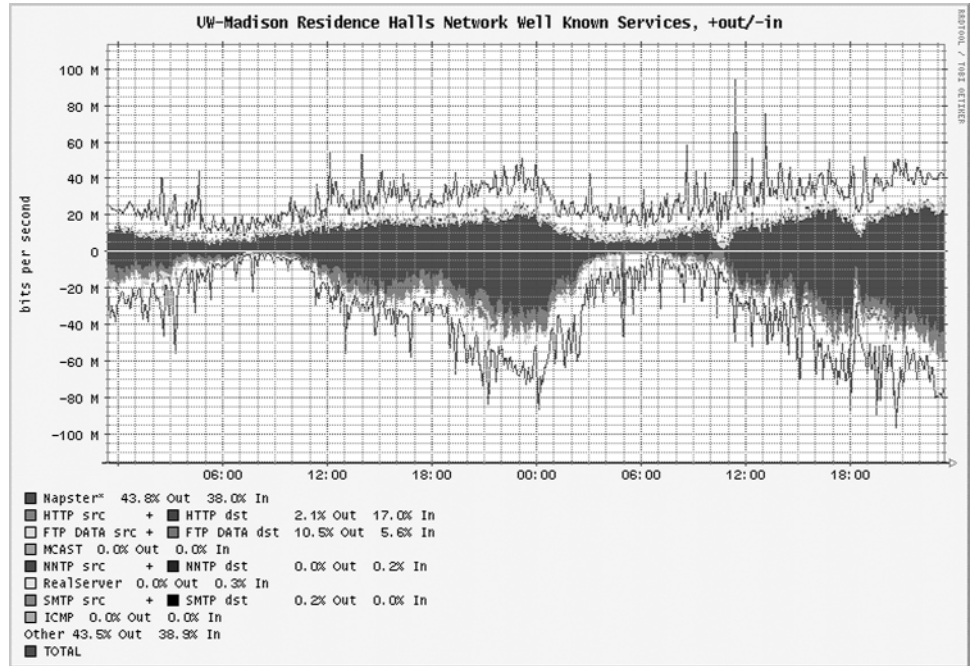
### **Example: Using LFAP to Deal with Napster Traffic**

With the data gathered with LFAP, service providers can create policies to influence end user behavior. For example, if Napster traffic is known to be prevalent on Friday evenings from 7pm to 10pm, network policies can be put in place to limit the amount of bandwidth dedicated to Napster traffic. The same can be done for streaming video or music applications. The key, however, is to be able to collect the data. While many network hardware vendor's products cannot perform the accounting feature in addition to wire speed routing and switching, the Riverstone architecture can account for every byte and packet.

Gathering and analyzing flow accounting data does not require a complex accounting system (though such systems are available for those who need them). Open source applications, used in conjunction with Riverstone's accounting architecture, will give an accurate graphical view of network traffic. For example, the FlowScan tool, developed by Dave Plonka from the University of Wisconsin-Madison, is a versatile tool that can meet many service providers' needs. FlowScan and the necessary tools required to gather flow accounting data can be found at Riverstone's open source network management Website at <http://www.nmops.org>.



Below is an example of FlowScan displaying the traffic of well-known services from the University of Wisconsin network. The graph clearly shows that Napster consumes as much as 20 Mbps of bandwidth on the outbound and over 40 Mbps of bandwidth on the inbound. If this were a service provider that was running out of bandwidth, a policy reducing Napster's bandwidth consumption during the hours of 12pm and 2am would increase the amount of available bandwidth for other applications.



### Valuing Peer Relationships

LFAP can also be used to gain a better understanding of the value of existing peering relationships. FlowScan provides tools that allow tabular or graphical views of the AS to AS traffic data. Tables 1 and 2 were created from LFAP data gathered, using the FlowScan tool. The LFAP data was combined with data from <ftp://ftp.arin.net/netinfo/asn.txt>. Table 1 displays the amount of incoming traffic from the top 10 origin ASNs. The ASNs are ranked by inbound bits/sec traffic. Table 2 shows the top 10 destination ASNs. The data in Table 2 is ranked by outbound bits/sec traffic. FlowScan includes a tool that will graphically display historical source and destination AS traffic data, similar to the Napster graph above.



**TABLE 1**

Top 10 origin ASNs by bytes in  
for five minute flow sample ending Mon Feb 19 13:35:16 2001

Rank	Origin-AS	Bits/sec in	% of total in	Bits/sec out	% of total out
#1	<b>UONET</b> (3582)	2.3 M	62.3%	54.9 k	9.4%
#2	<b>NSFNETTEST14-AS</b> (237)	241.3 k	6.5%	16.6 k	2.8%
#3	<b>AGIS-NET</b> (4200)	123.9 k	3.4%	75.0 k	12.8%
#4	<b>CRITICALPATH</b> (10627)	97.3 k	2.6%	3.2 k	0.5%
#5	<b>CNCX-AS1</b> (2828)	82.1 k	2.2%	2.2 k	0.4%
#6	<b>HURRICANE</b> (6939)	66.8 k	1.8%	2.1 k	0.4%
#7	<b>NETUSACOM-AS</b> (3967)	65.9 k	1.8%	8.2 k	1.4%
#8	<b>QUOTECOM</b> (11803)	41.9 k	1.1%	2.7 k	0.5%
#9	<b>MULTICASTTECH</b> (16517)	35.4 k	1.0%	1.9 k	0.3%
#10	<b>ALTERNET-AS</b> (701)	28.2 k	0.8%	7.3 k	1.2%

**TABLE 2**

Top 10 destination ASNs by bytes out  
for five minute flow sample ending Mon Feb 19 13:35:16 2001

Rank	Origin-AS	Bits/sec in	% of total in	Bits/sec out	% of total out
#1	<b>PLAYBOY-BLK-1</b> (14068)	0.0	0.0%	226.3 k	38.6%
#2	<b>AGIS-NET</b> (4200)	123.9 k	3.4%	75.0 k	12.8%
#3	<b>UONET</b> (3582)	2.3 M	62.3%	54.9 k	9.4%
#4	<b>VERIO</b> (2914)	21.7 k	0.6%	20.4 k	3.5%
#5	<b>NSFNETTEST14-AS</b> (237)	241.3 k	6.5%	16.6 k	2.8%
#6	<b>ALTERNET-AS</b> (703)	422.2	0.0%	14.6 k	2.5%
#7	<b>ATT-INTERNET4</b> (7018)	23.3 k	0.6%	8.7 k	1.5%
#8	<b>NETUSACOM-AS</b> (3967)	65.9 k	1.8%	8.2 k	1.4%
#9	<b>ALTERNET-AS</b> (701)	28.2 k	0.8%	7.3 k	1.2%
#10	<b>CISCOYSTEMS</b> (109)	10.0 k	0.3%	5.9 k	1.0%

The value of this data is that peering relationships can be analyzed. Furthermore, this information can be used to charge tiered rates for intra- versus inter-AS traffic to generate revenue more closely tied to service provider costs. The hardware data collection capability allows source and destination AS to be collected for every source IP address and every ingress/egress port.

#### Usage-Based Billing

Usage-based billing drives a clear and lucrative revenue model. Instead of charging all customers a flat rate, this system allows service providers to charge the high-usage customers more to gain additional revenue. Unfortunately, due to the lack of reliable data collection, most service providers in



North America have largely failed to implement true usage-based billing and have either used flat-rate billing or cost-averaging systems. Riverstone LFAP-based accounting system can be used to put in place a true usage-based billing system that provides both accurate billing and a solution for value-added services.

The most common usage-based billing method used today is called the “95th percentile” calculation method. The 95th percentile calculation involves taking a series of usage samples for fixed time intervals over a billing period. For example, the throughput for a given customer may be collected over a 100-minute timeframe for a billing period of 30 days. At the end of the billing cycle, the 100-minute bandwidth usage calculations are sorted from highest to lowest, and the customer is billed for the bandwidth consumption that falls in the 95th percentile.

This method is both inherently inaccurate and unnecessarily complicated. The customer’s final bill is based on peak usage (as defined by the 95th percentile) instead of actual usage. In addition, the data that is collected to calculate the bill is not accurate; meaning that any disputes between the service provider and customer cannot be resolved easily. Finally, for value-added service billing, such cost-averaged usage-based billing is completely inadequate.

Riverstone’s LFAP-based system provides a far better foundation for true usage-based billing. Using RS routers, a service provider can gather accurate usage data and charge customers for the exact amount of usage, just like a long-distance telephone call. In addition, any use of value-added services can be captured and directly billed for, without requiring user action (i.e., the application collecting credit card information). For service providers looking to transition to usage-based billing in whole or in part, a reliable accounting system is a must.

## **IMPLEMENTING AN ACCOUNTING AND BILLING SYSTEM**

Implementing an accounting system does not necessarily mean a huge dollar investment. Riverstone has created a set of open source and shareware applications that create a robust billing system at a fraction of the cost of commercial applications. For customers who do not have the necessary resources in house, Riverstone works with systems integrators who can take the open source tools and package them into a billing system that will fit the needs of many service providers.


While certain service providers will eventually need the capabilities of a commercial application, Riverstone’s open source model allows service providers to implement or try out a usage-based billing model with a far smaller capital outlay than would be otherwise required. In addition, if a service provider’s billing system requirements ever outgrow the in-house or open source solution, the open source experience results in a better understanding of what the new OSS requirements might be.

Riverstone also provides a full range of commercial applications through our software partners. Riverstone has significant experience working with two companies in particular: Xacct, an IP mediation application provider, and Portal, a company that offers a bill presentment and collection application.



### The Four-Part Billing Framework

A simple framework can be applied when designing or evaluating an accounting and billing system. There are four basic functions that must be performed:



#### 1. IP Flow Data Collection (“Mediation”)

The industry term for the collection layer in the OSS is usually referred to as “Mediation.” Although the function seems simple enough, the assumption is that the network equipment is able to collect the data and transmit it reliably. Without this key feature, there is no data to account.

Riverstone provides a set of shareware and open source applications that can be used to create the mediation layer in-house or through a systems integrator:

- Slate is a light accounting server that is useful for demo and lab purposes to understand how Riverstone routers send LFAP data. The application is able to read accounting data from one network device and there is no data persistency. More information can be obtained at <http://www.nmops.org/README-slate-1.1>.
- Mica 3.0 is a full accounting server that has been implemented in a production environment in other service providers’ networks. This package is able to collect data from several RS network elements and provides data persistency. The current version runs on Solaris 2.6 and above. When running on a Solaris 2.6 machine with 256 MB RAM, the server can handle over 8,000 flows per second. Most service provider networks have traffic with less than 1,000 flows per second. This package is available to Riverstone customers only.

#### 2. Data Aggregation

Data aggregation involves taking the flow data and summarizing it to the appropriate level of detail that is required to perform the type of billing desired. In the voice world, the output of the aggregation process is called a Call Detail Record, or CDR. The level of aggregation depends on how services are identified in the network. The aggregation function is useful since it reduces the amount of data that is collected and stored.

If the service were just basic access, the appropriate level of summarization would be to produce a billing record that represents all traffic that originates from a particular IP address, subnet, or physical port. If an additional service such as storage, online application, or content is provided, the flow data can be summarized by port number, IP address, or protocol to reflect the particular customer’s usage information.



Riverstone provides an open source aggregation application that can be used as part of an in-house billing system.

- Basalt is an application that will aggregate and process LFAP flow record files. LFAP produces two flow information files. The first file contains the flow setup information. The second file contains the actual usage information, which is sent from the router to Mica in increments as low as every 5 minutes. Basalt is able to integrate the two files and aggregate the flow usage data. Like Mica, this package is available to Riverstone customers only.

### 3. Data Rating (Convert to Currency)

Once the flow data is collected, the data must then be converted to a dollar amount. The industry term for this function is “rating.” Rating can become quite complex depending on the revenue model being pursued. For example, possible billing plans could be usage, per use, per transaction. Although the billing plan may be different, the flow data collected is the same.

### 4. Present Bill Data to Customer

Once the data is collected into a central “billing” database and converted to a currency, it can be presented to the customer. The most straightforward method is to provide a total bill amount and collect the payment.

**CONCLUSION** As the Internet continues to grow and additional IP services are introduced, service providers will need to be prepared to account for the traffic and bill for the corresponding services. This will require network equipment that is capable of collecting the data and a reliable protocol to transmit the data. Riverstone’s products meet these requirements on both fronts. The hardware-based architecture allows accounting to be a feature that does not affect performance, and we have implemented an advanced flow accounting protocol that is superior to any existing solutions.

Flow accounting data allows service providers to achieve two critical business objectives. The first and most obvious one is to provide the capability to do reliable usage-based billing. The second objective is to allow service providers to contain their costs and expand their networks in line with their revenue objectives.

**ONLINE RESOURCES:** <http://www.nmops.org/>

Open source network management Website. The Website hosts a series of mailing lists that are general discussion lists for topics relevant to service providers: accounting, provisioning, and monitoring. The Riverstone “Platform SNMP Management Guide” details the SNMP implementation specific to the Riverstone product line. The MIBs for Riverstone’s RapidOS are also posted on the Website.

**Internet Protocol Flow eXport Working Group** <http://net.doit.wisc.edu/ipfx/>

Riverstone is actively involved in a new proposed IETF Working Group that is developing a standard IP flow export protocol. The new group, which will be known as IPFX, will be proposed at the 51st IETF meeting.

<http://www.sisd.com/freeside/>

Open source ISP billing and account management software.



## Acronyms

ACL	Access Control List
ANSI	American National Standards Institute
ASIC	Application-Specific Integrated Circuit
ASP	Application Service Provider
ATM	Asynchronous Transfer Mode
CBR	Constant Bit Rate
CWDM	Coarse Wave Division Multiplexing
DS1/DS3	Digital Signal, Level 1 (1.54 Mbps) or 3 (44.7 Mbps)
DSL	Digital Subscriber Line
DWDM	Dense Wave Division Multiplexing
DVMRP	Distance Vector Multicast Protocol
E1/E2	European Trunk 1/2 (2 Mbps/34.3 Mbps)
ERP	Enterprise Resource Planning
HSSI	High Speed Serial Interface
ISP	Internet Service Provider
ITU	International Telecommunications Union
LAN	Local Area Network
LEC	Local Exchange Carrier
MAC	Media Access Control
MAN	Metropolitan Area Network
MDU	Multiple Dwelling Unit
MLPPP	Multi Layer Point-to-Point Protocol
MPLS	Multiple Protocol Label Switching. See "MPLS in Metro IP Networks," <a href="http://www.riverstonenet.com/technology/mpls.shtml">http://www.riverstonenet.com/technology/mpls.shtml</a>
MTU	Multiple Tenant Unit
OC-3/OC-12	Optical Carrier 3/12 (155 Mbps/622 Mbps)
PDH	Plesiochronous Digital Hierarchy
PIM	Protocol Independent Multicast
POS	Packet over SONET
PPP	Point-to-Point Protocol
PVC	Private Virtual Circuit
QoS	Quality of Service
RED	Random Early Discard
SONET	Synchronous Optical NETWORK See <a href="http://www.techguide.com/comm/sec_html/sonet.shtml">http://www.techguide.com/comm/sec_html/sonet.shtml</a>
SLA	Service Level Agreement
SPE	Synchronous Payload Envelope
SRP	Spatial Reuse Protocol See RFC 2892
T1	Trunk 1 (1.544 Mbps)
TCP/IP	Transport Control Protocol/Internet Protocol
TDM	Time Division Multiplexing
UBR	Undefined Bit Rate
VBR	Variable Bit Rate
VLAN	Virtual LAN
VoD	Video on Demand
WAN	Wide Area Network
WDM	Wave Division Multiplexing
WRED	Weighted Random Early Discard



**River  
STONE  
NETWORKS™**

### **Riverstone Networks, Inc.**

5200 Great America Parkway, Santa Clara, CA 95054 USA

**877 / 778-9595 or 408 / 878-6500 or [www.riverstonenet.com](http://www.riverstonenet.com)**

© 2001 Riverstone Networks, Inc. All rights reserved. Riverstone Networks, RapidOS, and Enabling Service Provider Infrastructure are trademarks or service marks of Riverstone Networks, Inc. All other trademarks mentioned herein belong to their respective owners.