

Customer Success Story: WillitsOnline.Com

Issues with Generic Linux PPPoE Termination

August 2005



Making Broadband Manageable: Be Empowered.

© 2005 Fine Point Technologies, Inc. All rights reserved.

The information contained in this document represents the current view of Fine Point Technologies, Inc. on the issues discussed as of the date of publication. Because Fine Point Technologies must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Fine Point Technologies, and Fine Point cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. Fine Point Technologies Makes No Warranties, Express or Implied, In This Document.

Microsoft, Outlook, Windows, and Windows NT are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Other product and company names mentioned herein may be the trademarks of their respective owners.

Fine Point Technologies, Inc. • 139 Centre Street • New York, NY 10013 • USA

THE COMPANY

WillitsOnline.Com is a local provider of high-speed Internet access, hosting, and network integration services in Willits, CA. Founded in 2002,



WillitsOnline was the first and only local provider of high speed internet access in Willits, and currently boasts many satisfied customers. Focused on delivering 100% customer satisfaction to residential users switching from dialup to their first broadband experience, and business users who require guaranteed service and support levels, their service has steadily expanded and today covers a wide list of cities and areas.

WillitsOnline was founded by a computer engineer with over 22 years of experience in systems programming, network engineering, security, and high speed telecommunications. When available, WillitsOnline opts to use Linux and other free software in their operation to reduce costs.

THE NEED

The Linux source for PPPoE termination WillitsOnline originally selected was the Roaring Penguin RP-PPPoE, a user-mode PPPoE client for many Linux distributions. It originally took WillitsOnline an estimated 200 man-hours to setup a production-ready router solution based on free a Linux distribution that was capable of routing, terminating subscribers and authenticating from RADIUS. Despite this large time investment, WillitsOnline still experienced many issues with the free Linux solution that was jeopardizing the reliability of their business.

The issues that WillitsOnline experienced are explained in greater detail below,

Total Cost of Ownership (TCO)

WillitsOnline estimates that over 200 man hours were spent building, configuring and managing their home-grown Linux PPPoE server solution. Based on this data, it is clear that some of the common perceptions associated with platform management costs for Linux PPPoE server solution simply do not reflect reality.

“The amount of time and effort that I spent [creating my own Linux based PPPoE router] was ungodly and I wouldn’t wish it on my closest competitor.”

– Mike Ireton

Studies from the META Group Inc. show that a major source that consumes administrative time in Linux is hardware driver compatibility and lack of administration tools. In this area, Linux administrators spend more than 30% more time on hardware compatibility issues than they do “pre-boxed” solutions. Overall it is estimated that WillitsOnline investment more than \$20,000 in labor and \$600 in hardware costs for a grand total of \$20,600 to implement a home-grown Linux PPPoE router solution that can support only 250 subscribers.

User-Mode Implementation

The Linux operating system divides its use of the virtual address space into two distinct ranges for mapping. The user virtual address space (also called user space) maps the current user process and the kernel virtual address space (also called kernel space or system space) maps the operating system code and structures. User space occupies the lower range of the address space and kernel space occupies the upper range.

In accordance with this model, Linux utilizes two processor operating modes that are defined as “user-mode” and “kernel-mode”. Applications and services that run in user-mode have access only to the user virtual address space. When a user-mode process requires data that resides in the kernel virtual address space, it calls a system service to return the data. The Linux operating system and most drivers run in kernel mode and have access to operating system structures, to all system memory, and to all processor instructions. Kernel-mode components implicitly trust each other—that is, they assume that addresses and parameters they receive from other kernel-mode components are valid. Kernel-mode components can not, however, trust user-mode components. To safeguard system security, kernel-mode components must validate all data and addresses that are passed to them from user-mode.

The RP-PPPoE Linux software is based on this user-mode type architecture. For day to day desktop related use, user-mode will work perfectly fine, however as WillitsOnline had discovered, the user-mode implementation was not reliable for commercial use. Due to the number of issues that WillitsOnline was experiencing they were having difficulty scaling their service to the growing demand of subscribers.

“Hacking” to Achieve Functionality

If anything RP-PPPoE was not “ready out of the box”. WillitsOnline had to engage their various in-house Linux expertises to result in a usable solution. This endeavor began with working to get RP-PPPoE in to production level, with a self booting CD-ROM version, working out various bugs in the RP-PPPoE - Server source code, diagnosing a race condition problem with the user-mode PPPoE (more about this below) and “hacking” around in the Linux code all while in a production environment with customers online. The challenges continued with learning that RP-PPPoE would not work over any device in the server not named “ethxx” such as bridged devices. After sporting through the numerous issues, WillitsOnline then continued to complete the fine tuning required for a production environment.

Bandwidth Contention

Some of the concerns with scaling being reported by WillitsOnline and reported with user-mode implementations of PPPoE pertain to “bandwidth contention”, where two or more users attempt to transmit a data across the same “pipe” at the same time. With the Internet and its bandwidth intensive content, the idea of bandwidth contention is a common concept and happens regularly. For example, take the average cable home user. At 4 AM, a time when few are on the Internet, you may experience up to 3 or 4 mbps download speed. However, the same user at 7PM when everyone is home from work and the children are researching school projects would be lucky if to experience speeds faster than ISDN. Since the user-mode PPPoE is performing all its data conversion and handling in the “user space” it creates a contention issue with minimal users online.

To further investigate this issue that WillitsOnline was experiencing, Fine Point Networks put RP-PPPoE through our Spirent PPPoE Certification and found that if as few as 64 simultaneous subscribers were connected to the RP-PPPoE server, the throughput available to those subscribers was less than 23 mbps (or 360 kbps per subscriber) on a 1 gbps interface.

PPP Daemon Issues

WillitsOnline found that in their environment, which is a wireless network, they had a high occurrence of certain Customer Premise Equipment (CPE) devices which would not terminate cleanly when disconnected from the RP-PPPoE server. This would cause the RP-PPPoE server to not be notified of the CPE disconnect, and when the CPE attempted to reconnect it would fail or launch two PPP daemons for the same user which can cause CPU and memory issues.

WillitsOnline was also experiencing various race conditions in Linux’s Point-to-Point Protocol stack (PPPd) when used in conjunction with the RP-PPPoE’s PPPoE plug-in. These race conditions would cause various problems such as two or more PPPd processes being launched for the same user. This resulted in reducing the performance of the router by ½.

CPE Compatibility Issues

There also was a verifiable compatibility problem with the RP-PPPoE implementation and a certain brands of wireless routers which could not successfully initiate a PPPoE session with the server.

Load Balancing

“Some other things were that I was beginning to get concerned about not having any failover redundancy (and none of the free tools would provide this feature).

– Mike Ireton

When it comes to delivering reliable service, load balancing and redundant systems, especially on a mission critical device such as routers and network infrastructure are critical.

To understand the limitations for redundancy and failover that RP-PPPoE presents, first we need to look at the limitations in the PPPoE RFC protocol itself, PPPoE relies on a broadcast of “PADI” data packets to the PPPoE servers in the network. The first PPPoE server to response back is the server that then accepts that PPPoE session. This creates a network scenario where one (1) PPPoE server is potentially loaded to near capacity while any others are at a low capacity.

A single fully loaded PPPoE server creates a high probability for failure, and if the unit fails all subscribers experiences an interruption in service, driving calls to the call center. A fully loaded unit also creates bandwidth contention and packet dropping scenarios. WillitsOnline understood these issues and needed to address them, recognizing that the free code does not have this support the risk of a single point of failure was one that WillitsOnline could not afford.

Limited Protocol Support

RP-PPPoE has not been updated since October 2003. While limited support for Virtual LANs (VLANs) is included, due to limitation in the user-mode implementation, it can only support 64. However, the RFC 3069 calls for 4096.

There is very limited support for L2TP (RP-L2TP), which has been found to limit bandwidth to less than 60 mbps on a gigabit Ethernet interface as well as crash the PPPoE-Server daemon after approx 24 hours of use.

These limitations are only compounded by the limited support for domain “Realms” and ATM interfaces which are a general requirement for any DSL reseller or broadband aggregator.

Bandwidth Abuse

Several months after moving to a PPPoE based broadband service, WillitsOnline began to recognize that their subscribers we utilizing file sharing applications. This cause a tremendous bandwidth strain no the service and affected the SLA of the other subscribers. However, when trying to solve this problem, WillitsOnline simply could not allocate the resources:

“I also really needed to have some sort of PPPoE enforced bandwidth limits in addition to the controls I have on most, but not all, of my subscriber’s equipment, and implementing this was yet another distraction I didn’t have time for.”

– Mike Ireton

THE SOLUTION

The ServPoET[®] Broadband Management Server (BMS) platform provides, the industry-leading family of PPPoE subscriber management products that offload and terminate PPPoE sessions for all types of broadband service at nearly 1/3 the cost and 3x the performance. The ServPoET BMS[™] platform offers carriers and service providers worldwide a solution for aggregating and managing broadband PPPoE subscribers and value-added services.

Kernel-Mode Implementation

ServPoET[®] eliminates the standard problems of RP-PPPoE by implementing all data processing at the Kernel level. ServPoET[®] using a Kernel mode driver for both the PPPoE as well as the PPP processes. By doing this, ServPoET eliminates the throughput issues as well as the memory usage issues.

No Bandwidth Contention

ServPoET[®] can support over 6,000 simultaneous subscriber sessions on a 1U server and establish combined subscriber session traffic of up to 893 mbps on a gigabit Ethernet interface. This means that if you had 6,000 simultaneous subscribers logged in to only one (1) ServPoET BMS 800 (a 1U rack mount device) and each user was simultaneously downloading online they could each still achieve 152 kbps.

PPPoE Race Conditions

ServPoET is engineered as a commercial grade application. Using our PPPoE Certification labs, ServPoET goes through extensive testing, that spawns thousands subscriber “real world” sessions per second to place extreme load on the unit to confirm that such race conditions don’t exist. Due to the limited funding of freeware and the lack of carrier grade scope, RP-PPPoE has not able to be placed under such loads such as that described above. Freeware applications such as RP-PPPoE are created by a community that does not have access to the tools and manpower to perform stress testing and engineer their solutions for mission critical uptime and applications.

Load Balancing

One of the largest advantages to ServPoET[®] is not just an advantage over the native RP-PPPoE in Linux but an advantage over all our competitors’ products, and that is the patent-pending SmartCluster[™] Advanced Clustering Technology. SmartCluster[™] Advanced Clustering Technology enables multiple ServPoET[®] BMS units to automatically load-balance subscribers, ensuring optimal distribution of subscribers among clustered units for improved service levels. SmartCluster[™] technology evenly distributes sessions among cluster members and automatically removes members that fail or become unreachable. New cluster members can be hot-added to expand an existing cluster for additional capacity and performance. Administrators can scale their services with unprecedented time and cost savings.

Full Protocol Support

ServPoET[®] supports full VLAN as well as integrated Kernel level L2TP support. Also ServPoET[®] (BMS 500 and above) support an optional ATM interface for PPPoE over ATM. ServPoET[®] has full Realm support via prefix, suffix and Service name. The BMS 500 and above models also have full support for Access Lists, and OSPF.

ServPoET[®] also offers the ability to throttle bandwidth globally for all users via Realms, or on a user by user basis using our proprietary RADIUS attributes. These features allow a service provider to limit in kbps both the download and upload, per subscriber, independently. Allowing a service provider to offer tiered services, or implement a Wireless infrastructure and limit the bandwidth available per subscriber to reduce contention.

Ease of Use

All configurations for ServPoET are done in a simple graphical user interface, no knowledge of any operating system or command line configuration needs to be known to operate and maintain the system, the average Service Provider can have the system up and running, terminating PPPoE sessions in as little as 20 minutes.

“I really haven’t had to do anything to ServPoET beyond its initial configuration and that’s the way I like it.”

– Mike Ireton

Total Cost of Ownership

Companies such as WillitsOnline are looking to vendors for better documentation and best practices, especially in the Linux environment. Part of the value of better documentation and best practices will be seen in server deployment, which, on average, consumes an elapsed period of five business days and more than 25 staff hours for every single deployed server.

Although organizations tend to be most cognizant of acquisition costs because they are tangible, they often do not consider the largest cost component of total cost of ownership (TCO) — the “people costs” (i.e., the costs of IT staff members who provide ongoing maintenance and support, network administration, ability to read TCP dumps and RFCs side by side, C programming and shell scripting). However, companies such as WillitsOnline have recognized that the key to lowering overall TCO is to decrease these IT staffing costs by making administrators more efficient. The ServPoET BMS is one such solution.

So where is time being saved? A look at the use of tools available in the ServPoET BMS environment provides some insight. In the ServPoET BMS environment, WillitsOnline has put various tools to work to try to reduce task time requirements. The simplified configuration interface of ServPoET enabled WillitsOnline to reduce the elapsed time to get up and running to under one (1) hour. At a conservative estimated cost of \$100 per man-hour, WillitsOnline was able to setup the ServPoET BMS for \$11,701 less than the cost to implement the Linux solution.

“[With the RP-PPPoE], I was able to deal with the technical issues and produce the solution we were running with. The time and effort cost of doing so, however, was

anything but free, and I think it underscores one of the most important reasons why anyone considering a PPPoE deployment and thinking about using the free code or solutions built on free code, should consider ServPoET instead: Your time is better spent elsewhere!”

– Mike Ireton

CONCLUSION

ServPoET[®] is as a commercial-grade Broadband Remote Access Server (B-RAS) with the features and support needed to give Service Providers the ability to offer 99.99% uptime to their subscribers. The ServPoET[®] Broadband Management Server (BMS) platform provides the industry-leading family of PPPoE subscriber management products that offload and terminate PPPoE sessions for all types of broadband service at nearly 1/3 the cost and 3x the performance. The ServPoET[®] BMS[™] platform offers carriers and service providers worldwide a solution for aggregating and managing broadband PPPoE subscribers and value-added services.

REGIONAL LOCATIONS WORLDWIDE

United States (Corporate Headquarters)

139 Centre Street, 6th Floor
New York, NY 10013
USA
+1.212.962.7410
info@finepoint.com

United Kingdom

Globix House
1 Olivers Yard
London EC1 Y1HQ
UK
+ 44.2075.264818

France (Europe, Middle East, Africa)

Les Algorithmes, Bât. Aristote A
2000, Route des Lucioles, BP 29
06901 Sophia Antipolis
France
+ 33.1707.18418