

ts-PWLAN: A Value-add System for Providing Tiered Wireless Services in Public Hot-spots

Arup Acharya, Chatschik Bisdikian, Archan Misra
IBM Corp.
Thomas J. Watson Research Center
Hawthorne, NY 10532, USA
{arup, bisdik, archan}@us.ibm.com

Young-Bae Ko¹
School of Information and Computer Engineering
Ajou University
Suwon, South Korea.
Youngko@ajou.ac.kr

Abstract:

Access to data services via wireless LANs at private (e.g., corporations or home) and public “hot-spot” (e.g., hotels and airports) settings is becoming commonplace daily. Data access via (for profit) public wireless LAN (PWLAN) installations is typically based on user subscription and pre-configured services profiles pertaining primarily access to the global Internet. The goal of the ts-PWLAN¹ project is to define an architecture and a prototype implementation that enables the provision of premium and non-premium tiers of services to transient and non-transient users. ts-PWLAN provides for dynamic renegotiations of tier of services and enables various billing modes, e.g., based on connectivity time and usage, thus enabling service providers to increase their revenue opportunities via multiple service offerings.

I. INTRODUCTION

We have been witnessing the rapid deployment of IEEE 802.11 [4] based wireless LANs in a variety of public hot-spots, such as airports, hotels, internet cafés, and so on. Public hot-spots allow mobile users to access the backbone networks and associated services using their personal devices, such as notebook computers and PDAs. Current deployments of (for profit) PWLANs, e.g., [7], [9], provide only a single service, namely Internet access. While some solutions do offer differential pricing and bandwidth guarantees for Internet access, they require the user to have a subscription (provisioned off-line) with the service provider established and activated before use. While such a subscriber-based access model is appropriate in many contexts, such as wireline ISPs, we believe that, in the PWLAN arena, it suffers from two important and logically distinct drawbacks:

- The PWLAN marketplace is still fairly fragmented. Different hot-spots being serviced by different service-providers. An access mechanism that is solely based on pre-configured subscriber profiles is clearly counter-productive in such an environment, since it prevents customers of one WISP (wireless ISP) from exploiting the public access infrastructure of another WISP.
- Even if the multiple-provider problem is eventually solved, e.g., through the establishment of roaming agreements, such subscription-based access paradigms

do not allow users to *dynamically* change or modify their service levels. Allowing users the ability to obtain services normally outside their current profile is especially important in hot-spot scenarios, where users may access premium services in an impulsive manner.

In this paper, we present the architectural framework and the current implementation of *ts-PWLAN*, a solution for supporting dynamic and differentiated access to network services over a public access infrastructure. *ts-PWLAN* allows mobile users with no prior subscription to discover the different service tiers or choices available at the current public access infrastructure and then select their desired service tier and associated service duration on a pay-per-use basis. On the infrastructure side, *ts-PWLAN* allows the WISP to register new users via a Web-based interface and then perform *access control* to ensure that a user accesses only those services in the tier that she has selected. While access control can, in general, be performed at the link layer (e.g., at wireless access points) or the application layer (e.g., at individual Web servers), we prefer to perform access control at the network layer by establishing filtering rules at the access routers. As we shall show later, our access control mechanism does not require any significant modification to existing network components and is able to ensure appropriate access privilege for mobile users in a scalable manner.

The rest of the paper is organized as follows: Section 2 describes the architectural framework for our system. Section 3 describes a reference implementation of our system plus some performance results relative to our implementation. We conclude in section 4 with a summary and a reference to related work.

II. ARCHITECTURE OF THE *ts-PWLAN* SYSTEM

In this section, we shall present the *ts-PWLAN* architecture and its functional elements. We shall also explain how standard browser features, such as cookies and HTTP redirects, are used to enhance the *ts-PWLAN* functionality without requiring any additional modifications on the client device.

A. The *ts-PWLAN* Architectural Framework

The wireless LAN infrastructure typically consists of a collection of access points (AP), which provide customers wireless connectivity to the *ts-PWLAN* access infrastructure. The

¹ The abbreviation stands for “tiered services PWLAN”.

infrastructure itself consists of certain networking configuration services (such as DHCP and DNS) that are not subject to any form of access control and are freely available to any device equipped with a wireless LAN card and an appropriate IP stack. For *ts-PWLANS*, we further envision a variety of *local* services, as well as *global* services such as Internet access, all of which lie behind an intelligent *gateway* that regulates access to these services. The local services may include some free Web services (such as local weather or a directory of local restaurants and shops), which are available to all users and do not require any explicit user registration. Other local services, such as local video (e.g., servers for downloading special movies) or VoIP, can be considered to be premium services provided by the local WISP. In addition to these local services, the WISP may also provide various global-connectivity-related services, such as Internet access (with possibly different levels of pricing and associated QoS guarantees) and remote VPN access. One or more gateways are responsible for ensuring that these services are only accessible by people who have registered for these services (and they can be charged for).

The primary intelligence in the *ts-PWLAN* system resides in the access control server, which we refer to as the *registration server* (RS), see Figure 1. This server presents a Web-based registration menu to clients who wish to access services via the WISP. Whenever a user wishes to utilize the services offered by the WISP, the client device must interact with the RS and select from a list of offered services. The RS registration menu is available at a locally unique URL, e.g., http://www.public_WLAN.com. It is important to note that the registration-process is purely browser-based and does not require any modifications or other software or hardware utilities to be installed on the client device for this purpose. As part of this registration process, the client may need to supply various authentication or payment credentials (such as a credit card) to the RS, which may authenticate these credentials using techniques external to the *ts-PWLAN* system. Once a client selects a particular level of service, the RS will “tie” the IP address of the client device with the selected tier of services. Then, the RS issues the appropriate remote configuration commands to one or more controllable network elements (referred to as access gateways), which then set up appropriate packet filters. The RS server is also responsible for coordinating the user de-registration process. When the RS detects that a client has either left the *ts-PWLAN* system or that its current service duration has expired, it is responsible for issuing the appropriate set of commands to remove the corresponding filtering rules from the gateways and thereby prevent further access to the controlled resources.

The gateway device acts as the policy enforcer regulating user access to the services offered by the WISP. This gateway enforces policies on a per-packet basis—every incoming packet (from the WLAN) is inspected against the appropriate access control list. While conformant packets are simply for-

warded over the internal *ts-PWLAN* network, non-conformant packets are forwarded to the RS, which can then decide to alert the corresponding user of an attempt to access services outside the currently registered profile. As a network-layer device, the gateway performs filtering based on a combination of the source and destination IP address, the destination port and protocol (UDP/TCP). The gateway is also responsible for maintaining a count of resource usage (in terms of packet and byte count) per user—this can be used in post-paid billing scenarios where the user is billed in terms of actual resource usage. The various steps in the client registration and access control process under *ts-PWLAN* are shown in Figure 1.




Figure 1: Registration and service access in a *ts-PWLAN*

1) Dynamic Service Re-negotiation

The *ts-PWLAN* system uses the HTTP-redirection feature to provide registered users with an easy mechanism for dynamically upgrading their service. The access gateway performs per-packet filtering against a per-user access control list to ensure that a user is restricted only to the authorized set of services. If a packet from a registered user is found to violate the currently configured profile, this packet is simply forwarded to the RS. Since the RS is aware (from the source IP address on this packet) of the identity and service selection of the corresponding user, it is able to respond with a *customized* Web page indicating an unauthorized access attempt and indicating the service upgrade needed to gain access to the corresponding service. *ts-PWLAN* also allows any registered user to both upgrade or downgrade their current service selection at any point. *ts-PWLAN* users simply have to point their browser back to the Welcome page of the RS and suitably modify their service selection; the RS will appropriately modify the access control lists on the access gateways.

2) Service Termination

While designing the *ts-PWLAN* architecture, we were faced with two challenges. Firstly, we had to devise a mechanism by which we could track (with reasonable accuracy) when users would leave the *ts-PWLAN* network, since this would clearly determine the charges in a time-based billing scenario.

While *ts-PWLAN* may easily provide an active de-registration mechanism (where a user explicitly interacts with a de-registration page on the RS) to terminate the service, we realize that most *PWLAN* users would simply walk away from the hot-spot or shut down their access device. Secondly, due to the possibility of only intermittent connectivity and user movement between different network segment with a single hot-spot location, the *ts-PWLAN* infrastructure must be able to provide continued access to a selected service even if the mobile device changes, say, its IP address.

The problem of passive de-registration was solved through the use of DHCP leases of a configurable granularity (e.g., leases of the order of a few minutes) and the addition of notification messages between the DHCP server and the RS. Whenever a client device fails to renew its lease, *ts-PWLAN* considers the corresponding user to have (possibly temporarily) terminated the service. The DHCP server immediately notifies the RS, which then removes the appropriate access control filters at one or more access gateways. This mechanism allows *ts-PWLAN* not only to monitor (with reasonably fine granularity) the time when a user stops using the public access infrastructure, but also prevents a different user from gaining continued access to the service after the departure of the registered user.

3) *Session management*

We have chosen the use of Web cookies to solve several problems related to session management without requiring any specialized code on the client device. Whenever user selects a *ts-PWLAN* service offering at a hot-spot, the RS generates a unique cookie (valid for the service duration) which is stored on the client device. On the RS-side, the cookie is tied to a user profile created at registration time, that includes information pertinent to the user's current session. When the client device attempts to access the URL of the RS again (for functions such as modification or termination of a service), the browser will automatically insert this cookie information in the corresponding HTTP request. By validating this cookie, the RS is able to correlate any request to a specific user and accordingly customize its response.

4) *Billing and Accounting*

Differential billing is clearly one of the primary business objectives behind *ts-PWLAN*—the WISP should be able to provide each individual user with a variety of pricing and payment options. The access gateways, which perform per-packet access control, are also key elements in the accounting infrastructure—as all authorized packets must pass through these gateways, they are able to obtain packet-level usage statistics for each user. The accounting information gathered at this packet-level granularity is retrieved by the RS, which can then be provided to a third party accounting and billing system. *ts-PWLAN* also allows various forms of time-based billing—on the expiry of the currently negotiated service duration, the RS removes the appropriate access authorization lists from the access gateways. For continued use of the

WISP access network, the user must then re-negotiate a new service tier selection.

5) *Security Issues*

For a solution designed for public *WLAN* access, *ts-PWLAN* is distinguished by the lack of any *ts-PWLAN*-specific security mechanisms—this is primarily due to our desire to avoid the installation of any *ts-PWLAN*-specific component on the client. The *ts-PWLAN* system is seen as a value-add solution to existing or new public wireless (or wireline) LAN installations. Accordingly, *ts-PWLAN* does not seek to replace existing network elements, but rather coordinate their capabilities to support access to dynamically selected tiers of services. By focusing on the service layer, *ts-PWLAN* can be augmented with any link-layer security mechanisms (such as WEP [4] or EAP [3] or the emerging 802.1X standard [5]) that have been proposed specifically for wireless LANs. Also, it may leverage current and future registration protocols (e.g., PANA [10] from the IETF), or simply use standard secure HTTP (HTTPS) specifications implemented on all standard browsers. Moreover, all control messages exchanged between the internal RS elements, such as the RS and the access gateways, can be secured by standard IP-based encryption and authentication mechanisms.

III. PROTOTYPE IMPLEMENTATION AND EXPERIMENTAL EVALUATION

We have implemented a prototype of the *ts-PWLAN* public access infrastructure and deployed it on our testbed, shown in Figure 2. Any incoming client is configured with an IP address and other network parameters by the DHCP server. We have two different modes in which the client device obtains the URL of the *ts-PWLAN* RS. In the simplest mode, the user can be provided with an offline listing of the URL of the local RS (e.g., http://www.public_PWLAN.com) the user can then simply manually direct their browser to the RS's Welcome page, which lists the range of offered services. We have also implemented an alternative approach (for both Windows and Linux clients), where the URL of the RS is supplied as a DHCP option by the DHCP server. A simple client script then fires up the default browser application to this URL. Clearly, this alternative is not a required part of *ts-PWLAN* but it provides for an enhanced user experience (since the browser now pops up whenever the client device detects the availability of a *ts-PWLAN* service offering). As a third alternative, Web-page redirection can also be used to redirect unregistered users to the RS.




Figure 2: Prototype testbed layout

The registration process on the RS navigates users through a set of pages that list the service tiers and payment/pricing options available to the users. It also accepts and validates a user’s set of choices and the associated payment credentials. For a gateway, we use a notebook computer running Linux (configured with routing capabilities) and with two network interfaces. Access control is affected through the use of the *iptables* [6] code that allows the notebook to function as a router performing per-host routing. Whenever a new user registers with the RS, the RS server issues remote configuration commands (using the *rsh* command in our prototype) to the *iptables* daemon on the gateway, thereby setting up the appropriate packet filters. Upon termination of a user’s access session, the RS is able to retrieve the usage statistics (in terms of packet/byte counts) from the *iptables* daemon on the access gateway and timing details (by combining the registration time with the lease expiration time provided by the DHCP server). Our prototype implementation includes two service tiers represented by a “Gold” and a “Silver” Web server. A user that has selected the Silver tier is unable to access the Gold server, while a user that has selected the Gold tier may access both the Silver and Gold servers.

A. Experimental Results

Deploying *ts-PWLAN* on our testbed demonstrated the feasibility of our architecture. However, in an actual deployment additional issues like scalability are equally important. While a dedicated server and routing devices could be employed, the simplicity of our testbed bears the question as to what is the performance penalty (if any) paid using simple general purpose computers to perform per packet filtering in our Linux-based router, how the complexity of the routing rules affects the performance, and so on.

We performed a set of stress tests on our routing infrastructure. The test results reported in this section are based on experiments of measuring throughput performance between a

standard ftp server and clients. The server and the client reside on separate subnets and interconnected by our Linux router along their path. In particular, the system serving as the ftp server was a 700 MHz Pentium III notebook computer with 128 MB RAM; the system acting as the router was a 365 MHz Pentium II notebook computer with 128 MB RAM. These machines are connected each other via a 10/100 Mbits PCMCIA Ethernet cards and Ethernet switch. The systems that we used as ftp clients were a collection of various notebook computers: (a) a 400 MHz Pentium II notebook with 64 MB RAM; (b) a 233 MHz Pentium II notebook with 128 MB RAM; and (c) a 365 MHz Pentium II with 128 MB RAM. The client machines and the router are also interconnected directly using an Ethernet LAN. All systems were running Red Hat Linux ver. 7.1.

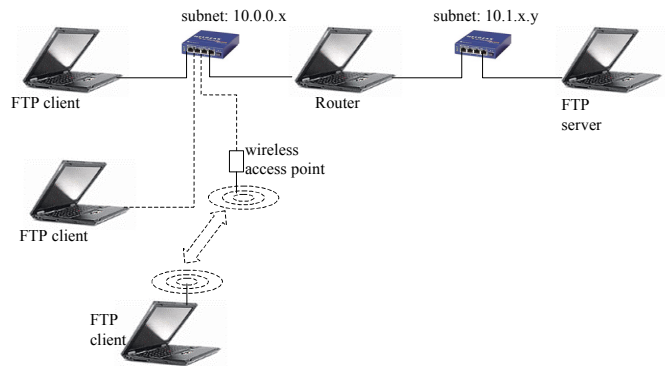


Figure 3: Set-up for performance testing

Figure 3 shows our test configuration that was used to establish a baseline throughput measurement; tests with both a single and multiple ftp clients were performed. Ftp throughput measurements were first taken between ftp server and client, being interconnected by a pure router (i.e., a router with zero packet filtering rules) in the middle. We ran the tests ten times and the average throughput result was about 24 Mb/s; this value has served as our baseline system bandwidth.

In our first experiment, we investigated the effect of the number of the IP packet filtering rules on the performance of the system. This test measures how scalable such a filtering-augmented router would be, by simulating the case where only one client’s flow is active while other multiple clients are inactive while, nevertheless, their presence impacts the decision time at the router. From this test, we found no noticeable impact on the router performance when the number of rules increases from 0 to 2000; the throughput did not fall below 23 Mb/s. These results indicate that any router configured with a number of packet filtering rules does not materially reduce the network performance, i.e., the filtering rules of our *ts-PWLAN* router may add very little overhead to the overall throughput.

Throughput (Mb/s)	1 wired client	2 wired clients	3 wired clients	3 wired (& 1 wireless) clients
per client	23.7	13.8	7.4	6.3 (4.0)
aggregate	23.7	27.6	22.2	22.9

Table 1: Throughput vs. number of ftp clients

Another experiment that we performed was to measure throughput when multiple clients are active and engaged in data transfer simultaneously, while varying the number of active clients from two to four. Although the actual number of clients varied during each run was limited by four, we believe that the results can still reveal any trends about the forwarding overhead when multiple active flows exist. The throughputs observed (per client and aggregate) are shown in Table 1. The last column shows the case with four clients, where one client is connected wirelessly through an IEEE 802.11b wireless link. It is worth noting here that even though the maximum link speed of an 802.11b LAN is 11 Mb/s, the effective speed after the communication protocol headers are removed is typically around 6 Mb/s.

IV. CONCLUDING REMARKS

While industry efforts focus on enabling basic access (like deployment and billing), we believe that dynamic service selection is a key differentiating feature for an attractive (for profit) PWLAN offering—users must have the capability to change their service selection to gain access to new services *on impulse*. In this paper, we have presented *ts-PWLAN* as a value-add proposition for enabling impulse provisioning of tiered services in PWLAN (hot-spot) deployments. *ts-PWLAN* facilitates the creation of new revenue streams to service providers (public space) property owners by creating an infrastructure for dynamic service offering differentiation. It provides dynamic access control to traffic streams from clients to services based on up-to-the minute user choices. Service providers and property owners can provide their own selection of premium (local) services allowing them to create customizable and personalizable service offerings with their own revenue streams. The *ts-PWLAN* system is based on open and existing standards, thus facilitating its easy incorporation to existing and new PWLAN installations without the need for modifications in client devices for their operation in *ts-PWLAN*-enabled installations.

In closing, the CHOICE architecture and system ([1], [2]²), one of the earliest implementations of an architecture for differentiated services over PWLANs, deserves special mention. The CHOICE architecture is very similar to *ts-PWLAN*: while

users interact with a network admission server (NAS) lying in the unprotected domain to choose their service tier and authenticate themselves, the access control is performed by the Traffic Control Gateway (TCG) on a per-packet basis using filtering rules provided by the NAS. Unlike *ts-PWLAN*, the CHOICE architecture does require the installation of a CHOICE client on the user device. This client is responsible for functions such as public WLAN detection (the NAS advertises the CHOICE network via periodic beacons), security (all packets are encrypted with a session key) and mobility management (the client is responsible for using the beacons to infer a change in the network connectivity and initiate re-registration), all of which are based on a protocol called PANS. In contrast, *ts-PWLAN* provides features such as mobility management and passive deregistration with no assumption other than the presence of a standard HTTP browser on the client device. The primary focus of *ts-PWLAN* is service selection and management allowing it to leverage any existing (or future) technologies used to control access to an PWLAN installation.

REFERENCES

- [1] V. Bahl, A. Balachandran, and S. Venkatachary, “The CHOICE Network: Broadband Wireless Internet Access in Public Places,” *Microsoft Technical Report*, no. MSR-TR-2000-21, Feb. 2000.
- [2] V. Bahl, W. Russel, Y.-M. Wang, A. Balachandran, G. Voelker, “PAWns: Satisfying the Need for Ubiquitous Secure Connectivity and Location Services,” *IEEE Wireless Communications Magazine*, Feb. 2002.
- [3] L. Blunk and J. Vollbrecht, “PPP Extensible Authentication Protocol (EAP),” *IETF*, RFC 2284, March 1998.
- [4] IEEE Computer Society. 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1999.
- [5] IEEE Computer Society, Standard for Local and Metropolitan Area Networks: Standard for Port-Based Network Access Control, IEEE Draft P802.1X/D11, March 2001.
- [6] netfilter/iptables Linux code, <http://www.iptables.org>
- [7] T-Mobile Wireless Broadband Network (formerly MobileStar Corporation), <http://www.tmobilebroadband.com/>
- [8] D. L. Wasley, “Authenticating Aperiodic Connections to the Campus Network”, White Paper, http://www.ucop.edu/irc/wp/wp_Reports/wpr005/wpr005_Wasley.html
- [9] Wayport Inc., <http://www.wayport.com>
- [10] A. Yegin, et al., “Protocol for Carrying Authentication for Network Access (PANA): Requirements and Terminology,” *IETF Draft*, Work in Progress, draft-ietf-pana-requirements-01.txt, March 2002.

¹ This work was performed while Dr. Ko was with IBM’s T. J. Watson Research Center.

² Additional references for research activities in this area can be found therein.