

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

Arup Acharya, Chatschik Bisdikian, Young-Bae Ko, Archan Misra

IBM Corp.

Thomas J. Watson Research Center

Hawthorne, NY 10532, USA

{arup, bisdik, youngko, archan}@us.ibm.com

Abstract:

Access to data services through wireless means gains is becoming commonplace daily. Wireless LANs are deployed in both private (e.g., corporate or home) and in public “hot-spot” (e.g., hotels and airports) settings. Accessing data services 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 ts-PWLAN project is to define a broad architecture and provide a prototype implementation that enables the provision of premium and non-premium tiers of services to transient users, provides for dynamic renegotiation of a tier of service and enables various billing modes based on connectivity time, usage, and so on, thus enabling service providers to increase their revenue opportunities through multiple service offerings to transient and non-transient users.

1 INTRODUCTION

We have been witnessing the rapid deployment of IEEE 802.11 [IEEE802.11] based wireless LANs in a variety of public *hot-spots*, such as airports, hotels, internet cafés, and so on. PWLANs provide broadband access at low cost and support the needs for an ever increasing mobile workforce. Public access services based on wireless LAN technologies can be viewed as either complementary or competitive to upcoming 3G cellular services. Public hot-spots allow mobile users to access the backbone networks and associated services from their personal devices, such as laptop computers and PDAs. Early deployments of public WLAN access solutions, e.g., [MobileStar], [Wayport], typically provide only a *single* service, namely Internet connectivity. 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 it suffers from two important, and logically distinct drawbacks, in the public wireless LAN (WLAN) arena:

- ? The wireless LAN marketplace is still fairly fragmented, with different hot-spots being serviced by different service-providers (ranging from the hot-spot owner to local Wireless ISPs (WISPs) to large WLAN aggregators). 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 from exploiting the public access infrastructure of other WISPs.
- ? 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. (For example, a mobile user with a

voice-only service profile may desire to download a video while waiting at the airport terminal for a delayed flight).

In this paper, we present the architectural framework and the current implementation of *ts-PWLAN*¹, 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 (HTTP) 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) as well, 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, even if they change either their network adapter or IP address.

In a later section, we shall survey and discuss several alternative proposals and architectures (both research prototypes and commercial products) aimed at providers of public access networks. We shall see that the *ts-PWLAN* solution is, however, unique in its support of all of the following three features:

- a) *Client Independence*: We do not require the user client devices to possess any additional software, other than a standard Web browser. All user interaction with the *ts-PWLAN* wired infrastructure at any location is through the use of standard HTTP-based mechanisms, such as cookies, and does not require the download of any proprietary components on the user's device. *ts-PWLAN* can thus be provided as a value-add service to existing PWLAN installations, with no additional client side modifications.
- b) *On-the-fly Service Re-negotiation*: Besides allowing users to specify their desired tier or set of services during initial sign-on, the *ts-PWLAN* infrastructure also permits users to dynamically upgrade or reduce their tier of service during their service duration. Standard HTTP redirection techniques are used to alert users who attempt to access services beyond their current profile and provide them the option of upgrading their service contract to the necessary level.
- c) *Intra-hot-spot Roaming Support*: Intermittent disconnection is a standard problem with wireless devices in hot-spot environments; as users move within a single hot-spot, they may lose connectivity to the infrastructure. The *ts-PWLAN* solution is able to preserve and migrate an existing service profile even when the mobile device re-connects with a new IP address or via a different network adapter (such as a switch from a WLAN interface to an Ethernet interface).

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.

2 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. Before describing the specific details of our current *ts-PWLAN* solution and its current implementation, we present an overview of the basic functions needed for offering tiered services in a public hot-spot. This discussion will also present the various design alternatives to *ts-PWLAN* and motivate our choices for the *ts-PWLAN* framework.

¹ The name is an abbreviation "Tiered Services Public Wireless LAN".

2.1 Generic Functions for Differentiated Hot-Spot Public Access

Figure 1 shows the various distinct steps involved in the process of public hot-spot access and lists the various possible alternatives. We consider these steps rather independent of each other, with each one been viewed on its own right. This “layered” view of accessing public services enables us to propose our system not as much as a substitute or a replacement of existing or planned installations, but rather as a value-add feature that enhances the capabilities of current and future WLAN installations.

Basic device *configuration* is clearly the first step in obtaining mobile access to networked services via a public WLAN infrastructure. This process involves functions such as detecting the availability of wireless access points and subsequently using configuration protocols such as DHCP to obtain parameters such as a host IP address and addresses of the next-hop gateway and the DNS server. Such functionality is typically present in all laptops and PDAs today.

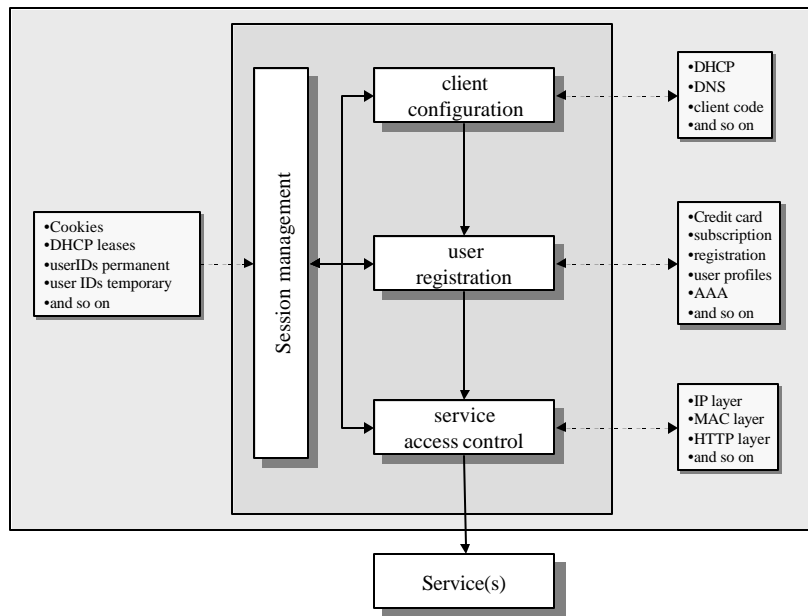


Figure 1: Functional Components of a Generic Public Hot-Spot Solution.

By itself, basic host configuration does not provide the network with the identity of the mobile user, nor does it clarify the level of service that the user is interested in. In a hot-spot environment, where different providers may support different service levels and access charges, a separate *registration* process is clearly needed to: a) let the network inform the client device of the list of available service choices; and b) let the client device subsequently inform the network of its choice and provide any supporting authentication and accounting parameters (such as a credit card). To perform this registration, the client device must be provided with the identity of the corresponding network registration entity (registration server), and a protocol for exchanging such registration credentials must be specified. While various groups are working on such a stand-alone registration protocol (e.g., PANA [Yegin:01] from the IETF), current client devices do not usually come pre-configured with any standard registration protocol. It bears saying that subscription-based access is simply an instantiation of this registration function, since it essentially involves the exchange of pre-configured authorization parameters between any client and the access network.

The user registration process must clearly be followed by a separate *access control mechanism*, which ensures that registered users are able to access only those services that were negotiated (and possibly paid for) during the registration process. Such enforcement involves the establishment of filtering rules at some

ingress-networking element; non-conforming packets can then be dropped at this element. Possible alternatives to our choice of network-layer enforcement include the establishment of filtering rules at the wireless access points or at the individual application servers. Our experience with commercial wireless access points showed that most of them currently perform all-or-nothing access control and do not possess the finer-grained filtering capability needed to allow users to access only a selected set of services (in terms of IP addresses, port numbers, protocols etc.); this however may change in the future. The alternative approach of configuring each service end-point individually is not very scalable. Furthermore, since the access point infrastructure is itself publicly accessible, this leaves the access network open to a barrage of potentially unauthorized packets traversing the network freely prior to blocking them at the end-point; denial-of-service attacks can become extremely easy.

In addition to the above functional components, a public access solution also needs to support dynamic profile re-negotiation as well as profile maintenance/migration for mobile nodes. Dynamic profile re-negotiation allows an already registered user to upgrade or degrade their selected service set during an active service duration. Profile maintenance/migration is another key feature in the public WLAN environment, since it is well known that users can experience intermittent disconnections in such wireless environments. Moreover, it is also possible for a user to change access points and IP subnets during the service duration. A well-designed public hot-spot access architecture should allow a registered user to maintain access to their configured services, even if the user changes access points or IP subnets within the access network.

2.2 *The ts-PWLAN Architectural Framework*

Figure 2 illustrates the logical layout of elements in 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 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. The *ts-PWLAN* vision also envisions 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 those services (and possibly paid for them).

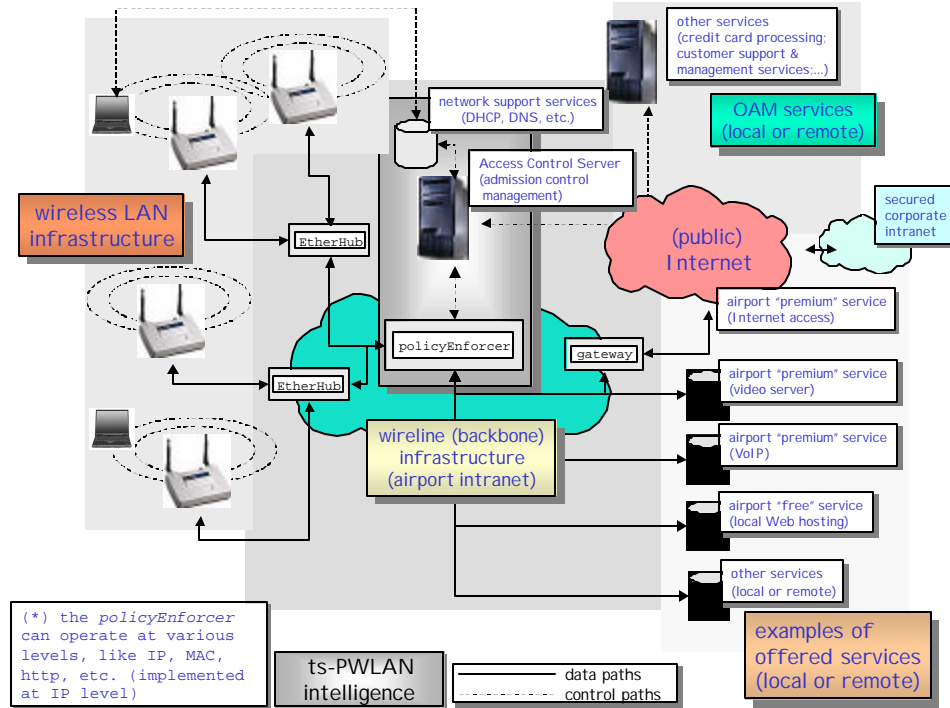


Figure 2: Functional Elements of the *ts-PWLAN* Architecture

The primary intelligence in the *ts-PWLAN* system resides in the Access Control Server (ACS), more commonly referred to as the *Registration Server (RS)*. 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-- this URL need not be the same in different hot-spots. 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. 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 and then issue the appropriate remote configuration commands to one or more access control gateways, which then set up appropriate packet filters. The RS server is also responsible for coordinating the user de-registration process. When the registration server 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.

As explained earlier, 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 forwarded over the internal *ts-PWLAN* network, non-conformant packets are forwarded to the Registration Server, 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 3.

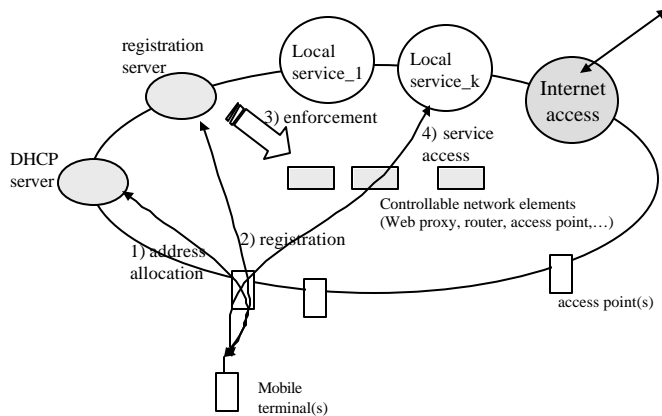


Figure 3: The *ts-PWLAN* steps for registration and service access

2.2.1 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. As explained earlier, 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 Registration Server. 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.² *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 direct their browser back towards 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. We shall shortly see how the use of cookies allows the RS to present a customized Welcome page for currently registered users who wish to modify their service selection. We believe that dynamic service selection is a key feature of any well-designed public hot-spot solution—users must have the capability to change their service selection to gain access to new services *on impulse*.

2.2.2 Service Termination and Mobility Management

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

² In principle the means of providing feedback to the user may vary depending on user desires and system capabilities. In our implementation, we use Web browsing applications, so the feedback is provided through a Web page.

intermittent connectivity and user movement between different subnets with a single hot-spot, the *ts-PWLAN* infrastructure must be able to provide continued access to a selected service even if the mobile device changes its MAC (e.g., switch from one WLAN card to another or use an Ethernet card instead³) or IP address (e.g., new IP address from different DHCP server).

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. If the disconnection is only temporary, we shall show shortly see how the use of cookies allows the registered user to regain access; the cookie mechanism solves the mobility problem as well.

2.2.3 Cookies and State Maintenance

Since the interaction between the client device and the *ts-PWLAN* registration and authentication infrastructure is HTTP-based, we use cookies to solve several problems related to mobility and intermittent connectivity without requiring any 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 registration server (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. The use of cookies also allows the *ts-PWLAN* system to handle configuration changes in a client device (such as the change in the IP address). Since a change in the host IP address causes an access control failure at the access gateway (where the filtering rules are based on the old IP address), the client browser is then re-directed to the RS. Since the Web page re-direction to the RS is actually achieved by "forcing" the client browser to issue a request to visit the Web page of the RS, the request will contain the cookie provided by the RS during the initial registration. Hence, the RS will be able to welcome back the user by associating the client device with an existing service profile. Note that this application-layer mobility management mechanism allows the *ts-PWLAN* user to maintain their existing service profile even if they switch their access media (such as a change from a WLAN card to an Ethernet one) or their IP address.

2.2.4 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 Gateway, which performs per-packet access control, is the key element in the accounting infrastructure—as all authorized packets must pass through this gateway, it is able to obtain packet-level usage statistics for each user. The accounting information gathered at this packet-level granularity is retrieved by the Registration Server, 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 gateway. For continued use of the WISP access network, the user must then re-negotiate a new service tier selection; we have already seen how *ts-PWLAN*'s Web-page redirection mechanism allows *ts-PWLAN* to accurately monitor the connectivity of an intermittently-connected or roaming user within a single service selection.

2.2.5 Security Issues

For a solution designed for public WLAN access, *ts-PWLAN* is distinguished by the lack of any *ts-PWLAN*-

³ It should be obvious from the presentation, that the proposed system may apply equally well to either wireline or wireless LAN technologies.

specific security mechanisms—this is primarily due to our desire to avoid the installation of any *ts-PWLAN*-specific component on the client. As mentioned earlier, 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. Thus, *ts-PWLAN* can be augmented with any link-layer security mechanisms (such as WEP [IEEE802.11] or EAP [EAP] or the emerging 802.1X standard [IEEE802.1x]) that have been proposed specifically for wireless LANs. Client devices perform the registration with the *ts-PWLAN* RS using the 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.

3 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 4. Any incoming client is configured with an IP address and other network parameters by the DHCP server, which allocates addresses in the range: 10.0.0.[5->250]. We have two different modes in which the client device obtains the URL of the *ts-PWLAN* Registration Server. In the simplest mode, which does not require any client code installation, the user can be provided with an offline listing of the URL of the local registration server (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 registration server is supplied as a DHCP option (option 96) by the DHCP server. A simple client script then fires up the default browser application to this URL. While this approach does require the installation of some scripting code on the client device, it is purely optional and merely serves to illustrate fancier client experiences (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 registration server.

The registration process on the Registration Server is written as a set of Java servlets, that navigate the user through a set of pages listing various services and the associated pricing options and that accept and validate the user's set of choices and the associated payment credentials (currently a credit card number). As a gateway, we use a laptop computer running Linux (configured with routing capabilities) and with two interfaces. Access control is affected through the use of the *iptables* [IPTABLES] code that allows the laptop 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 (as shown in Figure 4). 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 services, namely Gold and Silver: a Silver user is unable to access the Gold Server, while a user selecting the Gold service obtains access to both the Silver and Gold servers.

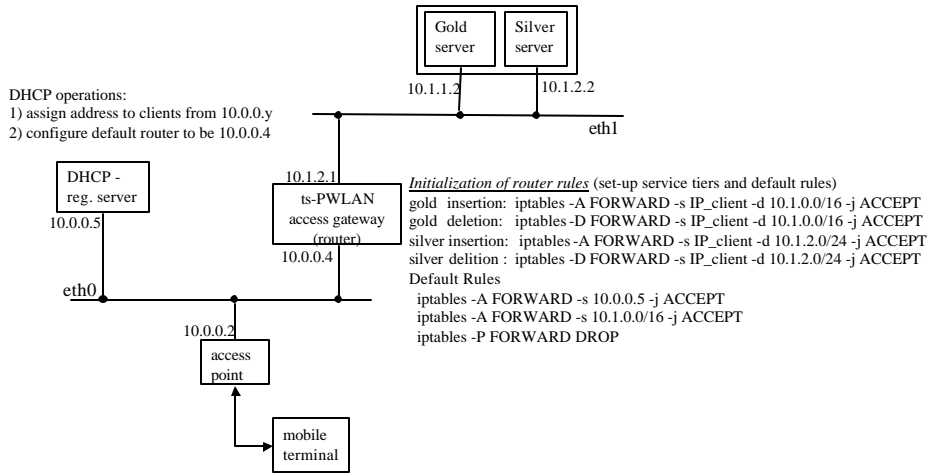


Figure 4: Prototype Testbed Layout and Filtering Rules

3.1 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 paid (if any) 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

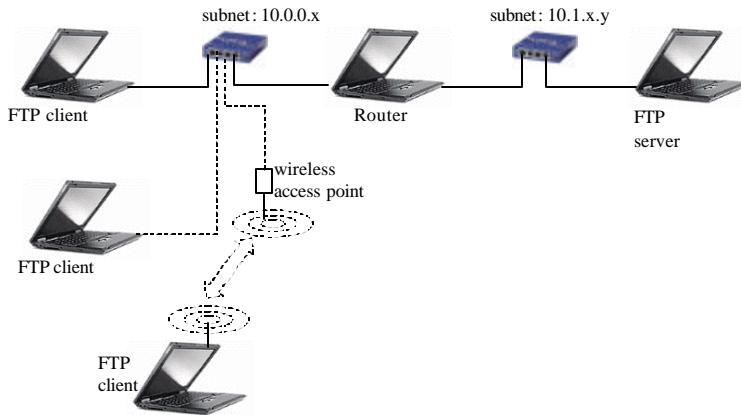


Figure 5: Set-up for performance testing

Figure 5 above shows our test configuration used to establish baseline throughput measurement; test with both a single and multiple FTP clients were performed. Thus, 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 Mbits per second; this value will serve as our baseline system bandwidth. The somewhat low value is attributed on the packet processing capabilities of our router that is implemented on a general purpose notebook computer⁴.

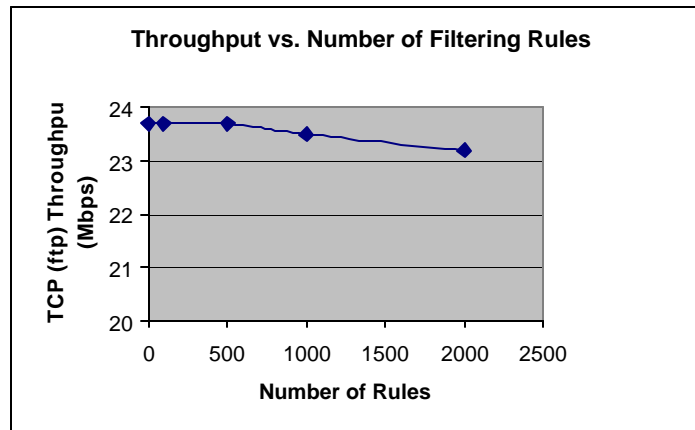


Figure 6: Throughput vs. number of filtering rules

⁴ In order to verify our argument here, we also performed throughput measurement tests between FTP server/client hosts residing on the same subnet and therefore in the absence of the intermediate router. We noticed that the average throughput went up to about 80 Mbps.

Figure 6 presents the measured throughput, averaged over ten runs, as a function of the number of IP packet filtering rules created on the router. Basically, this study can demonstrate 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. Observe that, from the plot in the figure, there is no noticeable impact on the router performance when the number of rules increases from 0 to 2000. 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.

Table 1: Throughput vs. Number of FTP clients

	With 1 wired Client	With 2 wired Clients	With 3 wired clients	With 3 wired/ 1 wireless clients
Per client Throughput (Mbps)	23.7	13.8	7.4	6.3/4.0
Aggregate Throughput (Mbps)	23.7	27.6	22.2	22.9

Another experiment that we performed was to measure throughput when multiple clients are being active and engaged in data transfer simultaneously, varying the number of active clients from two to four. Although the actual number of clients varied during each run were limited by four, we believe that the results can give a hint about the forwarding overhead while multiple active flows exist. The average throughput per connection is shown in Table 1. For two active clients, the average throughput per flow went down to about 13.8 Mbps and for three simultaneous connections to about 7.4 Mbps. We also ran the test with four clients, but here we made one client be connected wirelessly using the 11 Mbps IEEE 802.11 wireless card and access point. We found that the average throughput for three wired connections went down to about 6.3 Mbps and the one for the wireless connection was around 4 Mbps. It is worth noting here that even though the maximum link speed of an 802.11b LAN is 11 Mbps, the effective speed after the communication protocol headers are removed is typically between 6 and 7 Mbps.

4 CONCLUDING REMARKS

We have presented *ts-PWLAN*, a value-add system for enabling impulse provisioning of tiered services in (for profit) public wireless LAN (hot-spot) deployments. While industry efforts focus on basic access enablement (like deployment and billing), *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.

Among the various research prototypes and products designed for public hot-spot access related to this work, two deserve special attention for their similarity to the *ts-PWLAN* architecture. The CHOICE architecture and system from Microsoft Research [Bahl:00][Bahl:02]⁵ is, one of the earliest well-known implementations of an architecture for differentiated services over public WLANs. The CHOICE architecture is very similar to *ts-PWLAN*: while users interact with a Network Admission Server (NAS)

⁵ Additional references for research activities in this area can be found therein.

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 RAS 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.

Cisco's BroadBand Service Manager (BBSM) [BBSM] is a product specifically developed for user registration and profile management in public access broadband networks (such as WLANs, cable networks etc.) The process of user registration is very similar to that of *ts-PWLAN*: new users are directed to a Welcome page on the registration (Web) server to select the appropriate tier of service. BBSM also employs a per-user, per-packet access control mechanism to ensure authorized access. However, unlike *ts-PWLAN*, BBSM's registration server acts as a transparent proxy, so that client browsers are unaware of any packet redirection. Accordingly, the current BBSM solution does not exploit the use of HTTP cookies for service maintenance and is unable to handle changes in the client's attachment media (e.g., WLAN to Ethernet) or IP address. Moreover, the BBSM architecture suffers from a potential scalability issue since it combines both the registration and enforcement functionality in a single box. In contrast, both *ts-PWLAN* and CHOICE decouple registration from enforcement; the capacity of the system can be extended simply by having a single RS/ NAS can issue configuration requests to multiple access gateways/TCG.

References

[Bahl:00] 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.

[Bahl:02] 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.

[BBSM] "Cisco Building Broadband Service Manager: Technical Overview", *White paper*, Oct. 2001.

[EAP] L. Blunk and J. Vollbrecht, PPP Extensible Authentication Protocol (EAP), IETF, RFC 2284, March 1998.

[IEE802.11] IEEE Computer Society. *802.11 : Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1999.

[IEE802.1x] IEEE Computer Society, Standard for Local and Metropolitan Area Networks, : Standard for Port-Based Network Access Control, IEEE Draft P802.1X/D11, March 2001.

[IPTABLES] netfilter/iptables Linux code, <http://www.iptables.org>

[MobileStar] T-Mobile Wireless Broadband Network (formerly MobileStar Corporation), <http://www.tmobilebroadband.com/>

[Wasley:96] D. L. Wasley, "Authenticating Aperiodic Connections to the Campus Network", White Paper, http://www.ucop.edu/irc/wp/wp_Reports/wpr005/wpr005_Wasley.html

[Wayport] Wayport Inc., <http://www.wayport.com>

[Yegin:01] 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.