



SANS Institute

Information Security Reading Room

Preventing the fraudulent use of Internet DSL accesses by dial-up accounts: a network authentication issue.

Bruno Germain

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Bruno Germain, ccie
GSEC practical assignment version 1.4b – Research paper option
September 4th 2003

Preventing the fraudulent use of Internet DSL accesses by dial-up accounts: a network authentication issue.

Abstract / Introduction:

Access to the Internet by the masses has migrated from dial-up only connections to a combination of the former and so-called “broadband accesses”, typically Cable-Modem or Digital Subscriber Line (DSL). While most Internet Service Providers (ISP) now offer these services, many of them do not own the underlying telecommunications infrastructure and act as resellers to a larger provider, often a phone company. A consequence however of this network arrangement is the split between two distinct parties of the tasks required to allow a subscriber onto the Internet: the DSL provider who is responsible for the physical access and the ISP who is responsible for the authentication or “logical access”. This situation creates an opportunity for a fraudulent usage of the service and consequently some revenue losses for the ISPs.

This document will first describe such a scenario to put the following descriptions into context. Then we will look at the details of a typical deployment between DSL providers and ISPs in order to highlight the areas of vulnerability of the model. Finally we will suggest an approach to prevent this type of fraud with some other elements that could lead to tailored solutions for the ISPs: as the next sections will demonstrate, a unique overall solution is most unlikely given the number of ways each ISP could deploy its services.

We assume the reader has a general understanding of DSL and RADIUS technologies.

A scenario for a fraudulent usage of DSL access to the Internet

Consider the following points:

- ISP_A and ISP_B are both reselling the DSL services from the local phone company.
- ISP_A has many DSL service offerings, all based on a maximum amount of uploaded / downloaded bytes after which the subscriber pays for each megabytes extra.

For example, the basic package could offer 5 Gb upload and 10 Gb download per month after which the subscriber pays \$5.00 per extra Mb.

- ISP_B also offers dial-up account. These accounts are very cheap compared to the price of its DSL offerings. The prices of the dial-up packages vary according to the amount of time the user is connected in a month.

Since its dial-up offering is time-based and because of their low speed, ISP_B only monitors the cumulative connection time for dial-up accounts.

- User John Doe is a client of ISP_A and subscribes to the basic DSL service we described previously.
- User John Doe also subscribes to the “unlimited time” dial-up service of ISP_B for a fraction of the price he pays for a DSL service.

John Doe uses the DSL services of ISP_A for everything except his peer-to-peer download sessions as it would make him exceed his download limit.

So John reconfigures at night his DSL access device so that the PPPoE client now logs in with “username@ISP_B”. Doing so, John now has DSL access to the Internet via ISP_B’s network.

Since ISP_B is only monitoring its dial-up accounts for time usage, no alarms will be triggered and John can download all he wants without being bothered.

We realize that if John Doe was a hacker, access to any valid user account from ISP_B could possibly do the trick. However this research paper is not about how to secure passwords and user accounts but rather wants to look into why a user subscribing to a DSL offering from one ISP can use the same service from another ISP while subscribing to its dial-up service.

In order to understand how we can control this situation we will now look into the DSL wholesale environment and see why such a fraud can be so easily achieved.

Description of the DSL wholesale model

From a service perspective, the DSL wholesale model needs to answer three basic questions in order to work:

1. How will the Layer1 (physical network) to Layer2 (data link) or L3 (network) aggregation be done by the DSL provider?
2. How will the DSL provider support multiple ISPs and consequently identify and forward the individual subscribers to the proper one?
3. How will the ISPs authenticate and manage their subscribers?

To solve the first point, the DSL provider could potentially dedicate DSLAMs to each ISP it serves. This would hardwire each DSL subscriber to its proper ISP thus solving the issue. It would however have little practical and financial sense. The solution will therefore have to use technology residing higher in the protocol stack.

What is generally done is illustrated by Figure 1: all the physical connections are terminated on a common DSLAM, independently of the ISP to which the end user is subscribing, and the clients run a PPPoE client which can be aggregated further down the network path¹.

In the scenario illustrated, the CPE device is also the PPPoE client. This was made so to simplify the discussion but it is not unusual to see the PPPoE client reside on a device behind the CPE such as on the subscriber's computer. For a detailed look at the PPPoE specifications please refer to (Mamakos)

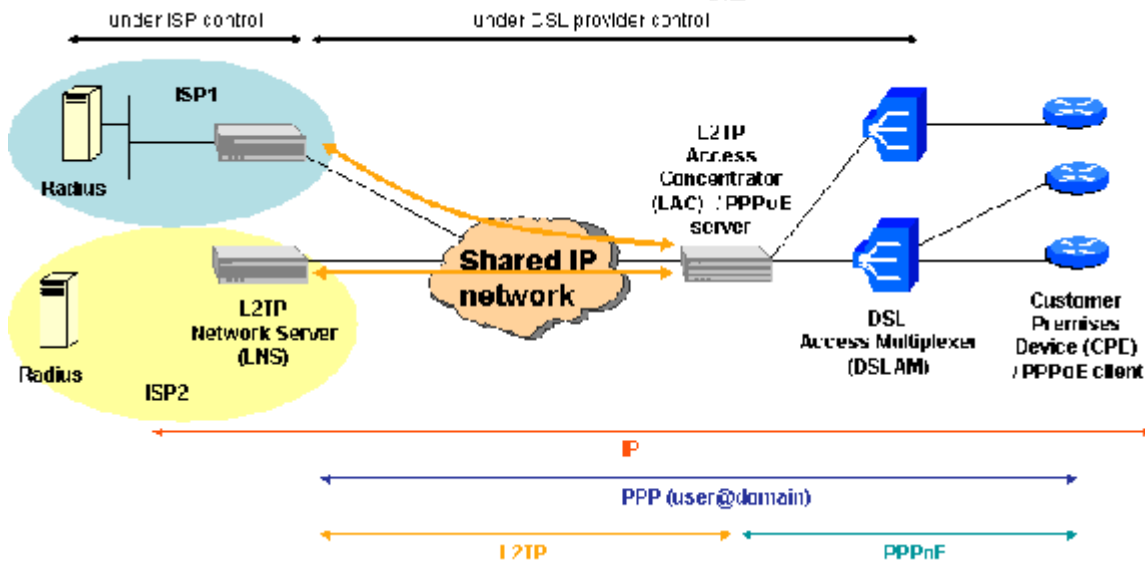


Figure 1: DSL wholesale configuration

In short the PPPoE technology allows the Point-to-Point Protocol (PPP) to work over a multi-access network by encapsulating it in Ethernet frames, which is one of the supported frame type for ATM / DSL technology as defined per RFC 2684.

So from a DSL provider's perspective, PPPoE assures the communication between multiple CPE devices and one or more PPPoE servers, thus solving the Layer 1 to Layer 2 aggregation problem.

¹ DSL lines can also be sent to the ISP on a one subscriber – one connection basis. This configuration is often used for “business” DSL subscribers since they have guaranteed bandwidth all the way up to the ISP, avoiding the Layer 1 to Layer 2 aggregator and the oversubscription normally done at that level. This arrangement is not covered in this paper as the network setup differs from the wholesale model we are describing.

Now the DSL provider needs to “switch” each subscriber to its proper ISP. This problem of sorting out multiple PPP sessions sharing a single access point but destined for multiple and mutually exclusive end points has seen numerous approaches put forth over time: in the same period PPPoE became adopted, the Layer-2 Tunneling Protocol (Townshley) was proposed and became the preferred method to achieve this goal.

In a nutshell, L2TP extends the reachability of a Network Access Server (NAS), typically found in dial-up networks to terminate PPP sessions, by breaking it in two server functions linked via an IP tunnel that can span multiple physical networks.

These components are²:

- The L2TP Network Server (LNS): the server functions under the ISP’s control which terminates the PPP sessions. It can maintain associations with multiple LACs.
- The L2TP Access Concentrator (LAC): the server function that is responsible to aggregate the end-users connections and for tunneling the traffic between the LNS and the end-users.

As illustrated in Figure 1, due to its role in the network, the LAC is often combined with other aggregation functions such as an IPsec gateway or a PPPoE server component.

- The L2TP tunnel itself: at least one tunnel must exist between the LNS and the LACs it supports. It carries the end-user’s PPP sessions between the LAC and the LNS. The tunnel is routable.

Since the LAC is the element which is responsible for sorting out the incoming PPP sessions once the Ethernet portion of the PPPoE packet has been stripped away, we need to understand how it makes its choice to switch them to the proper tunnel onward to the receiving ISP.

Two methods are possible for the LAC to make its decision:

1. Using a local database related to the domain name. In this approach, a static configuration links the domain portion of the fully qualified username (i.e. the @domain portion) to specific LNS entries.
2. Using RADIUS authentication. In this configuration, the LAC will ask a RADIUS server for authentication. This RADIUS server will acknowledge

² For a full discussion of PPP and L2TP, refer to (Black) which presents all the necessary information in the same textbook.

the request with the tunnel parameters. Then the LAC will dynamically use these values to tunnel the PPP session to the appropriate LNS.

The local database approach is by far the most widely deployed because it prevents the DSL providers from having to be responsible, in part or totally, for authenticating the end user.

In the RADIUS approach, the deployment faces a trust problem: ISPs don't like sharing their end users information with the DSL providers. So what ends up being deployed in this case is a form of RADIUS-proxy to the ISP's RADIUS³.

At this point in our description of the wholesale model, provided the PPP sessions have been authenticated, they would terminate on the proper LNS and the end user IP traffic would be forwarded to its final destination via the ISP network.

Moving to the LNS, we can see that it receives the traffic from one or many LACs and is seen as a NAS by the RADIUS server.

While we focus our discussion on DSL aggregation, L2TP and thus the LNS are also widely deployed for dial-up services, especially with ISPs covering large geographical areas, the Remote Access Servers (RAS) being the LAC in that case.

The ability to aggregate multiple access technologies through the LNS is an important feature of L2TP. However, this functionality comes at the cost of important security limitations as we will see in the next section.

What is left for us to describe in the wholesale model is the ISP method of authentication.

For historical reasons related to the use of dial-up lines and point to point communications, RADIUS became the preferred mean to authenticate, provide user specific information and cumulate basic accounting data needed for billing systems.

Since PPPoE provides the same access control, authentication method and accounting functionalities as traditional PPP sessions, ISPs could easily leverage their current RADIUS deployments.

³ Vendor-specific implementations of radius also offer functions such as pre-authentication messages and so-called pseudo-users to can be used in this context. While a discussion on vendor-specific radius options might be interesting, we will limit ourselves to the standard options defined in RFC2865, RFC2868 and RFC3437 as it is the basic common denominator that assures interoperability. Refer to RFC2882 (Mitton) for a discussion on the problems raised by vendor-specific implementations of the radius protocol.

Therefore RADIUS remained the de facto method of authentication for ISPs.

Areas of vulnerabilities of the DSL wholesale model

The first element of concern is the Layer 1 to Layer 2 aggregation done by the DSL wholesaler: once the PPPoE session gets to the PPPoE server / LAC, all information about the physical aspects of the connection is lost as the only connection known to the LAC is the interface aggregating the traffic (normally an Ethernet or an ATM VCC). So information such as the Dialed Number Identification Service or a circuit number that could be used by the ISP to discriminate the end-user is not passed along.

Consequently, the use of PPPoE as an aggregation technology does not provide any additional security hints to the ISP other than the username / password combination normally found in PPP sessions.

Therefore a first area of vulnerability can be identified: since nothing can be done at this stage of the network to discriminate the end-user, the DSL wholesale model makes the first step toward a successful misuse of the service, i.e. gaining physical access to any of the ISP's networks served by the DSL provider, a childish thing to accomplish once a user subscribes to any of the DSL offerings from one of the ISPs.

Also note that this can be achieved while maintaining a high level of anonymity.

A second element of concerns involves the LNS / NAS functions. As defined in (Townsend, section 4.1), a RADIUS Access-Request must contain a NAS-IP-Address or a NAS-Identifier attribute or both on top of the username / password combination.

While the LNS will send by default one of its IP addresses as the NAS-IP-Address, normally the one of the LAN interface pointing at the RADIUS server, this information is often of little value:

Either multiple types of services, such as DSL and Dial-Up, are aggregated via the LNS and therefore knowing from which LNS the connection came from has no value since it does not help discriminating the service through which the subscriber is accessing the network.

Or multiple NAS are present, potentially some dedicated to dial-up connections. This would require the ISP to configure the proper NAS IP addresses versus the service purchased in the RADIUS client definition file for every end-user. This could be the differentiator we are looking for.

However most ISPs offering DSL access also offer, as a complementary service, a dial-up access to their customers to provide them with the possibility of

accessing emails or the Internet in general while they are away from home. This means that even if the ISP could distinguish between the NAS that are used for DSL versus the ones used for dial-up access, the RADIUS server will not be configured with the NAS-IP-Address or NAS-Identifier attributes to authenticate the subscriber since he could potentially come in via any of them.

For those few ISPs that do not offer a dial-up access with their DSL service and could make use of these attributes, they will need to have deployed a RADIUS server that allows for multiple entries per attribute in their client definition file. Unfortunately, there are still many RADIUS implementations out there that do not support this.

One could think that other RADIUS attributes could potentially be configured on the LNS to help out during the authentication process as part of the Access-Request. Unfortunately, most of original attributes and the extensions proposed later are rendered useless once going through the LNS.

For instance, the NAS-PORT-Type attribute, which normally indicated an asynch or an ISDN type of connection in a dial-up scenario, is systematically set to “virtual” when sent by the LNS, regardless of the access technology.

So a second area of vulnerability is the lack of pertinent information provided at the time of the login by the LNS when compared to those that are obtained in a NAS when deployed for dial-up connections directly in the ISP network. This has led to minimal security settings in most RADIUS client configuration files: as long as the hacker has a valid username / password combination to use, he will succeed in getting authenticated by the ISP.

The last element we would like to cover in this section is the RADIUS server itself. As we have just seen, while RADIUS is capable of more complex schemes of authentication, the ways in which the DSL wholesale model is deployed and the minimal information received in the ACCESS-REQUEST packet have brought the authentication process to a simple formality with little security value.

A point which we have not covered so far is RADIUS accounting (Rigney): all ISPs use it in order to bill their clients. However few of them use it to correlate the information collected in order to identify abusive use of their network services.

As we indicated in the scenario at the beginning of this report, most ISPs collect the following informations:

- Time-based information - the cumulative period of time an end-user has an established session with the ISP.
- Usage-based information – the cumulative amount of bytes generated and received by the end-user.

This information is cumulated but it is normally the end-user who reaches the conclusion that he has been victim of a fraud upon receiving its statement. Then begins the difficult process for the end-user of proving that he did not transfer as much or stayed on-line as long as the statement says he did, the ISP hiding behind its accounting records.

This leads to the last area of vulnerability we wanted to highlight: the real opportunity to uncover hackers misusing someone else's DSL access resides in information correlation, as a single element is often not enough to raise suspicions. Few ISPs have implemented operational procedures that could flag suspicious accounts and therefore frauds are mostly uncovered after the fact and the hacker is long gone.

Potential approaches

Our intention in this section is to see what can be done to prevent our original scenario of a user subscribing to a DSL service from one ISP to use the same service from another ISP while subscribing to its dial-up service.

By no mean, however, would we like to give the impression we are exploring all the possibilities or even offering a definitive solution to the problem: as we have seen so far, things like network topology, operational procedures and service packaging, to name a few, are all variables that will vary from one ISP to another, thus making a "one size fits all" solution very unlikely.

The first things that should be looked at are the services offered by the ISP: by documenting what is offered by each service package and how it is implemented, the ISP will be able to identify the areas where there are overlaps and assess the security risks involved with the services. Once this exercise done, the ISP will be able to make an informed choice on the course of action.

Without this initial assessment, the ISP will spend money and resources trying to fix something that could have been solved more easily but in a different way: coming back to our example of an ISP offering DSL services with an amount of hours for dial-up access for the times the user is away from home, the solution might be to provide the subscribers with two distincts user logins, one for DSL and the other one for the dial-up, rather than to spend some time coding scripts interfacing with the RADIUS server.

From a network perspective, the ideal situation would be for the DSL providers to maintain a database of users versus the DSL lines provisionned. This would allow them the possibility of challenging the initiator of the connection with protocols such as the PPP Extendable Authentication Protocol (Blunk).

Unfortunately the DSLAM, which is Layer 1 device, does not have the “brains” for this and by the time the PPPoE traffic hits the aggregation device, all notions of a circuit is lost.

So the only place where a change in the network would provide some help with differentiating over which service a user is coming in, is at the ISP location itself: by making sure the NAS function for each service is located on separate devices, ISPs would at least be able to determine if the user is coming over a dial-up connection or a DSL connection.

For example, an ISP with its own RAS and a separate LNS for its DSL service has all it needs to associate a user with a specific service in RADIUS. If the RAS is also coming in via L2TP, then two separated LNS, one for DSL and one for the dial-up services, would achieve the same result. As we have demonstrated in the previous section, it is when a single box like the LNS is used to aggregate numerous types of services that we lose the ability to differentiate.

As we have also seen, the authentication capabilities of RADIUS depend largely on the information passed along by the NAS. However, for the case where the DSL and the dial-up connections would be coming in on separate devices, it would be possible to point them to separate servers while maintaining the same username / password combination⁴ for those users that are defined for both services.

Beside the inherent benefit of being able to match a RADIUS server with the NAS of a specific service, the client information file could now be customized for each service without having to resort to a minimal password validation as we have described previously. For example, the RADIUS server used for authenticating DSL users could have all of its entries set to a single connection at the time (how could a legitimate DSL user come in from 2 locations?) while the server used for dial-up authentication could make use of all the attributes pertinent to that environment such as the NAS-PORT-TYPE⁵.

Conclusion

With the two approaches described previously, separating the NAS and matching a RADIUS server with a specific service, we are able to prevent our original problem scenario of having a user subscribing to a DSL offering from one ISP ending up using the DSL service from another ISP while subscribing to its dial-up service: the user account would exist only on the RADIUS server tied to the dial-

⁴ The ISP would need to make sure the username / password combination is the same on both servers but this is a supported feature in many commercial RADIUS offerings.

⁵ Certain RADIUS implementations allow for nested client entries arranged in a “if / then” logic. Therefore the first entry could check for a dial-up connection and, if the verification fails, use the following DSL settings. This could eliminate the need for two servers provided all the services don't come over L2TP in which case all sessions look the same as we have described previously.

up service and therefore the authentication would fail if coming from the LNS aggregating the DSL service.

Even in the case where a valid DSL user account would be used to try to setup a session for which there is no subscription, setting the number of connections to 1 would result in one of two possibilities, provided the real user gets online frequently:

1. The user is already connected and the attempt will fail (too many connections)
2. The hacker is connected and the user will not be able to log for the same reason. However, chances are the user will phone the help desk to investigate the cause of the problem, thus raising the flag on the situation.

In the case where the approaches described so far would be unpractical, one area worth investigating is the accounting capabilities of the RADIUS server: through the corrolation of data such as the connection times, the cumulative connection times, the cumulative bytes counts over an accounting period, etc. it would be possible to identify suspicious accounts and possibly investigate them.

This however is beyond the scope of this research.

References

Black, Uyless. "PPP and L2TP: Remote Access Communications". Prentice Hall, 2000. 219 pages.

Blunk, L; Vollbrecht, J; "PPP Extensible Authentication Protocol (EAP)". RFC2284, 1998.
URL: <http://www.ietf.org/rfc/rfc2284.txt?number=2284>

Grossman, D; Heinanen, J; "Multiprotocol Encapsulation over ATM Adaptation Layer 5". RFC2684, 1999.
URL: <http://www.ietf.org/rfc/rfc2684.txt?number=2684>

Mamakos, L; Lidl, K; Evarts, J; Carrel, D; Simone, D; Wheeler, R; "A Method for Transmitting PPP Over Ethernet (PPPoE)". RFC2516, 1999.
URL: <http://www.ietf.org/rfc/rfc2516.txt?number=2516>

Mitton, D; "Network Access Servers Requirements: Extended RADIUS Practices". RFC2882, 2000.
URL: <http://www.ietf.org/rfc/rfc2882.txt?number=2882>

Palter, W; Townsley, W; "Layer-Two Tunneling Protocol Extensions for PPP Link Control Protocol Negotiation". RFC3437, 2002.

URL: <http://www.ietf.org/rfc/rfc3437.txt?number=3437>

Rigney, C; Willens, S; Rubens, A; Simpson, W; "Remote Authentication Dial In User Service (RADIUS)". RFC2865, 2000.

URL: <http://www.ietf.org/rfc/rfc2865.txt?number=2865>

Rigney, C; "RADIUS Accounting". RFC2866, 2000.

URL: <http://www.ietf.org/rfc/rfc2866.txt?number=2866>

Rigney, C; Willats, W; Calhoun, P; "RADIUS Extensions". RFC2869, 2000.

URL: <http://www.ietf.org/rfc/rfc2869.txt?number=2869>

Simpson, W; "The Point-to-Point Protocol (PPP)". RFC 1548, 1993

URL: <http://www.ietf.org/rfc/rfc1548.txt?number=1548>

Townsley, W; Valencia, A; Rubens, A; Pall, G; Zorn, G; Palter, B; "Layer Two Tunneling Protocol (L2TP)". RFC 2661, 1999.

URL: <http://www.ietf.org/rfc/rfc2661.txt?number=2661>

Zorn, G; Leifer, D; Rubens, A; Shriver, J; Holdrege, M; Goyret, I; "RADIUS Attributes for Tunnel Protocol Support". RFC 2868, 2000.

URL: <http://www.ietf.org/rfc/rfc2868.txt?number=2868>

SearchNetworking.Com "Dial Number Identification Service"

URL:

http://searchnetworking.techtarget.com/sDefinition/0,,sid7_gci213907,00.html

ISP Planet radius discussion list 2001 – 2003 archives

URL: <http://isp-lists.isp-planet.com/isp-radius/archives/>

Cisco Systems "Decoding a Sniffer-trace of RADIUS transaction"

URL:

http://www.cisco.com/en/US/tech/tk583/tk547/technologies_tech_note09186a0080093f42.shtml