



Understanding and using AEP Networks SmartGate VPN to access the IBM Innovation Center for Business Partners

Tim O. Robinson, Ph.D
Senior IBM Technical Consultant

Contents

2 Data communications for SmartGate

- *Network Application Communications Path*
- *VPN traffic communications path*
- *HTTP Proxy Support*
- *SOCKS Proxy Support*

5 Configuring and using SmartPass

- *Online Registration and Authentication*
- *Authorization and Network Traffic flows*
- *Telnet... again*
- *Accessing multiple services and locations via multiple AEP SmartGate servers*
- *Communication originating from the remote server*

8 Support for other applications

- *Configuring applications to use AEP SmartGate VPN*
- *A proxy? On my workstation?*
- *Citrix ICA Client*
- *Linux (and Solaris) SOCKS support*
- *Another technique... local port forwarding*

Introduction

This guide is designed to help you understand how the AEP Networks SmartGate VPN product operates and how it connects your workstation to the servers in our lab. Each section below provides technical details on a particular aspect of this product. Our objective is to provide a reference that you can scan through and learn about the architecture, the registration and authentication processes, and data flows required for access. With this information, you should also be able to configure a testing scenario that runs your Internet or other TCP/IP network based application through AEP Networks SmartGate VPN as well.

Data communications for SmartGate**Network Application Communications Path**

In order to communicate with servers over the Internet, network traffic flows between hosts using the Internet Protocol, or IP. It's outside our scope to describe all the fundamentals of IP, however we will describe a few aspects here that are important to understanding how to run networked applications via AEP SmartGate. For more information, consider the first 2 chapters of TCP/IP Network Administration by Craig Hunt (O'Reilly & Associates, 2nd Edition, 1997). Most IP networked applications use one of two communications protocols that manage the delivery of data, TCP or UDP. Common applications that use TCP are telnet, ftp, ssh and web clients and servers. Common applications that use UDP are domain name service (dns), network file system (nfs), and network time protocol (ntp). Both TCP and UDP use an addressing mechanism called port to coordinate the communications process and to standardize how specific services are accessed at the host. TCP and UDP do differ in the management of communication sessions between systems. Communication sessions using TCP have an intrinsic notion of direction of creation and once established, are retained until signaled to close by the applications involved. TCP also includes mechanisms to identify and retransmit data that is lost or dropped in the communication path. UDP does very little to manage communication sessions, but rather relies upon the applications involved to manage the session state (if any) and identification of lost data and its retransmission.

AEP SmartGate VPN provides a mechanism to send TCP data connections over the Internet or other unsecured networks in encrypted form. Since the data connections are TCP, there is a sense of direction for establishing the sessions. We will refer to the originator of a network request as the "network client" and the destination for the request the "network server". The VPN

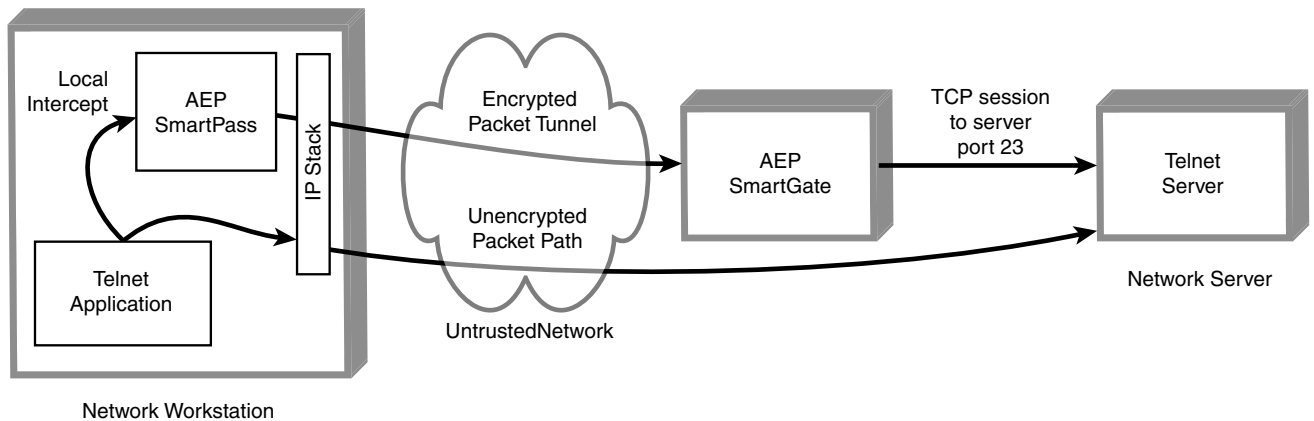


Figure 1

mechanism uses software installed on a network client (AEP SmartPass) and a server on the “other side” of the unsecured network (AEP SmartGate). In most cases, the AEP SmartGate is not the endpoint (network server) for the TCP data connections from the network client; it is just the endpoint for the encrypted data. This is much easier to follow with a practical example.

Let’s talk about the telnet application (see Figure 1). In the telnet application, a terminal program is launched on the network client system. In normal use on the Internet, this terminal program opens a TCP data connection on a remote server on port 23, where a network server application is “listening” for client connections. Once the session is established between the network client and network server, the server application prompts the user for login information, and the network client responds, etc.

With AEP SmartGate VPN, the TCP session coming from the network client is intercepted at the workstation. It is then encrypted and transmitted over the Internet to the AEP SmartGate. On that host, the encrypted data is decrypted and transmitted as a normal TCP session originating from the AEP SmartGate to the network server running the telnet service listening on port 23. Effectively, the TCP connection from the network client is forwarded via the AEP SmartGate to the network server. Once the TCP data session is established, data can move bidirectionally between the network client and the network server, the only key distinction of client and server is in the direction of origination of the communication session.

The same mechanism is used for other IP applications that use TCP -- for example web traffic to HTTP servers, or ssh sessions running over TCP port 22. So far, we have not discussed exactly how the AEP SmartPass and AEP

SmartGate communicate, just the services that the VPN transport provides. That's an entirely different topic, but an important one from a operational perspective -- let's cover that next.

VPN traffic communications path

In this section, we will describe how the AEP SmartPass itself communicates with AEP SmartGate. If you are looking for how networked application traffic and data is handled by AEP SmartGate VPN – that's in the section just above. By reviewing the communications mechanism between the AEP SmartPass and SmartGate, you should be able to learn enough information to successfully configure the AEP SmartPass so that it can work in your environment. Although AEP SmartPass can be configured in multiple ways to establish a communications path with AEP SmartGate, the default settings are usually acceptable for use immediately after installation.

By default, the AEP SmartPass will try to establish a direct TCP session with the AEP SmartGate. Immediately after installing, the AEP SmartPass needs to perform an online registration (OLR) to store the client authentication token on the server. The OLR process uses TCP connections to ports 80 and 443 on the AEP SmartGate. Once OLR is complete, the default transport for ordinary VPN traffic is TCP communication sessions to server ports 80, 443 or 3845 (the default is 443). The direction of the TCP communication session establishment is from the client to the server. If your firewall supports direct communication by TCP to servers at these destination ports, the AEP SmartPass should be able to connect to the AEP SmartGate without modification. For example, if you have a NAT firewall, and do not configure your web browser with special proxy settings, AEP SmartGate VPN should work without additional configuration.

HTTP Proxy Support

The AEP SmartPass supports and has been tested with firewall-based HTTP proxy servers that are used to provide a gateway between an enterprise and the Internet. To use these proxies, the AEP SmartPass may need to be configured. During installation, AEP SmartGate attempts to discover proxy settings from the configurations of installed web browsers. If this auto-detection is not successful, an HTTP proxy may be manually set. This is specified through the Options... panel in the AEP SmartPass. To enable network application traffic through the VPN, the HTTP proxy IP address and port need to be entered on the Web proxy and the SSL proxy tabs. Then to include all TCP network (non HTTP and HTTPS) traffic through the SSL proxy, check the "Use the SSL tunneling proxy defined on

the SSL proxy tab” check box on the ftp proxy and generic proxy tabs.

SOCKS Proxy Support

Since all communication between the AEP SmartPass and AEP SmartGate uses the TCP protocol, it is also possible for a SOCKS proxy at a firewall to be used for the VPN tunnel path. In current versions of the AEP SmartPass, there are no configuration options to enable the use of a SOCKS proxy. However, as with other TCP applications, a workstation that has been configured with a network stack “socksification” product (e.g., Hummingbird SOCKS, SocksCap32) will successfully run AEP SmartPass through an enterprise SOCKS server and communicate with the AEP SmartGate. In our testing, we have found that if the workstation network stack is “socksified”, that there can be problems when the winsock shim for AEP SmartPass is installed. If you are going to use a socksified TCP/IP stack on Microsoft Windows platforms, we recommend that you use the AEP SmartPass package without the winsock shim support package. Please inform your IBM Innovation Center technical consultant if your enterprise uses a SOCKS proxy to reach the Internet.

Configuring and using SmartPass

Online Registration and Authentication

The AEP SmartGate VPN product supports several different types of authentication mechanisms. Each type reflects different enterprise authentication needs for providing extranet access. In our environment, we are using the AEP FIPS software tokens as these provide for strong authentication and simple initial setup. When configured with a FIPS software token, AEP SmartPass, will require you to type in an access code to unlock the token on startup. During first time setup, you will be prompted for this access code after you complete the installation of AEP SmartPass. At a later time, the access code can be changed (only if it is known) using the AEP FIPS token control panel, or the ‘olr’ program on UNIX. The combination of the software token and the access code to unlock it characterizes this as a two-factor authentication scheme. It is based on something that you have (the token stored on the client) and something that you know (the access code).

The objective of online registration (OLR) is to securely transmit the key stored in the AEP SmartPass FIPS token to the AEP SmartGate. OLR uses a novel set of steps based on web technologies to securely perform this exchange. The first step in OLR is to launch the AEP SmartPass (or

the separate 'olr' program on UNIX), this initializes a local registration service. To register your client, open the URL that we have provided in the installation instructions in a web browser. This URL contains a special web form that is posted (via the HTTP POST method) to the local registration service running on your workstation. When that service receives the data from the form, the data and the key information from in the software token are forwarded to AEP SmartGate using Secure Sockets Layer (SSL).

The next phase in activation of the account is performed by our administrators on the AEP SmartGate in the IBM Innovation Center. We monitor incoming OLR requests and review the data that is provided in the web forms. If the data provided in the web form matches our records, we will mark your userid and the key associated with it as enabled for access. When your AEP SmartPass credentials are enabled, the account (user name) for your client is authorized with the Web and TCP port access permissions for you to use to reach the servers you are using in our lab. We will cover that, authorization and access, next.

Authorization and Network Traffic flows

This section is a follow up to the first section in the white paper where the basic communication path used in AEP SmartGate VPN is explained. Where that section talked about how the communication works, this section explains the mechanism whereby specific applications are allowed, and how to read the permissions that are listed by the AEP SmartPass. Immediately after authenticating (and by request by the Refresh option), the AEP SmartPass downloads a list of IP addresses, destination TCP ports and web URLs. These permissions tell the client what is allowed through the VPN and may include wild cards to specify connections for multiple host IP addresses or destination TCP ports.

There are two main categories of authorization, web permissions and TCP permissions. These are associated with different icons in the GUI display on Microsoft Windows platforms. Web permissions are evaluated when a web request is either handled implicitly by the winsock shim or when the local web browser uses the AEP SmartGate Web proxy running on port 2080. From an authorization perspective, a more granular level of control is available through web permissions as they control both the access to the web server (a TCP control) and the specific URL pattern that is permitted for the web transaction. TCP permissions define the mapping between a local TCP "proxy" port and the remote server IP address and TCP destination port. These permissions are evaluated when a request to open a TCP session

is seen by the winsock shim, sent directly on a local port, or when a remote TCP session establishment is requested via the SOCKS protocol on the local SOCKS server that is created by the AEP SmartPass on startup. Let's go back to the telnet example to see how this works.

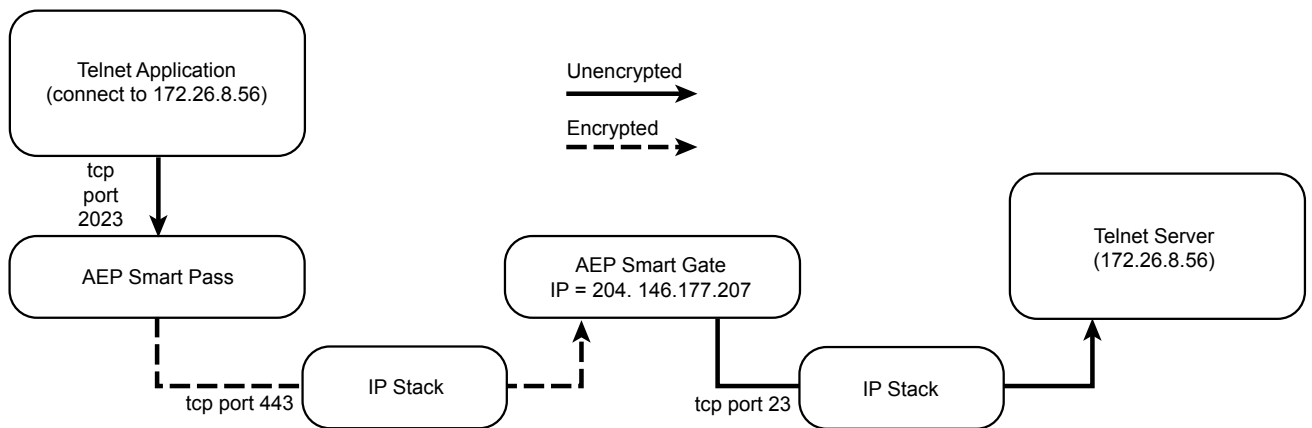


Figure 2

Telnet... again

Consider a telnet server with IP address 172.26.8.56. After the AEP SmartPass authenticates, it downloads a TCP permission that authorizes connections to the server address 172.26.8.56 on TCP port 23 (the telnet server port). This rule also includes two other ports that are shown in the AEP SmartPass display, the local port (lport) and the server port (rport). From a user perspective, the server port (rport) can be ignored. What is important is the local port. In our example of telnet, this port number will default to 2023. After evaluating this TCP permission, the AEP SmartPass will open a listening socket on the local system on port 2023. To connect to the telnet server, the telnet application is launched on the local workstation with localhost (127.0.0.1) on port 2023 as the destination. The AEP SmartPass takes TCP traffic coming to local port 2023, encrypts it, and sends it over the Internet to the AEP SmartGate. From there, the AEP SmartGate decrypts the TCP data, and forwards it to 172.26.8.56 on TCP port 23.

In the case where multiple telnet servers are allowed, there will be additional TCP access permissions defined as a list, an IP network range, or with a "*" wild card. All of the telnet servers will be available through connection to local port 2023 by default. This is coordinated through a prompting mechanism by the AEP SmartPass user interface. After the telnet application connects to localhost port 2023, a prompt is raised by the AEP SmartPass user interface where the correct destination server may be typed in or selected from a list. Alternatively, on the Microsoft Windows

platforms this process is simplified by a winsock shim. The winsock shim intercepts either the IP address or logical server name used by the telnet application. If that IP address or server name matches the TCP permissions, it automatically forwards the traffic over the VPN.

In a nutshell, the TCP and web permissions that the AEP SmartPass obtains form a policy that defines what network connections are forwarded from the local host ports to the remote server. This is a security model that supports most TCP applications in client/server and network computing environments. AEP SmartGate VPN also can support query/response UDP traffic through a similar model, and this is used to support Citrix Metaframe application browsing on UDP port 1604.

Support for other applications

Accessing multiple services and via multiple AEP SmartGate servers

In an Extranet application environment, it is often the case that the network client will need to access applications on network servers located in different areas of authority, or even in different geographies. A good example is a network client with the need to access both an intranet web-based human resources application and a hosted customer relationship management application at an application service provider. This type of mixed access is easy to support through multiple OLR steps, where each step adds a AEP SmartGate key to the AEP FIPS token on the network client. Since each of our labs that is connected to the Internet uses AEP SmartGate, you may

register with two (or more servers) to obtain seamless access to your testing resources. For example, you may be testing your application on an IBM System p at our lab in Australia, and an IBM System i at our lab in the United States. In this case, after OLR to the SmartGate servers in each geography is complete, the AEP SmartPass will display multiple TCP and web permissions when it refreshes. Using the example IP addresses shown in the 3, when a network connection to the

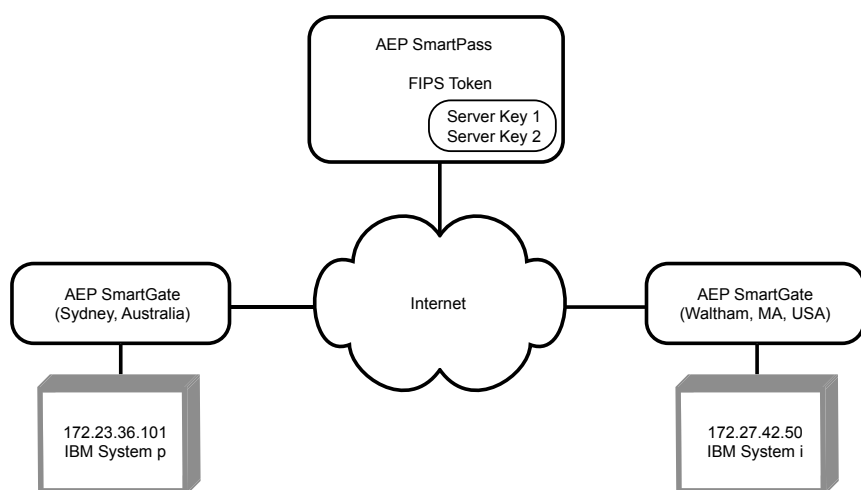


Figure 3

server at 172.23.36.101 is requested, the traffic is sent over the encrypted link to the AEP SmartGate in Sydney, Australia, whereas any traffic to 172.27.42.50 goes to the AEP SmartGate in Waltham, Massachusetts.

With the exception of the network latency delay associated with Internet transmission over an international link, the access to network servers in different geographies is transparent to the end user.

Communication originating from the remote server

What about communications (TCP or UDP queries and their responses) that originate at the remote server, for example, the server at 172.26.8.56? This traffic does not have an IP routing path back to the workstation running AEP SmartPass, so it cannot reach the workstation. This is ideal when viewed from the perspective of an enterprise user (behind a firewall) connecting to a remote server on an extranet. A downside is that X11 protocol applications on a remote server are unable to open up windows using that protocol on the local workstation. However, there are alternative ways to use graphical applications and retain this secure TCP flow architecture, for example by using VNC or Microsoft Remote Desktop Services.

Configuring applications to use AEP SmartGate VPN

In our telnet example, we provided a specific set of steps to force the telnet application to connect via the AEP SmartPass created local port 2023. This is possible because the telnet application can be easily configured to connect to a user-defined destination TCP port. Not every application is as accommodating, so other techniques need to be used to make sure that their traffic is forwarded by the AEP SmartPass over the VPN tunnel. On Microsoft Windows platforms, assistance is available in the form of the winsock shim. This shim intercepts network calls from applications and compares them to the TCP and Web access permissions known by the AEP SmartPass. When using the winsock shim, there is effectively no application configuration required so if you have installed the version of AEP SmartPass with this support, reading the rest of this section on configuring applications is not necessary.

For users of Linux and other UNIX platforms, or if you are using the non winsock shim version of AEP SmartPass for windows, the next few sections talk about how to configure applications to use some of the network proxies that are started by the client, and how to use other techniques to cause TCP application traffic to traverse the VPN tunnel.

A proxy? On my workstation?

Citrix ICA Client

The Citrix Independent Computing Architecture (ICA) technology provides a framework for access of applications from a number of device platforms.

The Citrix ICA client can be used to access a graphical user interface application from a number of operating platforms including Microsoft Windows and UNIX X11 platforms. If your workstation is using the winsock shim version of AEP SmartPass, the Citrix ICA client will work transparently to connect to applications through the VPN. Note: when defining an application for the first time, the name or IP address of the server hosting the application needs to be provided, instead of browsing through the Citrix Program Neighborhood. If you are not using the winsock shim version of AEP SmartPass, manually configure the Citrix ICA client SOCKS server preference settings to IP address 127.0.0.1 and port 1080.

Linux (and Solaris) SOCKS support

Although most of this guide is platform agnostic, we will talk here about SOCKS for Linux and Solaris. Consider this equal time to balance the earlier discussions about the winsock shim. In our lab we have tested AEP SmartPass with tsocks. This is a client-socksification library that supports Linux and Solaris, and can be selectively used by setting an environment variable. It can be downloaded from <http://tsocks.sourceforge.net/>, this software is licensed under the GPL. It also may be found included in the packaging for many Linux distributions.

We have tested the AEP SmartPass with tsocks currently recommended version 1.8 beta 5. The source for this package is extracted with the command:

```
# tar -zxvf tsocks-1.8beta5.tar.gz
```

After unpacking, change directory into the tsocks-1.8 directory and configure, build and install the package (there are instructions in the "INSTALL" file in this directory). You will need to create a simple /etc/tsocks.conf file. One of the easiest to use, which will only attempt to use the internal AEP SmartPass SOCKS server would look like this:

```
# simple /etc/tsocks.conf file

# define as local the path to the socks server
local = 127.0.0.1/255.0.0.0

# everything else can go locally to AEP
SmartGate, where it will
compare the
```

```
# request to the know TCP access permissions
server = 127.0.0.1

server_type = 5
```

For more information on creating complex configuration files, see the installed man page for `tsocks.conf`. To use this socks setting with your network applications, set the environment variable `LD_PRELOAD` to `/lib/libtsocks.so` (the default install location on Linux). After setting this environment variable, command line applications like `telnet`, `ftp` and `ssh` will communicate to the AEP SmartPass SOCKS server and be forwarded across the VPN when the destination matches one of the SmartPass TCP permissions. When these applications run, you will see connection information on the console or in the terminal window where you launched “`smartpass`”.

Another technique... local port forwarding

When configuring applications to use the local SOCKS proxy won't work, there is a technique that can be used when the workstation running AEP SmartPass does not run local services on the intended service port numbers. This trick works best when you are only connecting to a single server (for example an NNTP server to `get/post` usenet news messages, or a POP server to retrieve e-mail). For the server, a TCP access permission can be created that creates a local port with the same port number of the server's destination service. With the access permission in place, you can modify the server name in the application to `localhost`. This will result in the application connecting to the local port which is forwarded by AEP SmartPass to the actual server. A similar twist on this applies when you are using servers that are behind an extranet firewall. In this situation, the network names of those servers are usually not available for Internet DNS service. You can still communicate with those servers by name using this local port forwarding technique. To do this, add the name of the server to the local workstation's host file on the line for `127.0.0.1`, this name should be placed after any names like `loopback` or `localhost`. Warning – you should be careful with this type of change, especially if you mix between an intranet/extranet environment that does have DNS resolution when attached to the intranet.



© Copyright IBM Corporation 2008

IBM Corporation

11501 Burnet Road
Austin, TX 78758
USA

Printed in the United States

03-08

All Rights Reserved.

The IBM home page can be found on the Internet at **ibm.com**

IBM is a registered trademark of International Business Machines Corporation.

IBM and the IBM logo, IBM System p, and IBM System i are registered trademarks of IBM in the United States.

Citrix, Citrix Metaframe, Citrix Program Neighborhood and Citrix ICA are registered trademarks or trademarks of Citrix Systems, Inc. in the United States and other jurisdictions.

SocksCap is trademark of Blue Coat Systems.

Hummingbird is a trademark of Hummingbird Connectivity, a division of Open Text Corporation.

Linux is a registered trademark of Linus Torvalds.

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

Mozilla, Firefox and Thunderbird are either registered trademarks or trademarks of the Mozilla Foundation

SourceForge is a trademark of SourceForge, Inc.

Sun, Sun Microsystems and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc in the United States and other countries.

UNIX and X Window System are trademarks or registered trademarks of the Open Group in the United States and other countries.

AEP Networks, SmartGate, SmartPass, are trademarks or registered trademarks of AEP Networks.

Other company, product and service names may be trademarks, or service marks of others.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program or service is not intended to imply that only IBM's product, program or service may be used. Any functionally equivalent product, program or service may be used instead.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, IBM warranty terms apply.

This publication is for general guidance only.