An Often Overlooked Security Hole in Enterprise Extender and Mainframe Networks

By Stu Henderson and Peter Hager

I Introduction and Background on VTAM Networks

This article describes a common mainframe network security exposure, and how you can address it in your data center. If your installation uses IBM's Enterprise Extender to connect your network to another organization's network,then this article applies to you.

If you are already familiar with APPN, you may skip over this Introduction to section II. We start here by describing how **VTAM** (**Virtual Telecommunications Access Method**) works in a single computer, then how it works across different computers in your data center, and then how it works across different networks. We then show you how this architecture can introduce a security exposure, and the tools IBM gives us to close the exposure. (You will note that encryption, firewalls, and VPNs do not protect against this exposure.) We then describe how to investigate this all in your data center, whether you are the Data Security Officer, the CIO, or the VTAM system programmer.

VTAM is the system software on z/OS computers that controls all the terminals and all the telephone connections. Even TCP/IP (Transmission Control Protocol / Internet Protocol) executes under the control of VTAM. VTAM works by making logical connections between **logical units (LUs).** An LU is an entry point to the network, usually either a terminal, or a program like CICS or TSO. Such programs which talk to terminals through VTAM are called **applid**s or **application identifiers**.

When you first sit down at a terminal, it is controlled by VTAM. You type in a request for a connection to an applid like CICS or TSO. If VTAM knows of your terminal, and knows of the applid with that name, it can create a "**bind**" or logical connection between your terminal and the applid.

The result of the bind is that everything you type on your terminal gets sent directly to the applid. Likely the first thing you will type will be your userid and password.

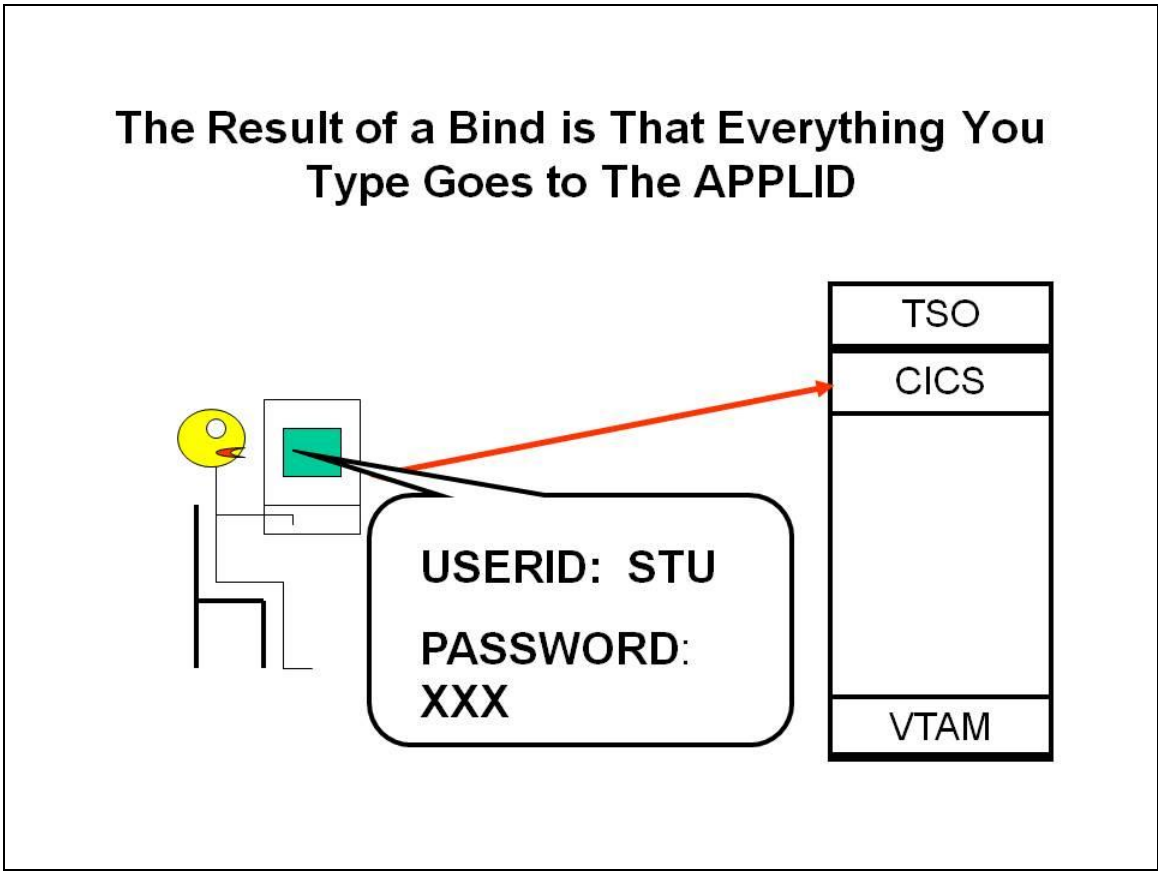

The Result of a Bind is That Everything You Type Goes to The APPLID

Now imagine two computers in your data center, perhaps one for test and one for production. A terminal controlled by VTAM on the TEST computer wants to connect to an applid (let’s say a CICS region named CICSP1) on the PROD computer.

VTAM on the TEST computer receives the request for a bind from the terminal, recognizes that it is for an applid on a different computer, and sends the request to VTAM on the PROD computer. (If you are a parent, think of the two VTAMs as parents setting up a play date for their children. Neither wants to allow the connection to happen unless it knows the other parent, and knows the other child.)

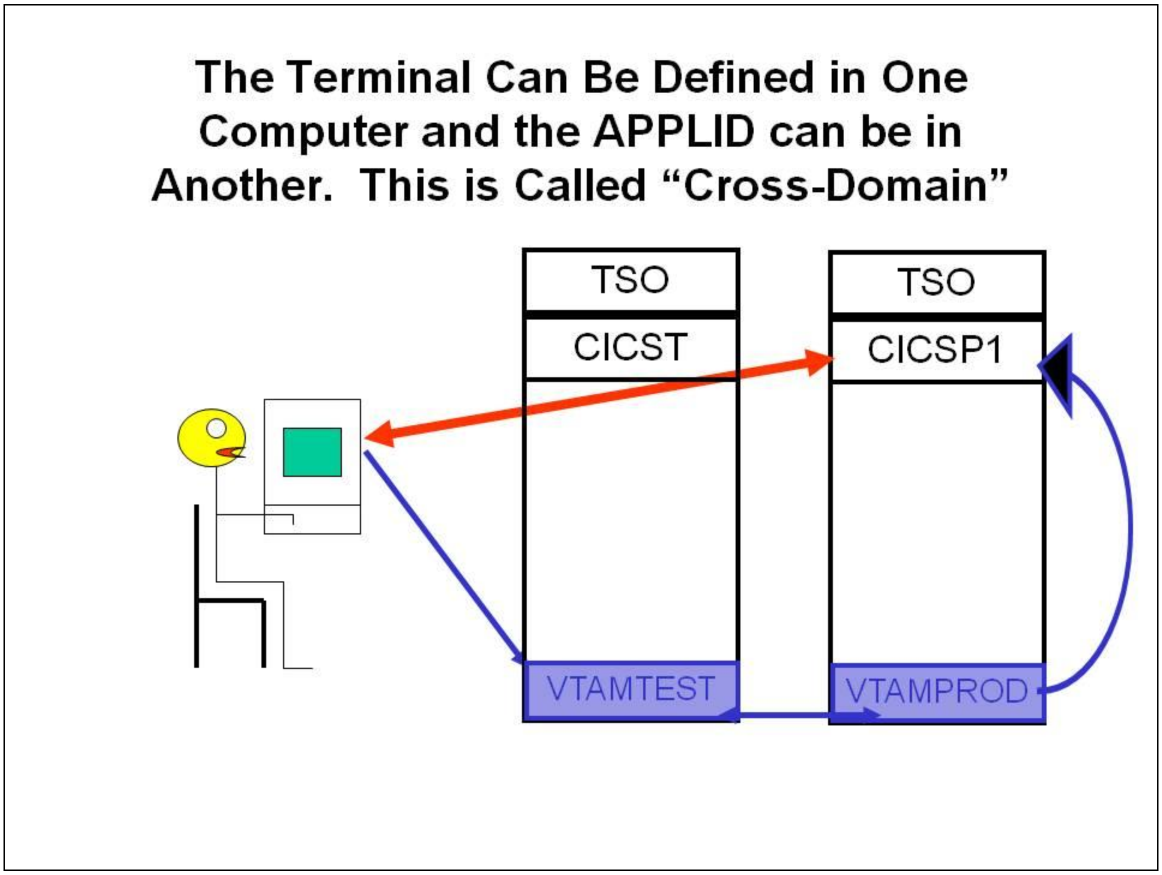

The Terminal Can Be Defined in One Computer and the APPLID can be in Another. This is Called “Cross-Domain”

Such a bind across different computers is called a “**cross-domain**” connection. It can is only permitted if VTAM on each computer knows about the other VTAM and its resources.

Now let’s complicate the picture by imagining a terminal in your data center wants to talk to an applid (again, let’s imagine a CICS region, perhaps named CICSPR2) in another computer, in a different network named NET1, belonging to one of your company’s trading partners. Perhaps you are a financial institution whose network is connected to another financial institution’s network, or you are a manufacturer whose network is linked to your suppliers’ networks.

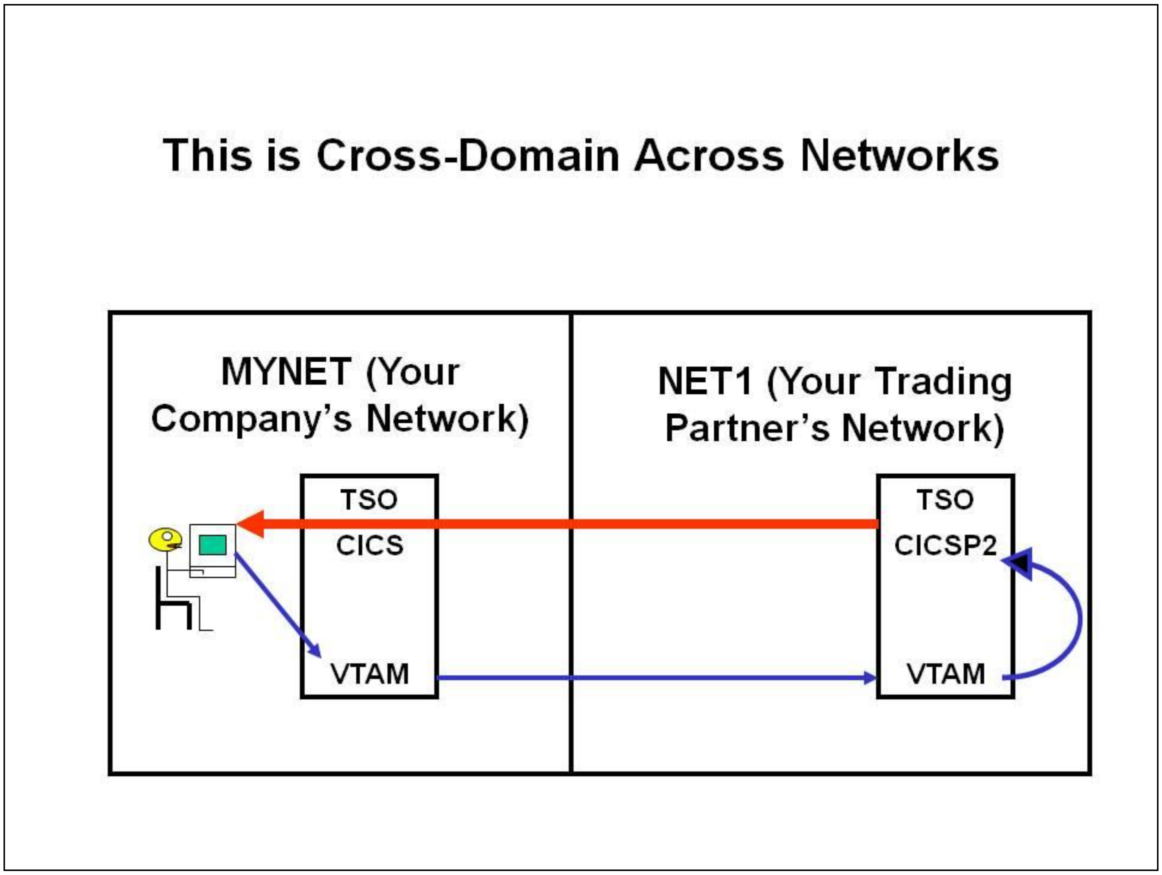

This is Cross-Domain Across Networks

MYNET (Your Company's Network)
NET1 (Your Trading Partner's Network)
TSO
CICS
VTAM
TSO
CICSP2
VTAM

To make this happen, the VTAMs in each network use a protocol called **APPN** or Advanced Peer to Peer Networking. (The word "**peer**" means an equal; the two networks are not master and slave, they are peers.) (You may be using an earlier protocol named **SNI or System Network Interconnection**, but what we describe here still applies.)

You can learn if this situation applies in your data center by asking your VTAM system programmer three questions:

1) What is the name of our network? (He will know this better than he knows his own middle name.)

2) What are the names of the networks it is connected to (and what organization owns each of them) (He will know this immediately, calling them "the adjacent networks". For our example, let's call them NET1, NET2, and NET3.)

3) What networks are the adjacent networks connected to? (He may not know these unless he has discussed them with the VTAM system programmer at each adjacent network, or if your contract with the other organization specifies it.)

II Description of the Security Exposure

This use of APPN or SNI is where the security risk we are describing was first encountered (although related risks have been discovered not involving adjacent networks, as we will describe below).

Imagine a hacker who controls a computer which is connected to a computer network named NETX which is connected to another computer network and so on, until it is connected to one of your adjacent networks. The hacker could send a broadcast from his computer announcing that his computer is NET1. Of course if the real NET1 (which belongs to your company's trading partner) is active, no one will believe this rogue claim, UNTIL the real NET1 comes down. In that case, the hacker's computer becomes the effective NET1. This occurs because the rogue computer's announcement that it is NET1 remains pending until the real NET1 comes down.

The hacker's computer can then put up a rogue applid (CICSP2). He is then spoofing (faking the identity of) both NET1 and the CICS region. At this point a terminal in your network trying to bind to CICSP2 in NET1 can be given the bind. Once the bind is established, your innocent user will likely type in his userid and password, which will be sent to the rogue CICS region. This makes it easy for the hacker to learn userids and passwords from your system.

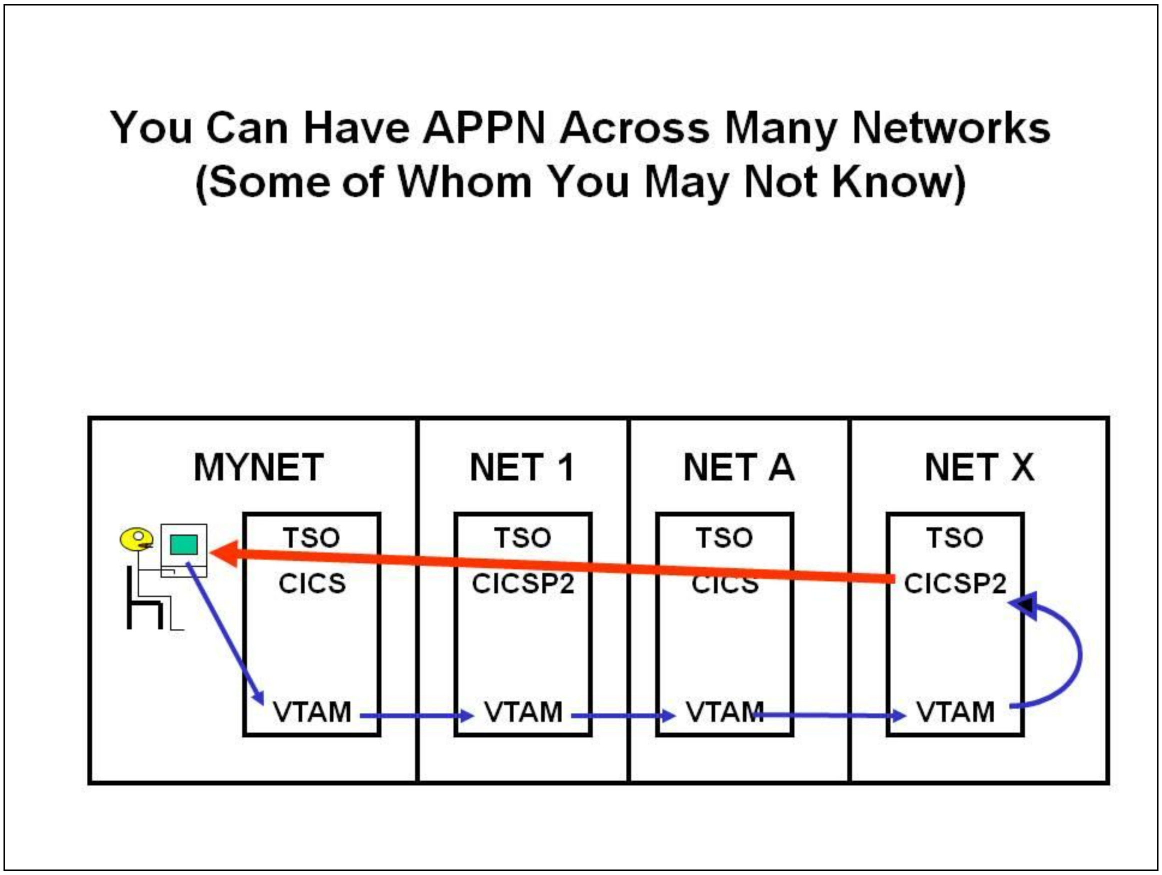

You Can Have APPN Across Many Networks
(Some of Whom You May Not Know)

MYNET

TSO
CICS
VTAM

NET 1

TSO
CICSP2
VTAM

NET A

TSO
CICS
VTAM

NET X

TSO
CICSP2
VTAM

As we will describe below, there are many variations on this attack, some of which are very similar to TCP/IP attacks (Denial of Service, Man in the Middle, etc.).

To see some of the coding involved in such attacks, please refer to the IBM manual: **“Communications Server SNA Programming ”,** number SC31-8829-01, downloadable from IBM z/OS Internet Library **at**

http://publibz.boulder.ibm.com/epubs/pdf/f1a1d930.pdf

There are steps your VTAM system programmer and your Data Security Administrator can take to protect against such attacks. However, unless they work together to address the problem, your network will likely continue to be at risk.

Why Encryption, Firewalls, and VPNs Won't Stop This

Say that you are connecting your network to an adjacent network, then you are probably using a product called **Enterprise Extender** or **EE**. Enterprise Extender actually sends the APPN messages as **SNA (System Network Architecture)** packets tunneled (contained within) UDP packets. (UDP is a protocol similar to, but different from TCP/IP.) Tools like encryption, firewalls, and VPNs all work on the UDP packet, which gets encrypted, decrypted, filtered, and so on, while these tools have no effect on the SNA packet itself. So the attack itself is passed right through inside the UDP packets.

But Aren't We Getting Rid of SNA?

Not really. Your terminals are almost certainly connected to the mainframe through a product called OSA or Open Systems Adaptor, which uses TCP/IP. However, the actual message between the terminal and the applid is an SNA message, which is tunneled (contained within the TCP/IP packet). This connection will continue to use SNA for some years into the future.

Your adjacent network connections might some day switch from SNA to pure TCP/IP. However for this to happen, the owners of all networks involved would have to coordinate their conversion efforts, which would be a major effort to undertake.

III Examples of Possible Attacks and Samples of Tools IBM Gives Us to Manage the Risk

One of the authors (Peter) has identified 20 variations on this attack. Some of them involve cross-network connections, some of them can be implemented with a personal computer plugged into a LAN (Local Area Network) in your building. (For a complete list, please see his website at www.net-q.com/ssl/security_issues.html) Here are a few examples from that list:

A) Spoofing of a network and applid as described above, in order to harvest userids and passwords

B) A personal computer plugged into a LAN connected to your mainframe can use software made available at no cost from Microsoft to perform similar spoofing from within your organization

C) A hacker spoofs a network and applid and a terminal, pretending to be the real terminal to the real applid and the real applid to the real terminal. The hacker then controls all messages passing between the two. The hacker can copy sensitive data from the messages without the two parties being aware. The hacker can also modify the messages maliciously before passing them on. (This is similar to a “man in the middle” attack in a TCP/IP network.)

D) A hacker spoofs a network and an applid (or some terminal), and then rejects all attempts to bind to it. This is a form of Denial of Service attack, similar to those on TCP/IP networks.

To investigate this risk in your organization, you need to answer the three questions listed above. If you do not use APPN (that is, you have no adjacent networks), then you are not exposed to this risk.

If you use APPN or SNI, then you need to examine several settings in the VTAM start-up parameters, resource definitions, Class of Service tables, and other places. You also need to review your security software (RACF, ACF2, or TopSecret) to see if it uses the VTAMAPPL and APPCLU resource classes, since these can provide protection against spoofing of applids and of networks. (In RACF, look at the DSMON report; in ACF2 look at the SHOW ALL output for the SAF mappings and definitions. For TopSecret, look at the RDT or Resource Definition Table.)

The VTAM settings include:

- specification of the number of cross-network hops permitted,
- whether to use the security software to verify the identity of networks and applids,
- whether to permit dynamic connections, and
- several dozen others.

You can review these yourself, use analytic software to evaluate them, or use fire-wall like software (actually called an “SNA Firewall”) to prevent such attacks. (Examples of such software are available at the Net-Q website. You may be able to identify additional tools from other sources.)

Once you know the security software and VTAM settings, you can adjust them to provide whatever level of protection you consider appropriate.

IV Summary and Call to Action

This is a real risk, often overlooked, and easy to minimize with the tools IBM gives us. The authors have encountered it in a large number of mainframe installations. While the probability of an attack occurring may be low, the amount of damage possible is very high, and the tools to protect against it are available now.

The most sensible first step is to invite your Data Security Officer and your VTAM system programmer to lunch to discuss whether any further investigation is needed. You should know by dessert whether this risk is negligible in your installation, or whether further investigation is justified.