

Effects of ISP Business Practices on the Versatile Home Network

Tom Schmidt
Schmidt Consulting
tom@tschmidt.com
2 July 2002

Abstract:

VHN was designed as an autonomous local area network. It did not require an ISP connection to operate. While VHN is capable of operating independently its main value derives from integrating it with Internet connectivity. ISP business practices have a profound impact on how VHN operates and what services it is capable of delivering.

Table of Contents

| | | |
|-----------|---|----------|
| 1 | OVERVIEW | 1 |
| 2 | THE IDEAL ISP | 1 |
| 2.1 | UNLIMITED IP ADDRESSES | 1 |
| 2.2 | PERSISTENT ADDRESSING | 1 |
| 2.3 | ALWAYS ON CONNECTION | 2 |
| 2.4 | AUTOMATIC CONFIGURATION | 2 |
| 2.5 | TRANSPARENT TRANSPORT | 2 |
| 2.6 | SPEED TIERS..... | 2 |
| 2.7 | QUALITY OF SERVICE..... | 2 |
| 2.8 | AUTHENTICATED SERVICES | 3 |
| 2.9 | CUSTOMER DNS | 3 |
| 2.10 | CUSTOMER FRIENDLY ACCEPTABLE USE POLICY | 3 |
| 3 | ACCESS ISSUES | 3 |
| 3.1 | NATED ISPs | 4 |
| 3.2 | PROXIED ISPs | 4 |
| 4 | HOW MANY IP ADDRESS..... | 4 |
| 4.1 | SINGLE ADDRESS | 4 |
| 4.2 | A FEW ADDRESSES | 5 |
| 4.3 | UNLIMITED ADDRESSES | 5 |
| 5 | BRIDGED VS. ROUTED NETWORK..... | 5 |
| 6 | ADDRESS ALLOCATION..... | 6 |
| 7 | ENCAPSULATION (PPPOE/PPPOA)..... | 6 |
| 8 | ASYMMETRIC SPEED..... | 7 |
| 9 | LATENCY | 7 |
| 10 | MULTILINK AND BONDING | 7 |
| 11 | MULTIPLE ISPS | 8 |
| 12 | AUTHENTICATION | 8 |
| 13 | OVERSUBSCRIPTION | 8 |
| 14 | QOS | 8 |
| 15 | RESTRICTIVE POLICIES | 9 |
| 15.1 | BANDWIDTH HOGS..... | 9 |
| 15.2 | PORT BLOCKING | 9 |
| 15.3 | PROHIBITION OF "BUSINESS CLASS" SERVICES (VPN)..... | 9 |
| 15.4 | PROHIBITION OF HOME NETWORK | 9 |
| 15.5 | PROHIBITION OF SERVERS | 10 |
| 15.6 | T&C CHANGE WITHOUT NOTICE..... | 10 |

| | | |
|-----------|----------------------------------|-----------|
| 15.7 | GEOGRAPHICAL RESTRICTIONS | 10 |
| 15.8 | DMCA SAFE HARBOR PROVISION | 10 |
| 15.9 | SERVICE LEVEL AGREEMENT | 10 |
| 16 | DEPLOYMENT ISSUES | 11 |
| 17 | TECHNICAL SUPPORT..... | 11 |
| 18 | CONCLUSION..... | 11 |

1 Overview

The main value of the VHN is its connection to the Internet. Even though the VHN is able to function as an autonomous isolated network the main reason users will deploy a VHN is to access the Internet. The ISP is the link between VHN and the Internet; as such ISP policies have a profound effect on which services can be implemented on the VHN.

The Internet has experienced explosive growth because it is a general-purpose bit delivery mechanism that allows end-to-end communication. The Internet makes few assumptions about the data it carries. New features can be deployed at the edge of the network, without requiring coordination or permission of the network owner. Unfortunately ISPs are trending in the opposite direction, especially the Cablecos. Service offerings are tailored to current usage patterns and restrictions put in place making it hard for new service to be developed. If this situation existed in the early '90s ISPs would have optimized their network for Gopher and discouraged use of HTML.

This paper discusses business and technical attributes of an ideal ISP and enumerates numerous areas where ISP practices diminish the value of the VHN and are at odds with Internet openness.

The [High Tech Broadband Coalition](#), of which CEA is a member, identified some of these problems. The coalition is lobbying the FCC and Congress to encourage customer friendly regulation.

2 The Ideal ISP

Before discussing restrictive business practice let's examine the ideal ISP service offering from the perspective of the VHN.

2.1 Unlimited IP addresses

A typical VHN will consist of a large number of devices, potentially hundreds of IP nodes. While this may seem ludicrous today one only has to look at the history of electric light and electric motors to see similar adoption patterns. What was once an expensive novelty became ubiquitous.

The ISP reserves a large contiguous address block for each customer. The block needs to be contiguous so each VHN acts as a single subnet. This allows local traffic to remain local. If the address range is not contiguous traffic has to flow across the first-mile connection to the ISP.

2.2 Persistent Addressing

Servers require a stable means of access either by address or name. TCP/IP tends to blur the notion of server. Many IP devices will implement server functions. Most users are familiar with Uniform Resource Locators. DNS translates URLs to IP address. It assumes a persistent IP address that changes infrequently. DNS changes take time to propagate so if the IP address changes the resource is temporarily unavailable until DNS is updated.

Persistent addresses are also useful in analyzing traffic and creating access policy rules. Event logs are reported by IP addresses. If the address is not constant interpreting the logs becomes difficult.

2.3 Always on Connection

One of the benefits of broadband is it is “always-on.” This is critical to maximize VHN benefit. It is not sufficient to merely reduce the connect time so it is invisible when an in-home device accesses the Internet. If the connection is not on 24/7 the VHN is not accessible remotely preventing the use of application such as Voice over IP telephony, instant messaging and personal web servers.

2.4 Automatic Configuration

Home user technical expertise varies from technically savvy to novice. Configuring arcane IP address settings is a daunting task for many customers.

Low-level settings such as IP address, subnet mask, gateway address and DNS server address should be performed automatically. Besides easing the end user task automatic assignment allows the ISP to modify these settings automatically as the need arises without having the customer manually effect the change.

2.5 Transparent Transport

The ISP is a transparent bit delivery mechanism. All IETF complaint traffic presented to the ISP is accepted. Unless specifically requested by the customer traffic is delivered on a best effort basis with equal preference for all.

2.6 Speed Tiers

First mile speed constraints will be around for the foreseeable future. A reasonable pricing model is to vary price based on offered speed. This allows the customer to make a price performance tradeoff decision. Once this decision is made service cost is unaffected by usage.

Most residential Internet access is asymmetric, download speed exceeds upload. This imbalance is the result of both technical and business considerations. While there is nothing inherently wrong with asymmetry the difference in download vs upload speed should not be extremely wide or used to effect defacto usage restrictions. Customers should be encouraged to use bandwidth, driving demand, similar to what has be done in the PC industry with CPU speed.

2.7 Quality of Service

Almost all residential broadband suppliers deliver best effort service. This causes problems with data that requires bounded latency, such as Voice over IP. Internal to the VHN various mechanisms can but used with either Firewire or Ethernet to guarantee maximum latency requirement are meet. For the foreseeable future the first mile connection will continue to be the bottleneck between VHN

and the Internet backbone. Therefore some form of preferential treatment of latency sensitive data, such as Multi Protocol Labeled Service (MPLS) is needed as long as the first mile connection is unable to delivery adequate bandwidth.

2.8 *Authenticated Services*

Some ISP services require authentication, such as mail and news. These services should require explicit authentication. This allows the user to authenticate to the service regardless of how they connect. If the ISP authenticates based on IP address if the customer uses a different ISP these services are inaccessible.

2.9 *Customer DNS*

One of the benefits of always on broadband is remote access to the VHN. This may include public resources like FTP or Voice over IP telephony or restricted to authorized users. Access to the home requires a user friendly URL so users do not have to enter the IP addresses directly. While many residential owners will likely register a personal domain name to decouple naming from that of the ISP the ISP should provide DNS mapping of customer name to IP address. Similar to how ISP based e-mail is handled today. This allows remote access into the VHN by URL rather than IP address. Since the ISP is responsible for changing customer IP addresses they can update DNS with the minimal amount of delay after a change.

2.10 *Customer Friendly Acceptable Use Policy*

The ISP should not be in a position to determine how and what type of data will be carried over its network. The ISP agrees not to prohibit particular services by blocking TCP/UDP ports or imposing other restrictive policies. This prevents the ISP from over optimizing the network for today's traffic usage and erecting barriers to new applications.

The ISP is neutral as far as VHN is concerned. Their responsibility ends at the customer premise equipment needed to interconnect to their network.

The scope of the customer's domain may become a controversial issue, especially with the popularity of wireless networks. Consumer ISP pricing assume typical usage patterns and high neighborhood take rate. Wireless networks threaten to upset that assumption by creating neighborhood ad hoc networks shared over wide area potentially reducing the take rate and increasing traffic.

3 *Access Issues*

Most ISPs allocate customer addresses from a pool of publicly routable addresses. This means the customer is able to both communicate with remote Internet host and act as a server if they choose to do so.

Unfortunately some service providers use techniques that limit customers to outbound connections by issuing private IP address or the use of a proxy server.

3.1 NATed ISPs

The IPv4 address shortage encourages ISPs to minimize the use of public IP addresses. One solution is to issue customers IP addresses from [RFC 1918](#) private address pool and use Network Address Translation to convert the address where the ISP connects to the interexchange carrier.

NAT is transparent to most outbound connections but effectively prevent remote access to the customer. The issue is no different then when NAT is implemented at edge of the VHN except the customer does not have to option of specifying mapping rules for inbound traffic, effectively precluding the use of servers.

3.2 Proxied ISPs

Instead of placing customers on private addresses the ISP can force customers to connect through a proxy server, as in the case of AOL. The proxy is able to precisely control how the connection is used preventing use of unauthorized services.

4 How Many IP Address

Due to the address shortage caused by the 32-bit limit of IPv4 addresses, address management has been raised to a fine art. What should have been a simple network identifier has become a valuable and scarce commodity – to be husbanded with great care. The ISP views this scarcity as an additional revenue source.

4.1 Single Address

This is the most common situation for residential customers both dialup and broadband. Each account is issued a single address. Various means have been developed to allow a single addressed to be shared with multiple devices on a LAN.

A common way to share a single address is to allocate private address from RFC 1918 and use Network Address and Port Translation (NAPT) commonly called NAT to translate between private and public address space. Proxy servers are another way to share a single connection and are often used with software implementations. The down side of using a proxy is each application must be aware of the proxy. This limits its usefulness and causes configuration headaches for the user.

Being limited to a single address does not preclude use of “server” functions. Most broadband routers and proxies have the ability to map private services to public addresses.

For example if one wanted to run a web server, the gateway device is programmed to forward all incoming TCP packets for port 80 to the web server. To the remote user the server appears to be at the public address. A limitation of this technique is only a single instance of each device can be

used. There is no way to run a second web server since TCP port 80 is already being forwarded to the first server. It is possible to use a different port to establish the connection, but this causes problems unless the remote user knows of the nonstandard port assignment.

As powerful and widespread as NAT is it is not without drawbacks. Protocols like FTP require special processing within NAT at the application level due to the way session ports are allocated. It also breaks some features of IPsec by changing address and port numbers. For more information on NAT refer to [RFC 2993](#) Architectural Implication of NAT

4.2 A Few Addresses

If one desires to operate multiple public servers most ISP charge for additional IP addresses so they are use sparingly. Having multiple public addresses allows multiple servers to be assigned unique IP addresses facilitating access from the outside world.

This is not without side effects. If the addresses are not in the same subnet inter device communication, which would normally be local, is transported over the ISP connection causing congestion and impacting speed. In this scenario NAT is still used for client services to minimize the cost of public IP addresses. This has the same ramifications as a single address, with the added complication of connecting to “public” devices locally.

4.3 Unlimited Addresses

This is the optimum situation but will require deployment of IPv6. Each customer is given a large contiguous address block. This allows each device to have a public address. As long as they are contiguous and part of the same subnet local communication stays local eliminating the degradation caused by forcing the WAN connection to carry traffic that should be local.

In this situation access policy is implemented by a firewall at the edge of the VHN rather than as a byproduct of NAT. This allows the user to control both inbound and outbound access for each device. NAT is often touted as a firewall because without forwarding rules remote access to the LAN is impossible. However NAT has no effect on outbound traffic, so a firewall is still desirable even in a NAT based network.

5 Bridged vs. Routed Network

The most common method used to allocate a “few” addresses to a customer is with a bridged connection. In effect the connection between the user and ISP looks like a LAN. VHN devices are connected to a hub or switch and the broadband modem connected directly to the hub or switch.

Bridged connections operate at ISO layer 2. Devices may be assigned addresses either statically or dynamically. If assigned dynamically the ISP operates the DHCP server. Because bridging works at layer 2 the ISP is privy to the MAC address of each bridged device raising potential privacy concerns. Another downside of bridged networks is access to the DHCP server is lost if the external network fails preventing local use in the event of an outage.

A better solution is to use a routed network. This is the norm for commercial customers but rare for home or small businesses. With a routed network the ISP allocates a block of IP addresses to the customer. The customer manages them as needed. The ISP forwards all incoming traffic bearing the customer's network prefix to the customer's router. The router in turn is responsible for deliver within the VHN. This minimizes the effect of remote network failures; the local network still works correctly of the ISP fails and it hides VHN specific information from the ISP.

6 Address Allocation

IP addresses may be assigned statically or dynamically. Static allocation is communicated between the ISP and customer through an out of band channel. The customer is responsible for configuring VHN devices. If the ISP needs to change this information the customer has to be notified and the changes updated manually. Dynamic assignment is much easier for the ISP to manage. Automatic mechanisms are used to issue addresses on an as needed basis, and if necessary the ISP is able to rebalance address utilization without involving the customer.

From the customer perspective a static address is more convenient for running servers. A static address provides a persistent address facilitating access by remote hosts.

For most home user the best of both worlds is a pseudo-static address as implemented by many Cable ISPs. The allocation mechanism is dynamic but bound to the MAC address of the user's device. This means the customer's address stays fixed for long periods of time, while still giving the ISP the ability to change them automatically on an as needed basis.

7 Encapsulation (PPPoE/PPPoA)

The Point-to-Point-Protocol is used to facilitate sharing the first-mile network by multiple ISPs. In effect the network is transformed into a virtual point-to-point connection between the customer and ISP. The ISP authenticates the customer, using the same RADIUS mechanism used for dialup. While purists cringe at the use of PPPoE Point-to-Point-over-Ethernet and PPPoA Point-to-Point-over-Asynchronous Transfer Mode (ATM) it is an effective tool to share a common physical network.

The down side of PPP is that it is an encapsulation protocol, each packet requires 8-bytes of overhead. This in and of itself is not a significant performance issue because maximum Ethernet packet size is 1500 bytes. What does cause problem, beside poor software implementation, is the reduction in maximum packet size. Applications are supposed to verify end-to-end limits on packet size. If this is not performed correctly and a 1500 byte packet is created when it is passed to the PPP layer the addition of 8 bytes will either cause the packet to be fragmented, impacting speed or be rejected preventing communication.

PPP is commonly used in conjunction with dynamic IP assignment, but it can also be used with static addresses.

A little known feature of PPP is that it allows the negotiation of multiple PPP sessions. Some carriers see this as a way to sell value add services bypassing the ISP. Separate PPP connections are

set up for say video streaming or digital telephony creating a virtual PPP connection between the customer and service provider. This has profound effects on VHN because it creates additional access methods that may or may not be IP based. It is also very much at odds with the notion of Internet end-to-end connectivity with services implemented at the network edge.

8 Asymmetric Speed

First mile access speed will likely be a precious commodity for the foreseeable future. One of the ways providers address customer preferences is tiered service. Faster speed costs more. Carriers have limited ability for downward price flexibility since most of the cost of broadband access is independent of speed. The motivation of tiered price is crude control of bandwidth consumption. This is especially important for Cable networks since each user shares a fixed medium with neighbors. The only way to provide acceptable service is to reduce the number of subscribers on a segment or reduce modem speed.

This is a reasonable mechanism for providers to segregate customers by willingness to pay.

Providers are using asymmetric speed to control how customers use the network. Due to technical and business consideration both Cable and DSL speed is asymmetric, download is considerably faster than upload. This is not unreasonable given the bulk of most residential traffic is weighted toward downloading. However an unfortunate side effect of this optimization is the growing emphasis on the residential market as a data sink rather than a network peer both “sourcing” and “consuming” information. This optimization bodes poorly for the emergence of cooperative distributed networks where data is distributed rather than located in centralized servers.

Asymmetric connection speed should not be used to thwart innovative use of the Internet and relegate the VHN as an information sink.

9 Latency

The notion of latency and bandwidth are often confused. High-speed connections can also experience high latency. A truck full of CD-ROMS is a good example. It takes a long time for the truck to arrive – high latency - but once it does it delivers tremendous bandwidth. More realistic is latency introduced by satellites in geosynchronous orbit. Another problem area is data interleaving error correction used by DSL this minimizes errors but adds substantial latency.

The most critical service from a latency standpoint is real time voice telephony. Round trip latency over 200 ms degrades quality and latency above 500 ms makes conversation virtually impossible.

10 Multilink and Bonding

Regardless how fast typical first mile connections become some customers want more. The most effective way to increase speed is to obtain multiple connections from the same ISP and bond or multilink them together. Bonding is performed at ISO layer 2, requiring complementary equipment

at both the ISP and customer. Because it is performed at layer 2 it looks like a single large pipe to IP traffic.

If the ISP does not support multilink the best that can be done unilaterally from the customer side is load balancing. A load balancing router alternates outgoing traffic between the various connections. This works well for outgoing traffic but cannot be used for incoming without a form of load balancing DNS to hand out different IP address to incoming traffic. This causes each remote host to use a different connection.

11 Multiple ISPs

As the price of broadband falls and the value of Internet access increases more residential customers will opt for multiple ISPs. This raises a number of interesting problems. Each ISP allocates addresses from a different IP address block. This requires a smart router at the edge of the VHN to handle multiple address blocks and the failure of one or more WAN connections.

Connections from multiple ISPs cannot be bonded together. They can be load balanced so traffic uses each optimally but a single session is limited by the speed of the particular connection.

12 Authentication

Many ISP services are restricted to customers, such as Usenet and mail. Access to these services may be controlled automatically or require explicit authentication. The ISP may configure these resources so that access is denied anyone not on an address issued by the ISP. This is convenient because it eliminates the need to be authenticated by the server. The side effect is if one uses an alternative ISP access to the service is denied. A better solution is to require explicit authentication to each resource. This allows the customer to access the resource regardless of how they connect.

13 Oversubscription

All ISP oversubscribe customers to control cost. Facilities are engineered to carry expected peak load. The underlying assumption is that every customer will not be using the connection at every instant. As long as facilities are properly sized this works fine.

When unexpected traffic volume spike such as during a major news event or capacity is lost as the result of a fiber cut performance plummets. If it is of short duration hardly anyone will even notice. When this type of problem occurs it is very difficult for the VHN user to know if the problem is local, the first-mile connection, the Internet in general, or the remote server. Automatic analysis tools are very useful to help the user determine the root cause of the problem.

14 QoS

The first-mile connection is likely to be a bottleneck for the foreseeable future. Various mechanisms can be used within the VHN to bound latency to acceptable limits. Improvements in the Internet backbone will address long haul congestion issues.

The first mile connection requires a mechanism such as Multi Protocol Label Switching (MPLS) to give latency bound services priority access to the link so it is not delayed by lower priority traffic.

15 Restrictive policies

In an attempt to maximize revenue ISPs have imposed a number of restrictive conditions. While the goal of each of these restrictions is to maximize short-term profitability they prevent users from obtaining maximum value from the service and discourage experimentation. It is likely such restriction will actually slow growth of high speed Internet access. In some case the restrictions become so onerous customers may even decide to forgo high-speed access altogether.

15.1 Bandwidth Hogs

The Cablecos seem especially fond of this. In an attempt to deter “bandwidth hogs” they are imposing maximum transfer restriction per month. It seems a counterintuitive marketing strategy to penalize your best customers for using the service.

Customers prefer guaranteed pricing by a large margin; predictability is more important then lowest cost. Monthly caps open the specter of a rude surprise at the end of the month. If transfer caps become widespread the VHN gateway should include monitoring capability to warn the user when they are reaching ISP imposed limitations.

15.2 Port blocking

Consumer class ISP routinely block particular ports, such as web servers or in an attempt to prevent peer-to-peer file sharing. Some ISPs block TCP port 110 and 25, preventing customers from accessing non-ISP supplied mail services.

In some cases a customer is able to select a different port to work around the bock. This becomes a game of cat and mouse between the customer and ISP. The ISP blocks ports they do not want customers to use, customers move service to a different port to bypass the block, ISP then blocks the now port, and so forth.

15.3 Prohibition of "Business Class" Services (VPN)

Cablecos are trying to establish the notion of different classes of service based on the type of traffic the network carries. One of the most controversial is prohibiting use of IPsec VPNs on residential accounts because it is a “business class” service.

15.4 Prohibition of Home Network

ISP tolerance of home networks ranges from outright ban to joint marketing deals with router vendors. Cablecos tend to be the most restrictive stemming from a business model that expects to generate revenue from each endpoint. CableLabs is working on technology to prevent customers

from using NAT based routers and place the home network under control of the Cable Company – for an extra monthly fee of course.

15.5 Prohibition of Servers

Many end user agreements prohibit servers. The notion of exactly what is and is not a server is somewhat vague. As a minimum the ISP takes a dim view of running a commercial web site on a residential broadband connection.

15.6 T&C Change Without Notice

Almost all ISP have customer unfriendly terms and conditions contracts. Basically the contract is a long list of acceptable and unacceptable uses. The definition of which may change at any time without notice. This makes it risky to depend on any particular ISP service since it may be canceled without notice at any time.

15.7 Geographical Restrictions

The Internet is a worldwide network. Since we are in the early stages it also has many of the characteristics of the Old West. Courts have yet to fully address the borderless nature of the Internet and all sorts of regional legal systems have attempted to prevent or restrict certain types of information flow.

To protect against foreign litigation ISPs require the user connect from within the US.

15.8 DMCA Safe Harbor Provision

The DMCA exempts ISPs from liability resulting from customer's unauthorized use of copyrighted material if the customer account is promptly shutoff after notification of volition by the copyright holder. The result is any copyright holder can shutdown any customer account, a court order is not required just a complaint of infringement by a copyright holder. The accused can fight the shutdown after the fact by proving ones innocent DMCA considers the threat of unauthorized use of copyrighted material so serious that is reverses the normal notion of innocent until proven guilty.

There have even been instances when the wholesale provider shut down an entire ISP because the contact person at the ISP was unavailable to respond to the complaint. To avoid potential liability the wholesale supplier pulled the plug on the ISP itself.

It is likely copyrighted material will be stored on the VHN. Accidental misconfiguration may result in exposing this material to the Internet running the risk of summary disconnection by the ISP.

15.9 Service Level Agreement

Current consumer ISPs contracts are best effort. The ISP makes no guarantee if and when data will be delivered. Commercial ISP typically provide contracts that include Service Level Agreements

(SLA). The ISP guarantees various parameters such as latency, dropped packets and time to restoration. These guarantee are rare in the consumer space since they add cost.

16 Deployment Issues

As with any new technology there have been many teething problems. This is especially true for DSL since it is designed to operate over the copper telephone plant used for the last 100 years.

Customers are frustrated when service is not delivered as planned or does not work as advertise. Identifying the root cause can be difficult and a rather daunting task for the typical user.

17 Technical Support

In general technical support is lacking and automatic diagnostics are not completely accurate resulting in lots of finger pointing. Worst since residential agreement only specify best effort service the carrier is not obligated to deliver anything to be in contractual compliance.

The ISP is faced with a heterogeneous environment. To the uninitiated user all problems look like ISP problems. In many cases the root cause is not within the province of the ISP. This can only be determined after extensive and costly troubleshooting which the low service revenues cannot justify.

18 Conclusion

The CEA and other members of the High Tech Broadband Coalition should continue lobbying for transparent broadband carriage – to maximize the value of the VHN. This should be accomplished by evangelizing ISPs that such openness is in their long-term best interest and through regulatory activity thought the FCC and Congress.

From a technical perspective the design for the VHN and access gateway needs to accommodate the wide diversity of ISPs from VHN friendly to extremely hostile.