The Only Really Open Net Is The Really Stupid Net !

The Only Really Open Net Is The Really Stupid Net !

(draft)

Nisse Husberg, Dr.Techn.

Many claim to have an “open access network” structure but in reality
this is not true in most cases. There are all kinds of limitations for
the openness.

Definition of Open Network

An open net cannot be compared to a telephone exchange (even with
digital packets). It must rather be compared to the Post office where
the packets are sent to the right destination no matter what they
contain. There are some technical limitations on weight and bulk (data
packet length) but in general the Post distributes anything.

An open access network should have no limitations for the access to
any service of any kind anywhere in the Internet. There must be no
limitations of IP addresses to connect to, protocols used or

ports
used.

Clearly, this is not true of virtually any of those networks claiming
to be “open access”. There is, however, no restrictions on the services provided to the net. If a service provider wants to limit the
access to his service in any way, it has nothing to do with the
network – as long as those restrictions are applied outside the net.

It is also possible to require any kind of special hardware or software to be installed at the customer end – as long as it is outside
the net.

“Outside the net” means often outside the terminal or router or
firewall connected to the network. Sometimes this can be difficult
to define exactly but as a general rule it must be possible to access
the net without any restrictions. If there is a customer terminal that
contains restrictions in any way or of any kind which cannot be
switched off or bypassed, then the net is not open. The necessity to
use for example Ethernet in accessing the network cannot be seen as
such a restriction because Ethernet is already such a standard protocol that the access of any service over the net hardly is restricted.

It is of course possible to have local nets which are very

restricted
in many ways but they cannot be considered part of the open net in
that case.

The Stupid Net

The only really open network is the “stupid” net. It means that the
network is ONLY moving packets to the right IP address. It does not
care at all what is inside the packets. There must not be any
“intelligence” inside the network – only at its edges. The concept was
first presented by David S. Isenberg in 1997: “Rise of the Stupid
Network”, Computer Telephony, August 1997, pp. 16-26. A later version
was published in 1998 “The Dawn of the Stupid Network”, ACM Networker
2.1, February/March 1998, pp. 24-31.

Basically it follows the KISS principle (Keep It Simple Stupid). The
original article is already over 10 years old but the ideas are even
more adequate today with very fast networks (optical fibre) and fast
and cheap electronics.

The main point is that the network should only move packets – it
should be “stupid”. If the intelligence is at the edges of the network
it is extremely flexible. Going into new applications or protocols

does not change the network at all, just the equipment at the edge of
the network. It is also possible to use different applications at the
same time without problems.

FLEXIBILITY is the most important feature of stupid networks and as
the applications change and new are invented all the time this is
really an enormous advantage. In fact we do not know what is behind
the corner in the development and the possibility to introduce new
applications very easily saves much time and money.

All kinds of control and optimisation must be outside the stupid
network because they destroy the flexibility. Optimisation is also a
work which usually is wasted in the long run. The capacity increases
so fast that no optimisation is needed. Just as memory size increased
from a few kilobytes to Gigabytes, the speed of networks is increasing
from kbits/s to Gbits/s. The limit of a single fibre is about 10000
Gbits/s which makes all optimisation quite unnecessary.

Control is also a wasteful undertaking. All kinds of checks in the
network can easily be fooled, even by schoolchildren. It is much
better to put the equipment and programs at the edges of the network,
This also improves flexibility – it is possible to use any method and

change it at any time.

This goes as well for security as for identification. Every network must be seen as a hostile environment and you cannot rely on anything. Thus building tunnels through the network with heavy coding and identification equipment is much better. Also when these methods change it does not mean that the network has to be changed. Again time and money is saved.

The stupid network (which to my mind is the only real data network as opposed to old-fashioned tele networks) of course consists of several small network – as Internet does. Especially for security reasons it is necessary to insert firewalls between the networks. This makes it hard for the crooks to get into the network but because they anyway can infect careless and unsuspecting users computers, it is necessary to protect any connection at the edge of the network. Even simple routers now have the basic functions built-in. But this is a field where improvements happen often and therefore it is a good idea to have a different router instead of a terminal with everything. Then it is easy to change the router only. Possibly this will change so that improvements can be downloaded easily.

Identification

One problem for services is how to identify a client. All ways of doing it by structuring the net are inefficient and impose unnecessary restrictions. And the worst thing is that they destroy the flexibility of the network.

Using VLAN for identifying (one client – one VLAN) is to utterly destroy the structure of the network. It is an extremely inefficient way. To use the MAC to identify a client is almost impossible because it is so easy to change the MAC that any schoolboy can do it. The same goes for using the IP-address of the client.

In a stupid net the identification must be handled outside the network. It is possible using passwords, programs or hardware and give much better security. And it is very flexible. The client can move around as much as he likes and the networks can be changed in any way without interfering with the identification. As long as the network passes the packets to the right place everything will work.

Security

Security is going to be one of the worst problems but it should not be

implemented within the network. Basic security can be handled with routers and firewalls at the edges of the network – both at the connection to Internet and at the connection of the user. Also local networks must be considered insecure.

For important tasks like banking or work over the net it is possible to use VPN tunnels or any kind of heavily coded transmission. This can be implemented in software like the Secure Shell or hard coded. A fairly secure and fast system is hardware at both ends of the connection.

One problem is, however, to achieve a common standard. It seems a little unnecessary to have different hardware for each connection. But that is a universal problem and not connected to the stupid net. It must be solved quite independently from the network design. Also, with more and more mobile users it is impossible to solve these problems within the net. They must work where-ever the user is in the whole world.