

Sharing is NOT Caring

Wireless Bridging Run Amok

A Case Study for E2Epi

Definitions:

DHCP = Dynamic Host Configuration Protocol – how a host finds its IP address

NOC = Network Operations Center

MAC = Media Access Control – the "hardware address"; like Ethernet address

LAN = Local Area Network – a high-speed network connecting computers in a small area; Ethernet is the most popular/famous example

ARP = Address Resolution Protocol – how to map an IP address to a hardware address

vLAN = virtual LAN – how to make a single physical Ethernet look like multiple Ethernets (to partition for security, for example); see LAN.

Thursday, September 4, 2003. The campus DHCP server runs out of addresses for a wired network in the school of business while, at the same time, all new connections to the campus-wide wireless network are failing. Inspecting the DHCP server's log files reveals a single MAC address on the lease-exhausted wireless network is repeatedly requesting a DHCP lease. Dwight writes: "There is some unknown problem with this address and it is implicated in the wireless problems we are having campus wide."

When the question is raised "How does one MAC address get multiple IP addresses from the DHCP server?," Haiyan replies that there is an option in the DHCP server config file (currently enabled on the server) that allows one-lease-per-client flag. If this flag is enabled (which is the default), whenever a client sends a DHCPREQUEST for a particular lease, the server will automatically free any other leases the client holds. This presumes that, when the client sends a DHCP REQUEST, it has forgotten any lease not mentioned in the DHCP REQUEST. An example might be that the client has only a single network interface and does not remember leases it is holding on networks to which it is not currently attached.

They looked to see if the requests for licenses were spaced out or were in short bursts; Haiyan also noted that the leases for dorms are all for 8-hour periods. Before it was blocked, there were many requests from that particular MAC address

– the lease file shows this as well, but it looks like different clients were all using the same MAC address! The short-term solution was to turn off the wireless vLAN to the affected area.

However, the problem reoccurred on September 24 to a *different* MAC address. This time, a single MAC address had 92 wireless leases and multiple IP addresses assigned to it! The support center received many calls about the problem – from other users who were unable to acquire licenses because they were all "in use" – which brought the problem quickly to Haiyan's attention. He immediately blocked the offending MAC address from the DHCP server in an attempt to stem the problem.

At this point, it was clear that what was occurring was an unauthorized bridge between the wired and wireless networks called Internet Connection Sharing. Dwight found the bridge in a specific room – the data jack was a new cat5 set for auto. He disabled the port, which stopped the problem, but, because the port was a wireless access point, they still did not know the room number of the person who had bridged the wireless network to the building vLAN. Dwight commented, "Turning off the wireless access point may be a good short-term solution but how long can we keep it off? We need better tools for this!"

Steve got involved, at this point, and suggested that they determine the subnet where the offender is located by "sniffing" the wireless LAN for gratuitous ARPs

from devices not using the wireless LAN addresses. He also suggested telling the DHCP server the range of MAC addresses used by wireless cards – this might tell which subnet the offender was on, as well as preventing this from affecting other wireless access points.

These were efforts to locate the offender but, as had already been discovered, there were multiple offenders, which most likely meant that it was a host configuration problem. What appeared to be happening is that, when the wired and wireless networks were bridged, the DHCP request sourced on the wireless side goes to the DHCP server from *both* the wireless *and* the wired networks. If the response to the laptop is from the packet that traveled over the wired network, the laptop will have no connectivity – and this will happen campus-wide for the wireless network!

On September 25, Steve came up with a potential solution: put a BSD (Unix) machine in front of the DHCP server; intercept every DHCP request and hold it for some small amount of time (0.2 seconds). If a second request comes from a different network but using the same Ethernet MAC address, delete the request coming from the non-wireless LAN, forward the wireless DHCP request, and alert the NOC via e-mail as to which subnet has the bridged device.

After putting this stop-gap measure in place, Steve began contacting members of the Internet2 community who might have had similar experiences to see if they had other solutions to offer. William responded that he had discovered that Microsoft Windows ARPs for “a 192.something address” before it brings up a dialog box asking about bridging. Apparently, the address being “ARP’ed for” signals that the machine is part of some Windows home network that wants

to participate in Internet Connection Sharing. To protect against this, William’s campus has PC’s with these assigned addresses connected to each of its wireless vLANs in the broadcast domains. When the clients ARP for the address at startup and hear an answer, they do not bring up “connection sharing.”

William’s campus has found this problem most prevalent in Windows XP (Microsoft changed the behavior of bridging) without deliberate user configuration (a default). He believes that Mac OSX requires users to actively do something to turn on bridging, but this has still found its way onto campus. William provided the offending addresses:

Microsoft: 192.168.0.1

Apple: 10.0.0.1

Recommendations

Steve recommended that his campus set up a secondary IP address on the wireless vLAN at 192.168.0.1. This is a fairly inexpensive and easy-to-implement solution.

Dwight, Steve, and Haiyan were members of a technical support group at their campus; members of the Internet2 community can join the **E2Eperf Interest Group** for access to a broader range of technical expertise.

What is the E2Eperf Interest Group?

The E2Eperf Interest Group is an email list for those interested in end-to-end issues. It was started to help form Internet2’s End-to-End Performance Initiative (E2Epi) and continues today as a place for discussions of end-to-end performance problems. To subscribe to the mailing list:

- Send a message to the address: listproc@internet2.edu
- Do not include anything in the subject line
- In the body of the message include the line: subscribe e2eperf-interest FirstName LastName [i.e., subscribe e2eperf-interest John Smith]

Once you are subscribed, you may send a message to the list by using this address: E2EPERF-INTEREST@Internet2.edu

Archives of the E2EPERF-INTEREST list can be found at: <http://archives.internet2.edu/>

What is E2Epi? E2Epi is Internet2’s End-to-End Performance Initiative. For more information about us, see our website at: <http://e2epi.internet2.edu/>. There you’ll find links to projects (such as the E2E piPEs, OWAMP, and H.323 Beacon) and presentations, as well as other information related to performance issues, measurement tools, upcoming events, and related activities.