

Troubleshooting Cisco Routers

It's all too easy for a simple configuration error to lead to problems on a routed network. With a little knowledge, tracking down a fault on a Cisco router needn't be daunting.

By Neil Briscoe

In PCNA 105 [File T1822] I wrote an article about configuring Cisco routers. That article largely assumed that once you'd decided on the equipment you required, and configured it appropriately, you could plug it into your network and your wide area connections and everything would work nicely. That's the theory, but, as anyone who has ever done any router configuration will tell you, theory and practice are two parallel lines.

If things do go wrong, hopefully the tips and tricks in this article will help you isolate and fix the problem.

As before, this article is for those providing routed connections for their company, rather than those troubleshooting Internet connection links - since the ISP will provide you with much assistance in getting your connection set up.

One area where you have no control whatsoever, other than in the ordering, is in the provision of the link itself. Once you've ordered it, your local telephone company (telco) will send in their engineers to provide the link. Once they've done their end-to-end tests (if it's a leased line) or made a network call (if it's an ISDN line) they'll tell you that the line is ready for you to plug your Customer Premises Equipment (CPE) into.

Normally, once they've done that, they'll expect you to prove that any faults you might find are theirs rather than yours.

ISDN

Let's start with ISDN Basic Rate Interfaces. In my experience, these are the biggest cause of problems - espe-

cially if the site on which you're installing it ordered the line from the telco months before they asked you to install the routers to communicate over the link.

In the United Kingdom, initial diagnosis is easy. The ISDN2 lines provided by any of the UK telcos end in a termination box which contains a green LED in roughly the centre of the box. It's known as the "Customer Confidence Light" and, if it's lit, it is supposed to indicate a functional line. My view is that it indicates an operational line 99% of the time, since I have come across those situations where the LED is lit but the line still doesn't work. You have to be sure of your diagnosis before you call the telco, though, since they will swear that if they find the fault is yours, you'll be charged. If, on the other hand, the light is out, head straight for the faults number. A non-lit LED means your line isn't going to work, until the problem is resolved.

In North America and other countries, unfortunately, such initial diagnostics are not available at all, and so you will only be able to rely on such diagnostics as your router may give you. The most important thing to remember there is that you have to pro-

vide the terminating electronics, either built into your device or as an external extra. The device in question for ISDN2 lines in North America is called an NT1. Make sure it's either built into the router you bought or go out and buy an external one.

Similarly with X.21 links, at least those over 64 Kbits/sec, the boxes provided in the UK by the telcos contain lights, all of which should actually be non-illuminated, apart from the power light. There are lights for self-diagnosis, network (by which it means the telco network) and equipment (the CPE). If any of these is lit, there is a problem with either the box itself, an inability to see the telco network, or an inability to see the router you think you've just plugged into it. The same box is used for Primary Rate Interface ISDN lines, and the same rules apply.

So, if you're in the UK, first, check you have a power light on the box, and then check that you have no other lights lit. If you do have other lights lit, for network, or electronics faults, phone the telco and tell them.

If it's the CPE light that's lit, then initially you have to do some troubleshooting yourself. Is your router plugged in? Have you switched it on?

"There are lights for self-diagnosis, network (by which it means the telco network) and equipment (the CPE). If any of these is lit, there is a problem ."

If yes to those two, does the router you've plugged in have any diagnostic lights? If so, check them. On Cisco boxes, you normally get a green LED on each interface to indicate that it's functional. It will be out if it's damaged, or if it's not been set to come up.

Finally, in the initial diagnostics stage, we have the cables. Are they plugged in securely? If so, you might want to meter them, or feel for a break

in the cable. Are they wired correctly?

I once had a case where I'd ordered the correct Cisco cable for an ISDN PRI and the cable was supplied, with the correct connectors, and was even labelled with the correct part number. When, after much diagnosis, I prevailed upon Cisco support to send me a new cable, we found the Cisco-supplied one was wired entirely differently to that supplied by our

distributor. Our distributor, we found, had wired the cable themselves, and got it wrong. Always get your Cisco connectors supplied in official Cisco blue and accept no substitutes.

Next Steps

So, that's the basic diagnostics covered. Now we're at the point where we have the router plugged in, all the indicator lights appear to be correct, but things are still not working.

First things first. From a PC on the same LAN segment as the router, check to see that you can ping the router's LAN interface. If so, at least that's fine. If not, start by checking the PC's setup - it's quicker. Ensure it's configured with an IP address in the correct subnet, is using the correct subnet mask, and, whilst you're at it, check that you have configured the stack with a default gateway address - which should, of course, be the address of the router's LAN interface.

If that looks right, it's time to start doing some diagnostics on the router. First check, of course, is to ensure that the router's LAN interface is configured with an address in the same subnet as the PC, and that the interface has been enabled.

I once made the mistake of forgetting to type "no shutdown" on a Cisco router, such that the interface was simply not listening to packets. That kind of mistake should be caught by the ping test mentioned earlier. It can also be caught by doing a "show int e0" (or whatever name the Ethernet interface happens to have) and if this shows "Interface Administratively shut-down" then you know that you've made the mistake I once made.

If the Ethernet interface appears to be up, make the same check of the WAN interface. If it's an ISDN line, this might show "Line up, Protocol up (spoofing)", which doesn't tell you an awful lot. If it's a serial line, however, what you should see is "Interface up, Line up". If you see "Interface up, Line Down", then you have a problem with your telco and it's time to hit the fault line again, possibly after checking that you've plugged the X.21 cable in correctly. Another anecdote: a customer of mine managed to plug an X.21 cable

Debugging On A Cisco

Cisco's IOS, which stands for Internetwork Operating System, and runs on most, although not all, Cisco products, has an absolute wealth of debugging commands, and I most certainly won't be covering them all here. Instead, I'll just give a few examples relating to the tests discussed in the article.

First of all, you need to ensure the router is actually logging the debugging output somewhere. Some people like to use the "logging console" command to enable output on their terminal session - but I prefer to use "logging buffered", which causes the router to write the output to a circular in-memory buffer. That doesn't clutter up the screen whilst you're typing pings or telnets, and you can also use the following little ploy. If you type "no logging buffered" followed by "logging buffered" you turn off, and then turn on, logging, but in the process you clear the buffer of the previous attempt's output. This can be useful when, based on the first attempt to fix something, you find you still haven't fixed the problem. Without the aforementioned ruse, you suddenly find yourself wading through the previous attempt's output before getting to this attempt's output, and you can get confused as to whether anything has changed from last time to this.

Note that all of the commands discussed so far have to be typed from configuration mode.

Once you've arranged for debugging output to be logged somewhere, you leave configuration mode and actually enable the debugging output you want. Before you do that, a quick use of "show debug" is useful, as it will list what debugging output you have currently asked for. It's best to do one bit of debugging at a time. If you don't and you try to debug both ISDN Q931 packets and PPP authentication packets at the same time, outputs from both will be displayed/logged but what you'll get is a confusing mess which is less help than the two bits of debugging done separately.

So let's assume we have an ISDN line that our telco assures us is functional - they always say that - and we haven't been able to make a connection. First thing then, is to see if the router is even attempting to dial out to the number it's supposed to, and if so, if we can see any reason why it isn't working. So the command you might type would be "debug isdn q931". After that, try pinging a device at the other end of the dial-up line. You can do that from either the router session or from a machine connected to the same LAN segment as your router. If you used the "logging console" command, you'll get the output displayed at once, interspersed with either full stops ("."), a failed ping) or exclamation marks ("!", a successful ping).

If, on the other hand, you used the "logging buffered" command, you can now examine the buffer with "show log". You may discover the router attempted to dial out a few times - this is quite common if it doesn't get a connection the first time. Unfortunately the reason codes given for the line being cleared aren't always easy to understand, but you should be able to see the number it tried to dial, and whether it received a "normal call clearing" or whether there was some other reason for the failure.

Cisco Routers

in upside down. No, they're really not keyed. Turning it up the right way fixed the problem.

Delving Deeper

Now we're getting deep into the realms of difficult territory. We've checked the telco indicators, checked our router interfaces, checked the cables, re-checked our interfaces, and still the router is not working - we still can't see the other end. It's at about this time that you really wish you could be at both ends of a routed link at once. If you happen to have a colleague who is at least reasonably knowledgeable and can be talked through commands, and if you're responsible for both ends of the link, it's time to ship someone up the far end. Not quite as useful as cloning yourself, but it will have to suffice.

Next we come to encapsulation. If you were using anything other than a Cisco, this wouldn't be a problem, since most other products only sup-

“If you're dealing with a Cisco at each end, you need to make sure that both ends are using the same encapsulation. Normally for ISDN links, you want to see an “encapsulation ppp” sub-command under the relevant interface”.

port PPP encapsulation on their links. But if you're dealing with a Cisco at each end, you need to make sure that both ends are using the same encapsulation. Normally for ISDN links, you want to see an “encapsulation ppp” sub-command under the relevant interface. If it's Cisco at both ends of an X.21 link, you probably don't want to see this (if you don't it will default to

Cisco's HDLC encapsulation), but, if you have a Cisco at one end and a non-Cisco at the other end of your X.21 link, you probably do want to see the encapsulation command. It's simply a case here of, if you spot mis-matches, fixing them by re-configuring one or both routers.

If things still aren't working, and you're on an X.21 link, it's time to go check your routing commands (see below). If you're on an ISDN link, then the next thing to check is that your authentication methods match.

With PPP links, you can choose to use either PAP or CHAP authentication. You need to check that a) you're using the same authentication at each end, and b) that you've got the username and password entries correct at each end.

For a CHAP authentication, normally, router1 will send “router1” followed by the password. So on router2, there needs to be an entry for user router1 with the password. On router1, there needs to be an entry for user router2 with the same password.

For PAP authentications, you just need to ensure that the correct username and password are sent to the remote end.

As a side note on authentication, you might appreciate this anecdote of recent experience. If you happen to have a suitably large network, or support many incoming lines, perhaps because you're an ISP, you might be using a Radius server to manage the database of all the user/password combinations (together with certain other information, including routing information and an indication of how

Debugging On A Cisco (continued)

Things to check at this point include whether the router attempted to dial the correct number. Typos in a configuration might mean it's not dialling the correct number for the remote end, which is another good reason for your pings failing.

For another example, we'll assume that your connection dialled the correct number, and received a normal call clearing message, but nevertheless the ping was unsuccessful. That's nicely absolved the telco of blame, so the next thing to check is your authentication.

Turn off your Q931 debugging command with “no debug isdn q931”, if you're using buffered logging, then recycle the log in the manner described earlier, and then it's time to debug the PPP authentication procedure. Type “debug ip ppp authentication” and try your ping again. I once used this to good effect when assisting a client with their unmanaged connection to an ISP. The ISP support desk had assured me that they used PAP authentication, and when I did this bit of debugging I was able to phone them back and inform them that in fact their router was doing CHAP authentication and was sending out CHAP authentication requests which, of course, I had configured our router not to use.

Debug Commands

On Cisco kit there is an extensive set of commands, and it's impossible to remember them all. Fortunately, the whole of IOS has a helpful little facet. Type part of a command followed by a question mark and IOS will list the options for the next word.

So, if you type “debug?” you'll get a list of all the first operands, together with a description of what they're for. This means that with a knowledge of networking, and a brief bit of training on Cisco routers, if you want to debug ISDN Q921 packets, as opposed to the Q931 packets we discussed earlier, you can type “debug isdn?” and find that there are actually three choices.

many ISDN channels can be used by that user).

At one site I worked on, the ISP in question had its Radius server hosted on a Unix box. Most of the operators doing the configuration had Windows/NT workstations. One person had used Windows cut and paste to copy the password from a CPE router's configuration file into the Unix editing session they were telnetted into. This led to an invisible character being pasted into the Radius file and compiled into its database. It took me and a person from the core routing team some time to debug that particular one.

Once you've ensured that you're happy with the authentication matching at both ends, it's down to your router's diagnostic capabilities as to whether you can diagnose what's going on, or just check and see if it works. It is a sad fact that many low-end routers simply don't possess very much in the way of diagnostics, and about all you have are traceroute and ping.

Another thing you have to consider when dealing with ISDN links is what are known as "Packets of Interest." A Packet of Interest is one which will cause an ISDN line to raise if it happens to be down when the packet is sent.

In the UK at least, call charges being what they are, it is a common practice for router installation engineers to configure things such that ICMP ECHO_REQUEST packets (a ping to most people) will not cause the line to rise if it is currently in a down state. This is good during normal operation - the last thing you want is a vicious user raising your line unnecessarily - but it's anathema to the person tasked with debugging a problem.

So, whilst in debug mode, ensure you have no ACL in place which causes your line not to come up - because it can really foul up the diagnostic process. If you do have an ACL - remove it - but keep a note of it so that you can put it back later. On Cisco kit it's easier than that, as you can change one command so that it no longer uses an ACL but permits any IP packet to raise the line.

For the purpose of this narrative, we're now down to having checked everything but the routing. I tend to prefer EIGRP routing to communicate

“Once you’ve ensured that you’re happy with the authentication matching at both ends, it’s down to your router’s diagnostic capabilities as to whether you can diagnose what’s going on.”

between routers, and RIP routing to talk to the servers on the network segment which should be able to understand the protocol.

Check Your Routing

Check your configuration, and ensure that a) you're using the same routing protocols at each end (this might involve a phone call to the person at the other end) and, where EIGRP is concerned, ensure you are using the same ASN at each end. In addition, at your end at least, check the routing table and make sure that you can actually see a route to the far end. You should certainly see one if you've used any static routing commands, but may not, if you can't "see" the other end.

On a large network, something may have gone awry with one of any number of routers and you may be in the midst of a routing transition. On the Internet, things are arranged such that eventually, this sort of thing heals. On a private network, it's entirely down to how well designed the network is in the first place, as to whether the routers will be able to "heal" the temporary outage once they've all told each other about the failed link.

Where RIP is concerned, ensure that you're using the same version. Either both ends should show that they're using RIP version 2, or both ends should show no such sub-command. Mixed versions can cause problems. If you have a non-Cisco at one end of the link, try to find out which RIP version it uses, and configure your Cisco to suit.

Once you've checked all of these things out, you should have been able to diagnose the problem, but despite

the fact that I've been doing this sort of work for some time now, I'm consistently amazed by just what else there is to go wrong.



The Author

Neil Briscoe (neil.briscoe@itp-journals.com) is a networking consultant specialising in Cisco kit.