

Bulletproof your iSeries network connection

By Larry Bolhuis

Have you ever lost the TCP/IP connection to your iSeries, prompting a torrent of calls to your help desk and sending jobs spinning into joblog land? If you haven't, you're just lucky. If you don't take precautions, however, your time will come!

Many things can cause your iSeries to disconnect from the LAN. Because your local users, and probably your remote sites, connect through the network, upon failure, your server will disappear completely from corporate radar. The most likely reasons for this unwelcome disappearance are:

- A bad or disconnected cable
- A bad switch or hub or port
- Lost power to the switch or hub
- Improper configuration of the network switch or OS/400 line description
- A bad card
- A failed IOP or bus
- The accidental end of a TCP/IP interface
- A variance off the LAN line
- An OS/400 crash

Proper configuration and implementation can prevent all of these, except the last one, but fortunately, OS/400 crashes are as rare as 80-degree days in January in Michigan!

Let's start by looking at how each of the items in the list above can be avoided or at least minimized. Later we'll see how to configure your system so that, if and when these events do happen, the help desk phone will stay quiet.

A bad or disconnected cable

If you have heard me speak at COMMON or other groups, you know that one thing I strongly advocate is doing this right the first time. So don't purchase the cheapest cables available and—I beg of you—do not make your own! Unless you are extremely good at making cables and have the correct type of wire, ends, and tools, you are basically just installing a land mine that will go off—it's just a matter of time. Also, you should always properly route and label your cables. Poorly routed cables are stepped on, hooked, pulled, and otherwise abused, eventually causing failure. Improperly labeled cables just whisper, "Pull me, pull me!" It's also not a bad idea to color-code your wires—just select a color for all critical servers and don't use the extra wires of that color for other uses.

Finally, be sure you have securely connected the cables. I've often traced intermittent connection problems to a wire that is not fully inserted into the switch or NIC. As soon as you insert it, try to pull it out to be sure it is secure—*before* you put it into production!

A bad switch, hub, or port

If your switch, hub, or port is bad, there are several things that can help. Start again with good equipment, preferably something with SNMP-monitoring capabilities so that you can be alerted when things begin to fail, instead of when they've gone away altogether.

Use good equipment. Maybe you don't need a Cisco 6500, but that no-name, no-management, no-redundant power unit is much more likely to fail than one built to handle core network loads.

Avoiding auto detect isn't just for low-end equipment—well-known brands such as Dell and Cisco are also affected by this problem. Some switches also support “portfast” (or fast negotiation) of spanning tree. This should also be enabled if it is available so that the port switches from the blocking state to the forwarding state in as little time as possible.

Remember the keys here are:

- Use switches
- Avoid auto detect
- Match the OS/400 line description to the switch
- Trust but verify

A bad card

Although it is rare to get a bad card, it does happen. There is only one way to prevent a bad card from sending your data to the bit bucket, and that's to have more than one card and to have it properly configured. We'll get there in a few paragraphs.

A failed IOP or bus

A failed IOP or bus is even less likely than a bad card, but it also happens. I ran into an IOP failure on an i810 just last month. It turned out to be a PTF fix, but failed is failed. The solution here is more than one card, on different IOPs, on different buses, properly configured. Sounds like the same solution as above, and it is. Patience!

The accidental end of a TCP/IP interface

A TCP/IP interface doesn't usually end accidentally either, but it can be fatal if you have only one interface. Clearly, you need to restrict the number of people with *IOSYSCFG and *JOBCTL authority. The solution to this problem will also be addressed shortly.

A variance off the LAN line

It is hard to vary off the LAN line because OS/400 warns you frequently that the line is in use. But if you work hard enough at it, you can vary off the LAN line. And if you have only one line, it is fatal.

An OS/400 crash

An OS/400 crash hurts no matter what. Until we get complete stateful replication (where failover is completely seamless, even to the application user) from the high availability solution vendors, or until it is built into OS/400, this is also fatal. Fortunately, it is extremely rare.

The solution to these network problems

What is the silver bullet that solves all these problems? Redundancy, of course! And how do we achieve redundancy? Is it as simple as multiple cards? Not quite.

Before you start you need:

- Multiple LAN cards
- Multiple switches
- A cable for each LAN card

- A new IP address for each LAN card, even the current production one
- OS/400 V5R2 for complete implementation
- A PC equipped with a full iSeries Navigator complement with current fix pack
- A dedicated system

Begin with multiple LAN cards—two at a minimum, three is better. Make sure they are located on separate buses to ensure that they are also on separate IOPs, and if you have multiple towers to spread them around in, so much the better. You may need help from IBM or your local business partner; just make sure your helper understands your requirements and is qualified to work with your hardware!

Next, connect the multiple LAN cards to *different* switches with good-quality, properly routed cables. All the line redundancy in the world is worthless if you leave a single point of failure in the network. Configure the switches for full duplex and speed to match your iSeries LAN cards.

With the physical requirements out of the way, follow the steps outlined below.

First, create a new LAN line description for each LAN card. Remember to use full duplex and the highest speed supported by your hardware. Vary them on and check the switch and OS/400 for the reported speed and duplex. Do not continue until these match! See the example in **Figure 1** (of course your resource name will almost certainly be different).

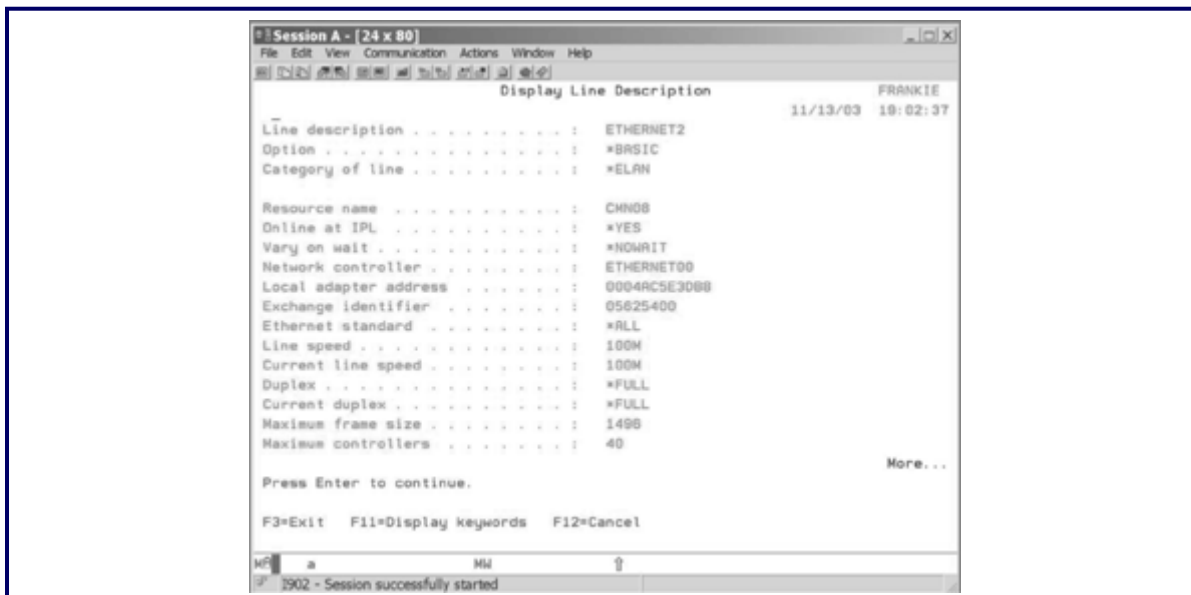


Figure 1. An example of a new LAN line description

Figure 2 shows the starting point in iSeries Navigator. Under Network, open TCP/IP Configuration, and you will see IPV4, IPV6, and Lines. You cannot maintain lines from here, but you can see status and line throughput.

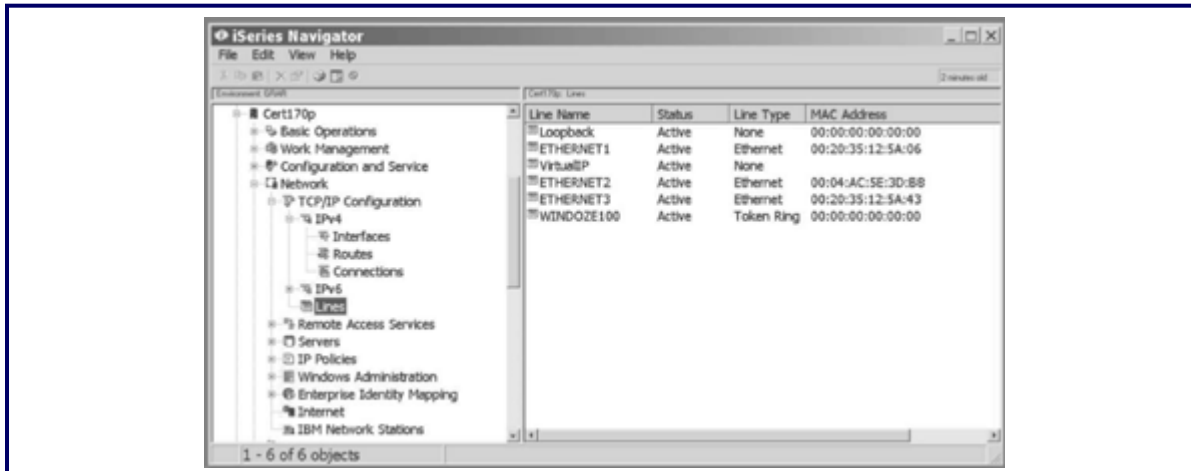


Figure 2. The starting point for IP interface configuration in iSeries Navigator

Second, add a new IP interface (ADDTCPIFC command or right click on Interfaces) for each LAN card and line that you now have available. This address must be in the same subnet as the production IP address that you want to protect from failure. This address will not be used for production and likely won't be added to your DNS for that reason.

The next thing you need to do is bring these new interfaces online and create a connection in iSeries Navigator to one of them. We are about to delete your production IP address! Once you have used the new address to reconnect to iSeries Navigator, you can proceed. Third, end your production IP interface (ENDTCPIFC or right-click and Stop) and delete it (RMVTCPIFC or right-click and Delete). Because it is currently tied to only one LAN card, we must recreate it to fix our redundancy issue.

Fourth, recreate the production IP interface. You can use the ADDTCPIFC command with the following parameters:

- Line name of *VIRTUALIP
- Subnet mask of *HOST or 255.255.255.255 (these are equivalent)

Or you can use iSeries Navigator. Right-click on Interfaces, select New Interface, and then select Virtual IP. The wizard will come up and ask you for the IP Address, subnet mask, and a description. When you create this as a virtual IP interface, the address is associated with your system rather than with a specific piece of hardware.

Regardless of which method you use, you will need to use iSeries Navigator to complete the recreation of your production interface. Locate the new production interface in iSeries Navigator, right-click on it, and select Properties. When the panel comes up, as shown in **Figure 3**, select the Advanced tab.

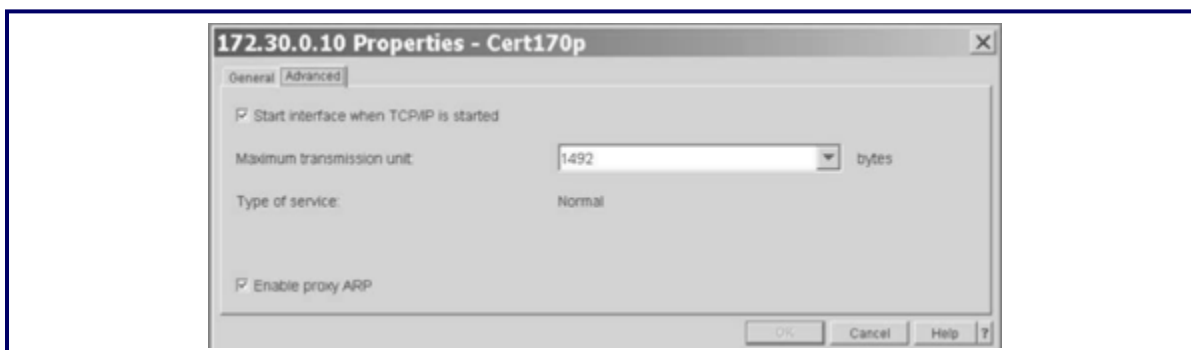


Figure 3. TCP/IP interface Advanced properties

There are two important items on this pane. The frame size must match the LAN line descriptions created above. For 10/100 interfaces, it's usually 1496 and for Gb interfaces, 8996. The system can't retrieve this value from the line description (the default for non-Virtual interfaces) because this interface isn't associated with any particular line description.

The most important piece is to check the box to enable Proxy ARP. The reason this is so important is that, since your production IP address is no longer assigned to a physical line, the system would normally not respond to any ARP requests for that address. This would make your system unreachable by this virtual address. With this box checked, the system will use one of the physical interfaces with an IP address in the same subnet to respond by Proxy for the virtual address. This means your system is reachable again!

If the interface that's being used to proxy the virtual address fails, the system will immediately toss a gratuitous ARP for the virtual address out on one of the remaining interfaces. Any clients communicating with that address (which should be all of your clients) will see the new MAC address and use it for all subsequent communications. This happens in milliseconds and you should never know.

Of course you'll want this interface to start when TCP/IP is started, so check that box too (if it's not already checked). Finally, press OK to complete the interface.

The fifth step is to add special routes to the local network for each of your LAN cards. Do this to let OS/400 load balance outbound traffic on all LAN cards. For each interface added in step two, add a new route with the following parameters:

- Route destination is the local network, so if your network is 172.30.0.0, that's what you enter here.
- Subnet mask matches that in the interfaces created. For a 172.30.0.0 network, it would typically be 255.255.255.0 since this is a class B network.
- Next Hop is the address of the interface this route is to use. This matches each interface from the second step. So if you created 172.30.0.230 for the first interface, that's what goes here.
- Preferred binding interface must match the Next Hop address (i.e., the 172.30.0.230).
- The duplicate route priority would be anything except 5 and must be the same for all routes added in this step. I usually use 6.

You can do this by using the ADDTCP RTE command or by right-clicking on Routes in iSeries Navigator. The resulting route in iSeries Navigator will look like that shown in **Figure 4**. Also be sure to check the Advanced tab for the rest of the parameters as shown in **Figure 5**.




Figure 4. The TCP/IP network route properties in iSeries Navigator

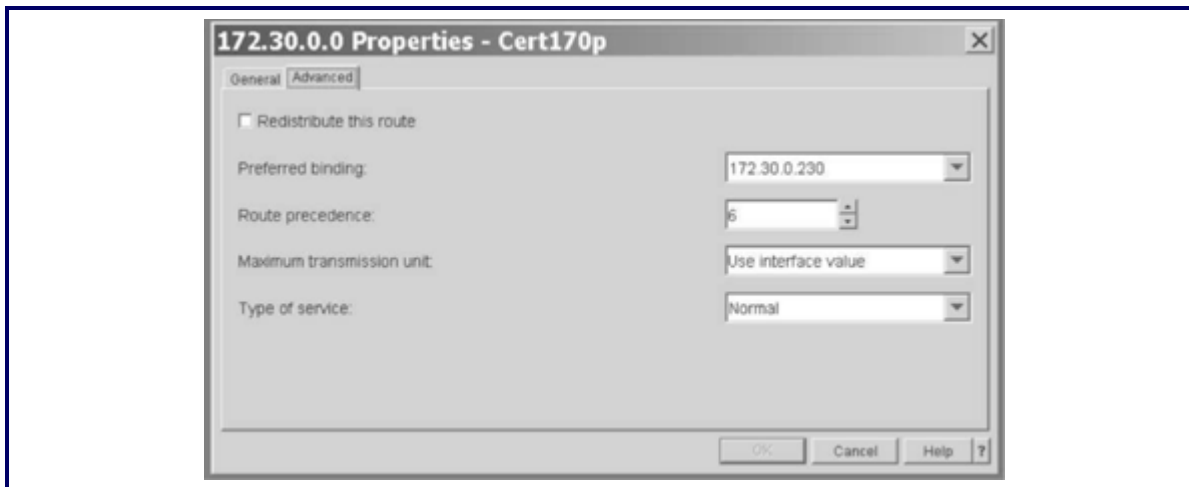


Figure 5. The TCP/IP network route Advanced properties in iSeries Navigator

Finally, bring up TCP/IP and start the rest of the interfaces created in step two. Be sure to start the production IP address. Your redundant connections are ready to keep your traffic safe.

We're in the home stretch now. The next step is to test the redundancy before you release the system to the users. To do this, connect to your system at the production IP address and start up at least one 5250 emulation session. One at a time, disconnect your LAN cards and watch the connection stay up! Of course, you will want to make sure that each line and interface is back online before you pull the next wire. The magic is that gratuitous ARP mentioned above, and it happens so fast that communications never fail. Do not go live until the test is completely successful. You just have to know that it is working, or you'll always have that nagging thought in the back of your head that it just might fail.

You must do one last thing to complete the package. With this redundant connection, you could have a failure and not even know it—that is the idea, after all! But just as with RAID disks, while one failure is annoying but tolerable, two failures will be fatal. So the last piece is to ensure that you have a way to know when the first failure occurs. This can be one of the iSeries products that watch QSYSOPR or an external product that monitors the availability of the interfaces. iSeries supports SNMP monitoring for those shops where this capability exists, or you could use a simpler network monitoring product that pings interfaces periodically. An important feature, no matter how you do it, is the ability to page a key person when the failure occurs. Note that if you are watching the interfaces from the outside, you'll want to monitor the IP addresses assigned to each physical line in the second step. Monitoring your production IP address won't tell you about the first failure until the last one, and it's too late then!

You can surely see that the cost involved with getting redundant is minimal compared with the grief of experiencing a failure and all the hassle of recovery. Now that you know how to keep your iSeries communicating on the LAN through almost every possible failure, you no longer have an excuse for that next LAN failure.

EJ