

Tip: IBM i Aggregate Ethernet Lines

Creating and managing aggregate Ethernet lines on IBM i 7.1

Jared Draper

July 03, 2012

IBM i can now support aggregate Ethernet lines, also known as *EtherChannel*. EtherChannel makes your system more resilient to failures and adds performance to its network capabilities. This tip describes the prerequisites and implementation for taking advantage of this new feature.

Link aggregation

802.3ad and 802.1ax are the IEEE designations for link aggregation; 802.3ad was adopted in 2000, and 802.1ax was adopted in 2008 to resolve some discrepancies in the 802.1 layer.

What is aggregation?

In the past, the IBM i operating system has only provided redundant Ethernet capabilities through the proxy Address Resolution Protocol (ARP) or "automatic failover" between line descriptions. Unfortunately, this implementation is neither the industry standard nor practical in an enterprise environment. Link aggregation, defined by the Institute of Electrical and Electronics Engineers (IEEE) as 802.3ad or 802.1ax, provides both redundancy and performance advantages. When properly implemented, link aggregation can both increase the resiliency of your system to network failures and provide a significant performance benefit. This tip refers to the technology as *aggregation*; however, it is also known as *EtherChannel*, *teaming*, or *trunking*.

Advantages of aggregation

There are three major advantages to aggregation over redundancy:

- **Resiliency.** Aggregation increases your system's resiliency by eliminating three single points of failure. First, the Ethernet port on your system can fail. Second, the Ethernet cable itself can fail. And third, the switch or switch port your system is connected to can fail. Aggregation can overcome all of these failures without any impact to your system or its users.
- **Performance.** With aggregation, TCP/IP traffic is allowed to traverse any of the available paths to the switch. The traffic is spread across the resources according to a configured preference, which means that each 1 Gbps line adds to the overall throughput capabilities of your system (for example, two 1 Gbps lines equals 2 Gbps of theoretical bandwidth). You can add up to eight ports in an aggregated line configuration. This is also true of 10 Gbps Ethernet

lines, meaning that your maximum throughput could *theoretically* be as much as 80 Gbps of bandwidth. This is a big advantage over redundancy, in which traffic would simply flow over one or the other physical connections, limiting you to the bandwidth of that single connection.

- **Routing simplification.** With redundant Ethernet lines on IBM i, routing was difficult, because the setup required an interface to be assigned with the subnet mask 255.255.255.255, or 32 bit. This interface acted as the master and would use proxy ARP to point to one of two physical interfaces, each with its own IP address. This configuration becomes a problem, because you cannot route traffic out that "master" interface because of the 32-bit subnet mask. Often times, that "master" interface is the IP Domain Name System (DNS) point that caused confusion or even made the scenario impossible, because inbound traffic would come across one IP address, but outbound traffic could possibly come from one of the two other IP addresses. With aggregation, this problem is solved, because the IP address points to a new media access control (MAC) address that is unique and therefore can be assigned to whatever subnet is necessary to route traffic correctly. The TCP/IP setup is exactly the same as if the underlying line description were a physical device, making routing and IP address assignment much simpler and cleaner.

Implementation

There are four prerequisites necessary for implementing Ethernet aggregation on IBM i version 7.1 with technology refresh 3:

- At least two gigabit Ethernet physical ports assigned to the partition. This assignment can include one host Ethernet adapter (HEA) port if it's the only logical port assigned to that physical HEA port.
- The following program temporary fixes (PTFs) must be applied: MF53900, MF54074, MF54188, MF54229, MF99003, SI42593, and SI42997.
- The ports to be used in the aggregate line must be connected to an EtherChannel capable switch or switch pair. When using a switch pair the attached switches will need to be configured in a Virtual Link Aggregation Group (VLAG) also known as a *stacked switch pair*. Your network administrator will need to enable the switch ports to use EtherChannel in a **static** configuration with Link Aggregation Control Protocol (LACP) off.

Implementation steps

The first step in implementing Ethernet aggregation on IBM i is to identify the communication resources you'll use as part of the aggregate resource list in your new line description. To do so, run the `WRKHDWRSC TYPE(*CMN)` command to list all available communication resources. Look for resources with the text description "Ethernet Port," and record the resource names to be used (for example, CMN01). For demonstration purposes, resource names *CMN01* and *CMN02* are used in this tip.

Now, create your new line description with the CMN resources specified in the `AGGRSCL` parameter. Other parameters you need to be aware of are `LIND`, the name of the line description; `RSRCNAME`, which should be `*AGG` to specify that this is an aggregate line; and `AGGPCY`, which is the type of aggregation standard and policy to use. As of the time of writing, the only standard supported is `*ETHCHL`. The policy you choose is up to you, but IBM recommends `*SRCDESTPRT`, which uses

the source and destination port of the TCP/IP traffic to determine which physical Ethernet port to transmit on—essentially using both ends of the conversation to determine which link to use.

Here is an example of the command to create an aggregate line description called *ETHERLIN01* using these parameters with CMN01 and CMN02 in the aggregate resource list:

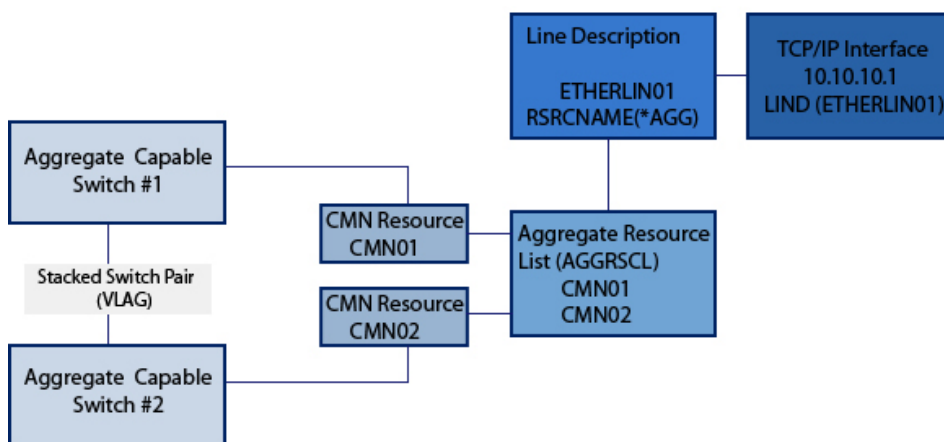
```
CRTLINETH LIND(ETHERLIN01) RSRCTYPE(*AGG)
AGGPCY(*ETHCHL *SRCDESTP) AGGRSCL(CMN01 CMN02)
```

The last step is to configure your TCP/IP address to use the new line description. To do so, run the command `ADDTCPIFC`, like so:

```
ADDTCPIFC INTNETADR('10.10.10.1') LIND(ETHERLIN01)
SUBNETMASK('255.255.255.0')
```

And that's it: You now have an interface that is redundant and aggregated. [Figure 1](#) provides a visual representation of the necessary components for the link aggregation.

Figure 1. Drawing showing the steps necessary to create an aggregate interface



Management and testing

Now that you have created your new line description, there are a couple of new things to be aware of. First, if you run the command `WRKHDWRSC *CMN` again, you will see a new device listed with a device ID of 6B26 and a description of AGGxx: This is the logical representation of the new device that you created.

Also, if you run `DSPLIND LIND(ETHERLIN01) OPTION(*AGGRSCL)`, you will notice the CMN resources that you identified earlier in your aggregate resource list and their current status. Output should look similar to [Listing 1](#).

Listing 1. DSPLIND sample output

```

                Display Line Description
Line description . . . . . : ETHERLIN01
Option . . . . . : *AGGRSCL
Category of line . . . . . : *ELAN

-Aggregated Resource List--
Name                Status
CMN01                LINK UP
CMN02                LINK UP
    
```

Testing

You can test the aggregate feature in several different ways. Physically unplugging the Ethernet cable from one of the ports that you defined in your resource list causes that link to go down but not the entire interface. You can also use a dynamic logical partitioning (DLPAR) function to unassign one of the cards. Your network administrator could shut down one of the ports on the attached switch, as well. None of these tests should affect traffic to or from your IBM i system, but during the test, you should see one of the CMN resources in the aggregated resource list change from **LINK UP** to **LINK DOWN**.

Conclusion

You can increase your system's resiliency and performance by using the steps outlined in this tip. By doing so, you can prevent unwanted downtime to your users.

Related topics

- Visit the [IBM i information center](#) to learn more about the commands referenced in this tip.
- The [IEEE 802.3 Ethernet Working Group](#) provides overwhelming information on Ethernet standards that apply to this tip.
- [Evaluate IBM products](#) in the way that suits you best: Download a product trial, try a product online, use a product in a cloud environment.

© Copyright IBM Corporation 2012

(www.ibm.com/legal/copytrade.shtml)

[Trademarks](#)

(www.ibm.com/developerworks/ibm/trademarks/)