LTM: Interface Failsafe

Deb Allen, 2008-04-03

Nathan McMahon and Kirk Bauer bring us another slick solution for a couple of common requirements in high performance environments: To fail over immediately when a single physical network connection loses connectivity, or if the number of active members in a trunk falls below the configured minimum.

For example, many customers want to fail over immediately when the egress router-facing interface loses connectivity. Others have minimum throughput requirements on multi-port trunks and need to monitor the trunk & fail over when more than 1 port in the trunk fails.

Existing failsafe mechanisms don't quite address these requirements. Interface Failsafe is the answer.

Failsafe Definitions

As of BIG-IP LTM version 9.0 – 9.4.x there are two primary mechanisms built into LTM to determine if it has lost connectivity with the network. Once determined, the LTM has the option of initiating a reboot or simply failing over to its peer in the redundant pair. The two built in mechanisms are Gateway Failsafe and VLAN Failsafe.

Both are excellent solutions, however they don't necessarily cover all scenarios. For instance, Gateway Failsafe requires that it actively ping or 'touch' an IP on the network. This can be difficult as there often isn't a stable IP address on an internal server network behind the LTM. The typical scenario is to have the LTM query the default gateway, but on the server network the LTM is the default gateway of the servers. It can't reliably query the IP address of its peer because if the standby LTM loses network connection, then the active LTM would believe it had lost connection as well and would initiate a failover.

VLAN Failsafe is a very robust solution relying on a combination of passive listening for traffic and generating or soliciting traffic if it hasn't received any data for a specified period of time. While robust, this also tends to translate into relatively long intervals between the time the LTM became disconnected from the network and when it signals a failover. Due to potential spanning tree convergence times, this timeout can be as high as 90 seconds, or as optimal as 10 seconds if Port Fast is supported by the switches.

What Is Interface Failsafe?

Interface Failsafe is a custom health monitor coupled with a method of configuring the LTM in order to allow the system to fail over based on the health of the physical link status. Interface Failsafe is meant to solve a very specific problem: When an interface becomes unplugged or otherwise loses connectivity to the network, the LTM needs to signal a failover immediately rather than try to wait for traffic to occur on a disconnected interface. It also has the ability to support M-of-N interfaces within a trunked VLAN.

Interface Failsafe provides:

- A lightweight mechanism to respond to a failure at Layer 1
- Very quick response to a physical change. In a system that is not overwhelmed by health checking servers or other management tasks, it is possible to provide significantly faster failover times than previously available with VLAN Failsafe
- A mechanism to determine link status without requiring an active interaction with other devices at Layer 2 or Layer 3 (such as ARP or Ping of an IP address).
- Ability to support M-of-N redundancy: If 2 out of 4 ports in a VLAN trunk are down, then signal a failover. See *M-of-N Redundancy* section for more details.
- A complementary addition to the capabilities of VLAN Failsafe. Interface Failsafe can run concurrently with VLAN Failsafe to provide the Layer 1 and Layer 2/3 checking.

failure to provide the Layer 4 and Layer 2/3 checking.

Interface Failsafe *does not* detect a port that has been disabled or blocked in software on the switch. It only checks physical link layer discontinuity such as the cable being unplugged.

M-of-N Redundancy

A common request is to health check a trunked VLAN. If the trunked VLAN has 4 interfaces then it might not be appropriate to signal a failover if only one of the interfaces is down. However, if two of the four are down, then you may want to initiate the failover. The Gateway and VLAN Failsafe options do not support this capability, as they would require all four interfaces to fail before signaling a failover. It is possible to monitor the trunk state and fail over via iControl, but a solution which doesn't depend on an external system to monitor the trunk state is preferable. With Interface Failsafe you can incorporate this ability.

Enabling Interface Failsafe

1. Copy the appropriate script from the codeshare to */usr/bin/monitors/interface_failsafe.eav*

2. Type: *chmod 700 /usr/bin/monitors/interface_failsafe.eav*

3. From the web UI create a new monitor called *interface_failsafe*

   - Type: *External*
   - Interval/Timeout: Recommended starting values are 5 second interval, 11 second timeout.
     This monitor is fairly lightweight, so it can be run at relatively aggressive intervals compared to typical server health checks. Keep in mind that a heavily loaded LTM may still require backing off of the interval. Timeout must be greater than the interval.
   - External Program: */usr/bin/monitors/interface_failsafe.eav*
   - Arguments:
     - For interface failsafe only (no trunk support), enter a list of the interfaces for which Interface Failsafe should be enabled. Use spaces to separate multiple entries: *1.1 1.7 2.1*

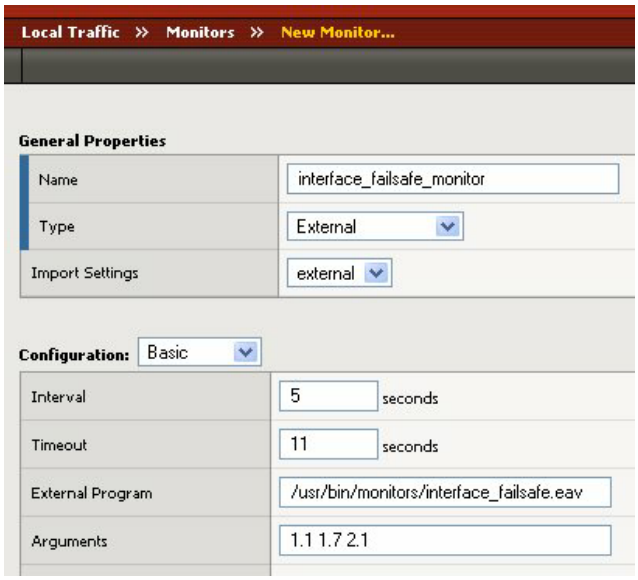

     - For interface failsafe w/Trunk Minumum Active Members support, you can also include in the list of interfaces value sindicating the name of a trunk and the minimum number of interfaces which should be up to avoid failover. For example, the following arguments will make sure that the 1.1 interface is up and test_trunk has at least two interfaces in it that are up: *1.1 test_trunk=2*

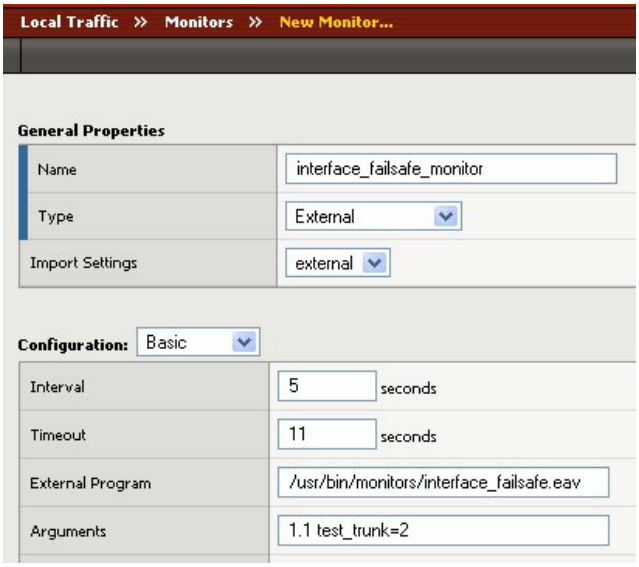

4. Create a new pool
    - Name: *Interface_Failsafe_Pool_1*
    - Health Monitor: Select the *interface_failsafe_monitor*
    - Address / Port: Pick any IP / Port. This data is discarded later and never used. This does not need to be a real server or device.
5. Create a second new pool
    - Name: *Interface_Failsafe_Pool_2*
    - Health Monitor: Select the *interface_failsafe_monitor*
    - Address / Port: Pick any IP / Port.

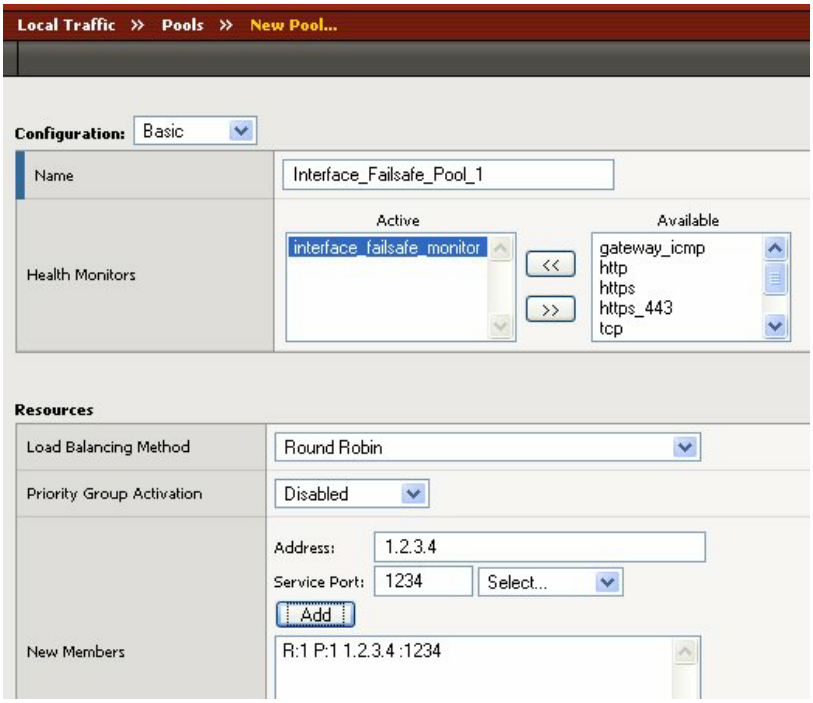

6. Create the gateway failsafe: System > High Availability > Fail-safe > Gateway
    - Gateway Pool: *Interface_Failsafe_Pool_1*
    - Unit ID: *1*
    - Threshold: *1* (number of members which must be available in the pool to avoid failover)
    - Action: Best practice is to *Fail Over*, not reboot.
7. Create a second gateway failsafe
    - Gateway Pool: *Interface_Failsafe_Pool_2*

- Unit ID: *2*
- Threshold: *1*
- Action: *Fail Over*

8. Synchronize the configuration to the LTM's peer and test.

Best Practices

- As with any failover mechanism, it is prudent to regularly test and ensure that it is configured correctly and will behave as you expect. While the Interface Failsafe can provide a greater level of availability insurance, it is important to understand exactly what it is and isn't providing.
- Setting both units to reboot is generally never recommended. If both LTMs are plugged into the same switch and the switch dies, then both LTMs will continuously reboot until an administrator manually intervenes. Consider setting only one LTM to reboot or preferably just setting both to the Fail Over option.
- Interface Failsafe can also be used in conjunction with VLAN failsafe to support both Layer 3 (network traffic connectivity) and Layer 1 connectivity. This provides the robustness of the VLAN Failsafe with the fast detection of link down status by the Interface Failsafe.

Links

Monitor Script Source

Get the Flash Player to see this player.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com