

Cisco Local Director Abstract

Stephen Gill
E-mail: gillsr@cymru.com
Revision: 1.0, 04/18/2001

Contents

Introduction	2
Dispatch v. Directed.....	2
Network Configuration Options	3
Switched Environment Implementation.....	4
Stateful Failover.....	4
Troubleshooting Aids	5
Undocumented Commands	6
Increasing Performance.....	7
Conclusion	7

Introduction

The capabilities of the Cisco Local Director are typically not disseminated in their entirety. This brief treatise expounds upon the lesser-known options for network implementation, and reviews some helpful troubleshooting aids. In most cases it assumes version 3.1 and above of the LocalDirector microcode.

The LocalDirector is a smart layer two bridge that provides colorful server load balancing capabilities. In line with how a standard bridge operates, it will bridge traffic between interfaces only if the source and destination addresses reside on separate interfaces. All traffic that is being load balanced should pass through the LocalDirector in both directions to ensure proper load distribution. The LD must see the entire dataflow in order to properly monitor connections and to compute the necessary packet modifications. Unlike an eNetwork Dispatcher that operates as a router, a LocalDirector only operates as a bridge.

Dispatch v. Directed

Typically bridges do not inspect or modify packets at the IP layer. This is not true with the LocalDirector. The bridging function of the LocalDirector supports two major modes of operation: dispatch and directed. The “redirection” command is used to enter either of these two modes. Directed mode is the default method that uses NAT (Network Address Translation) to modify IP headers. When in directed mode, the LD will inspect and modify the IP headers in packets. The destination IP of incoming connections will be translate from the virtual to the real server IP addresses; the source of outgoing replies will be converted from the real to the virtual IP address. NAT provides the quickest setup with no network address changes, and keeps setup and administration time at a minimum.

One downfall of using NAT is that it requires a bit more overhead to maintain, thus adding an extra load to the LocalDirectors. It also cannot always guarantee that payloads with embedded IP addresses will be

translated properly. Thirdly, it is incompatible with the ASLB feature of Cisco 6000 series catalyst switches.

Better performance and higher compatibility can be achieved by using dispatch mode. At the expense of additional setup requirements, dispatch mode allows for higher traffic throughput. The virtual IP address is aliased on the loopback interface of each real server, eliminating the need for address translation. For inbound connections the destination MAC address is replaced with the selected real server's MAC address, and the packet is forwarded retaining the IP address. In outgoing replies, real servers utilize the virtual IP alias on the loopback as the source address making it appear as if the response came from the LocalDirector. MAC translation is more efficient than IP translation because only inbound flows require modification. Individual site requirements should dictate which mode is chosen as they pose unique benefits and drawbacks.

Network Configuration Options

Aside from the two typical two-port configurations, several implementation possibilities abound. Among others, a few configurations that have been individually lab tested are described below.

Private VLANs were deployed on a pair of Catalyst 6509 switches to confirm the ability to load balance across multiple server clusters with only two ports. The internal interface of the LocalDirector was placed in a promiscuous port of a private VLAN. Community VLANs successfully provided website isolation whilst individual site clusters hung off of the same physical LocalDirector interface. In a lab environment both NAT and MAC address translation methods functioned as expected.

For architectures requiring visibility between server clusters, one might deploy a three port, three VLAN LocalDirector configuration. In this scenario load balancing across server clusters is still maintained while internal websites are allowed to see each other. Once again, a pair of Catalyst 6509 switches were deployed to provide VLAN separation between all three LocalDirector interfaces. The LD was configured with two internal server clusters, each hanging off of their own individual LD interface and separate VLAN. This configuration provides the unique ability to load balance requests coming from the external interface as well as any of the internal interfaces. All traffic destined to a virtual interface on a LocalDirector must pass through it in both directions, thus requiring that separate VLANs be deployed per internal LD interface. Both NAT and MAC address translation methods were successfully tested for this design.

A common misconception is that the LocalDirector must be on the same network as the servers being load balanced. The LD does in fact support private addressing for internal servers by utilizing the "alias ip address"

command. Assigning an alias to the LocalDirector has the effect of placing it in that network. This allows it to communicate with servers on a different network than it's own without the use of a router. This feature facilitates conservation of public address space, as only one public IP is necessary for the entire site. Real servers are not limited to only receiving incoming connections. The "static" command may be used to translate a real server address to a virtual server. This translation enables outbound connections to be made from a real server while hiding its internal IP address.

Switched Environment Implementation

When connecting a set of LocalDirectors to a Cisco switch, special care must be taken to guarantee proper site operation. A special feature of the catalyst switch named "portfast" is generally enabled on ports that belong end stations. This has the effect of disabling spanning tree and allowing a port belonging to a host to enter the forwarding state more rapidly. However, "portfast" can pose a problem to a LocalDirector when enabled on its switch port. Traffic that is bridged from one LocalDirector port to another can confuse a switch and cause it to think there is a spanning tree loop. The switch will then attempt to place the port in "errdisable" mode and not allow it to communicate. In reality there is no spanning tree loop, only the appearance of one since the interfaces are in separate VLANs. The best remedy is to disable portfast for all LocalDirector switch ports. Also, disabling CDP (Cisco Discovery Protocol) on those ports is recommended so that the switch will not be confused and spit out console error messages regarding a "Native Vlan Mismatch."

There exists an option on the LocalDirector to disable bridging per interface. The "secure" command can be used to block bridged traffic bound for a specific interface in LocalDirector without affecting traffic that is load-balanced through a virtual server. Only traffic being serviced by a virtual server traverses the interface, and no traffic is bridged to or from the interface. Enabling this option would take care of removing the appearance of a spanning tree loop, but at the cost of severely limiting traffic that will pass through the LocalDirector.

Stateful Failover

The ability to maintain connection state on a per-virtual server basis can be enabled using the "replicate" command. A dedicated interface must be reserved for this purpose in order to replicate established connections with the standby unit. However, state is not maintained for proxied connections such as ftp-proxy and SSL sticky. By default, replication is disabled such that all active connections are dropped during a failover. Clients must then re-establish their connections through the newly active unit.

Though possible, it is not recommended that the LocalDirector maintain state for short-lived connections such as HTTP. The increased overhead usually does not justify the minimal stability gained by enabling replication for transitory flows. Since the command is enabled on a per-virtual server basis, it is possible to turn it on for more crucial protocols while leaving it off for short-lived connections. Individual site requirements should dictate whether this feature should be enabled.

Troubleshooting Aids

A few documented commands that are essential to every LD troubleshooter's toolkit are listed below. The list is certainly not exhaustive, though the commands might prove valuable in debugging some of the thornier issues. Note that some of these may or may not work depending on the hardware platform and microcode revision in operation.

show blocks – A block (buffer) is the resource used to store packets from the network. This command will give you a good idea of memory usage for each range of packet sizes. In the event of memory resource exhaustion, the low column will display a 0. This means that LocalDirector ran out of that size block at some time since reboot. Use the number of *No Buffer* packets from the “show interface” command output to view the number of packets dropped. Also note in the output that the same source MAC address is used for every interface of the LocalDirector (since microcode 3.x).

show error – Be prepared for an extremely verbose, though useful output of all error counters of the LocalDirector. It can be useful for pinpointing particular problems with a site such as checksum errors, port errors, failover indicators, servers not available, resource exhaustion, fragmentation issues, and many more.

show syn – This may give you a different perspective on the number of connections per virtual server as well as TCP SYNS that remain outstanding.

show statistics – For a more general overview, this command will display the number of bytes, packets, and connections sent to virtual and real machines.

show processes – Information about running threads can be obtained using this command. This may show what processes are taking up most of the CPU cycles.

syslog <host ip | console | output facility.level> – It may be desirable to configure a syslog server and increase the facility-level (up to debug)

to obtain more detailed logs when problems are lurking. The “syslog console” command is always only valid for the current session and does work over telnet. Logs might prove useful in determining bridging loops and capturing informational errors.

<show | clear> bridge – The bridge command can be used to view or clear the bridging table. Bridging information is sometimes more useful than viewing an ARP cache because entries last longer and rx/tx statistics are listed per MAC address. As traffic is received, a bridge table is populated for each interface, showing the MAC addresses that are accessible through that interface. It may be necessary to clear the bridge table if any physical changes occur such as a server move from one interface to another. This will ensure that traffic is sent over the correct interface.

Undocumented Commands

By searching through the binaries, I’ve uncovered several hidden commands that may or may not aid in problem determination on a LocalDirector. Keep in mind that these commands have not been fully tested and are likely not documented for a reason. Do not use these commands on a production site unless you are certain of what they will produce.

To enable the following commands, you must first enter utilities mode by typing “utils”. When prompted for a password, enter “xyzy”.

<show | set> profile [start | stop | clear] – Control system profiling.

show ipfrag – Show IP fragmentation information.

show vmachlu [hash]– Show virtual machine hash index.

<show | set> watchdog – Show/Set watchdog limit in milliseconds.

show <sight | whereb> – Print last known location of blocks (similar to show block)

show state – Show state of connection objects. Use of this command is not recommended.

show chinfo – Show channel information.

gdb break – Manipulate remote debugger.

[no] debug block sub-block – Enable debugging functions for options listed below.

show debug – List available debugging functions.

The following output has been modified to improve readability.

```
LocalDirector# show debug
ci config      rip open      rip close     rip get
snmp pdu       snmp rslv     snmp if       snmp ioctl
snmp trap      snmp ae       snmp ld flash gen
flash intel    flash atmel   fover open    fover tx
fover rx       fover txdmp   fover rxdmp   fover cable
fover switch   fover get     fover put     fover verify
fover rxip     fover txip    fover ifc     fover fail
fover fliprx   fover fliptx  fover flifcrx fover flifctx
icmp open      icmp close    icmp get       icmp in
icmp out       ip config     ip open ip close ip put
ip get         ip ioctl     ip arpin      ip arpreq
ip in         ip route     ip pkt        ip rif
ip frag       bridge get    bridge put     bridge pkt
ip5511 rcv
LocalDirector#
```

Increasing Performance

Accelerated Server Load Balancing (ASLB) is a new feature of Cisco Catalyst 6000 series switches that works in conjunction with the LocalDirector. When used properly, ASLB can greatly increase the performance of a site by bypassing the LocalDirector for certain flows. The concept is very similar to layer three switching where the switch bypasses the MSFC (router) for flows that it knows how to handle. It should be noted that ASLB is only compatible with dispatch mode of the LocalDirector. Discussions regarding the implementation and configuration of ASLB shall be left to another paper.

Conclusion

The Cisco LocalDirector is a load balancing bridge with a wide range of capabilities. Its basic modes of operation include the ability to perform NAT or MAC translation of dataflows. A few configuration options have been explored with suggestions on how to implement LDs in a switched environment. Finally, several troubleshooting resources were highlighted that can greatly assist in problem resolution.