



**GUIDELINES FOR USING
DEVICE LEVEL RING
(DLR) WITH
ETHERNET/IP™**

Guidelines for Using Device Level Ring (DLR) with EtherNet/IP

Contents

1.	Introduction	3
2.	Overview of DLR	3
2.1.	Ring Supervisor	3
2.2.	Beacon-based Ring Node	3
2.3.	Announce-based Ring Node	4
2.4.	DLR Operation	4
2.5.	Devices which do not support the DLR protocol.....	4
2.5.1.	Non-DLR Device.....	4
2.5.2.	Non-compliant Device.....	4
3.	General Considerations.....	5
3.1.	Simple Stand-Alone DLR Ring	5
3.2.	DLR Attributes, status indicators and diagnostics	6
3.3.	DLR, High-Performance applications and IEEE-1588.....	8
3.4.	DLR and QuickConnect.	9
3.5.	Inserting Non-DLR Devices in a DLR Ring.....	9
3.5.1.	Requirements for Devices not supporting the DLR Protocol	13
3.5.2.	General Configuration Requirements for devices not supporting the DLR protocol.....	14
3.5.3.	Configuration Requirements for a Ring Network	14
3.6.	DLR and Resiliency Protocols.....	15
3.6.1.	A non-DLR loop passing through a DLR ring.....	16
3.6.2.	A DLR Redundant Gateways	17

Guidelines for Using Device Level Ring (DLR) with EtherNet/IP

1. Introduction

Ethernet is growing in popularity across all levels of enterprise including industrial processes and applications. There are a variety of Industrial applications in which Ethernet ring topologies are preferable to the star topologies common in enterprise networks. Ring networks provide inherent single-point fault tolerance. Ring nodes that include embedded switch technology reduce the need for infrastructure switches and simplify network cabling. Device Level Ring (DLR) protocol provides a means for detecting, managing and recovering from faults in a ring-based network.

Implementation of DLR imposes certain requirements upon the supporting network infrastructure. DLR does not inherently exclude the use of devices which do not support the DLR protocol in a DLR-enabled network. It is expected that legacy devices and other considerations will frequently dictate the use of such devices in a DLR network. However, the use of these devices in a DLR network may significantly affect DLR operation and performance. This whitepaper is intended to provide an overview of DLR and to provide guidelines for implementing a DLR network comprised of DLR devices and devices which do not support the DLR protocol.

2. Overview of DLR

As previously noted, the DLR protocol is intended to support simple ring topologies. DLR supports three classes of devices, the Ring Supervisor, the Beacon-based ring node and the Announce-based ring node.

2.1. Ring Supervisor

A DLR network requires that at least one device be configured to act as the Ring Supervisor. The Ring Supervisor is responsible for verifying the integrity of the ring; reconfiguring the ring to recover from faults and collecting diagnostic information for the ring. The active Ring Supervisor blocks traffic on one of its ports (with the exception of a few special frames) and does not forward traffic from one port to another, thus avoiding a network loop. It is strongly recommended that at least one additional device be capable of acting as a back-up ring-supervisor. Each supervisor is configured with a precedence value. The device with the highest precedence value will become the active Ring Supervisor. In the event that two supervisors on the ring have the same precedence value, the device with the numerically highest MAC address will become the active Ring Supervisor.

2.2. Beacon-based Ring Node

Generally speaking, a ring node is any non-supervisor device that operates on the ring and participates in the DLR protocol. Ring nodes participate in fault detection (Check Neighbor Process, Fault Detection Process). When a fault is detected, a ring node will reconfigure appropriately and relearn the network topology (Ring Recovery Process). These beacon-based ring nodes are required to process beacon frames within a specified beacon interval. The default beacon interval is 400 us. The minimum beacon rate is 100 us. The default beacon interval allows for ring recovery times on the order of 3 ms for a 50 node ring. Faster recovery times are possible with smaller beacon intervals.

2.3. Announce-based Ring Node

An announced-based ring node differs from a Beacon-based ring node only in its ability to process beacon frames. These devices are not required to process DLR beacon frames, but must be capable of processing announce frames. Announce frames are also generated by the Ring Supervisor. The default interval for announce frames is one second or immediately upon change of ring state. Ring recovery times for an announce-based node are on the order of 4 ms for a 50 node ring as opposed to 3ms for a Beacon-based Ring.

2.4. DLR Operation

A DLR network consists of an active Ring Supervisor and any number of Ring Nodes. Ring nodes incorporate embedded switch technology with at least two external ports. The Ring Supervisor is responsible for generating a “beacon” at regular intervals. These beacons traverse the ring in each direction. The Ring Supervisor also sends announce frames on both ports once per second. Announce frames allow ring nodes that are unable to process the high-speed beacon frames to participate in fault detection and ring recovery.

The Ring Supervisor must be capable of blocking DLR and other network traffic to avoid infinite propagation of these frames through the ring (Network storm). Faults are detected when beacon traffic is interrupted and link/node failure is detected by adjacent nodes. The DLR protocol contains a number of fault detection and ring recovery mechanisms.

2.5. Devices which do not support the DLR protocol.

Although not recommended, it is possible to insert device which do not support the DLR protocol. For the purposes of discussion, two such device types are defined: the Non-DLR device and the Non-compliant device.

2.5.1. Non-DLR Device

A Non-DLR device is any device not supporting the DLR protocol but which complies with all of the guidelines in sections 3.5.1 and 3.5.2. These devices are not recommended because they make finding the location of ring faults more difficult.

2.5.2. Non-compliant Device

A Non-compliant device is a Non-DLR device that fails to comply with one or more of the guidelines in sections 3.5.1 and 3.5.2. Non-compliant devices are prohibited in a DLR ring as they can have unpredictable and adverse effects on fault detection and ring restoration.

3. General Considerations

DLR is generally intended for a simple, single-ring topology requiring fast recovery from network failures. The protocol does not support the concept of multiple or overlapping rings. While a DLR ring can contain an arbitrary number of nodes, recommended ring size is less than 50 nodes. As the number of nodes grows, the time required for DLR frames to traverse the ring increases, leading to increased fault detection and recovery time (see Table 1). In addition, in a larger ring the probability of a fault increases, including double faults in which a segment may be lost from the rest of the network. Ultimately the number of nodes a user selects for a given ring depends upon the performance requirements the user has set for that ring.

Number of Ring Nodes	Beacon Interval (usecs)	Round Trip Time ¹ (usecs)	Beacon Timeout (usecs)	Physical Layer Faults Recovery Delay ¹ (usecs)	Non-physical Layer Faults Recovery Delay for Beacon Frame Based Nodes (usecs)	Non-physical Layer Faults Recovery Delay for Announce Frame Based Nodes (usecs)	Ring Restore Delay for Beacon frame Based Nodes (usecs)	Ring Restore Delay for Announce frame Based Nodes (usecs)
25	400	905	1380	980	1858	2335	1808	2260
50 (nominal network size)	400	1810	1960	1885	2890	3820	3165	4070
100	400	3620	3120	3695	4955	6790	5880	7690
150	400	5430	4280	5505	7020	9760	8595	11310
200	400	7240	5440	7315	9085	12730	11310	14930
250	400	9050	6600	9125	11150	15700	14025	18550

¹ Same for Beacon and Announce frames based nodes.

Table 1 - Example Ring Configuration Parameters and Performance

3.1. Simple Stand-Alone DLR Ring

Figure 1 shows a conceptual diagram of a simple stand-alone DLR ring. The application includes various I/O modules, control in the form of a PAC, an HMI to monitor device activity and an administrator station for network configuration. Note that each ring node supports two EtherNet/IP ports. It is expected therefore, that each ring node or supervisor will incorporate embedded Ethernet switch technology. For the purposes of this example, assume that no safety or performance critical constraints are placed upon the network. Further assume that, for this application, desired fault detection and recovery times are on the order of tens of milliseconds. Therefore, the presence of announced-based ring nodes in the ring would be acceptable from a performance and recovery viewpoint.




Figure 1- Simple Stand-Alone DLR Ring

Figure one also introduces a component for consideration when implementing a DLR ring: the DLR Tap. A DLR tap is simply a 3 port device which allows network infrastructure switches, Non-DLR devices, or devices without embedded switch technology to connect to a DLR ring. The DLR tap is a ring node which may or may not support the capability of acting as a Ring Supervisor.

3.2. DLR Attributes, status indicators and diagnostics

The DLR object attributes can provide helpful information when a fault occurs. It should be noted that not all attributes are available at all nodes in the ring. It should also be noted that while attribute 1, Network Topology, refers to both linear and ring topologies, this does not imply that DLR supports a line topology. Attribute 1 is a status bit indicating that the supervisor has either blocked traffic in one direction on the ring and the ring is in its RING_NORMAL_STATE (linear topology) or it has not blocked traffic in one direction because the ring is either in startup or a faulted state (ring topology). Many of the attributes, such as the participants list, depend upon a beacon frame traversing the network to gather the attribute's data. Some of the services provided by the DLR object are also useful in ring diagnostics. In particular, the "Verify Fault Location" service will update the attributes 6 and 7 so the user can determine the location of a fault.

Attribute ID	Attribute Name	Description	Available from Ring Supervisor	Available from Ring Node
1	Network Topology	0 indicates "Linear" 1 indicates "Ring"	Yes	Yes
2	Network Status	0 indicates "Normal" 1 indicates "Ring Fault" 2 indicates "Unexpected Loop Detected" 3 indicates "Partial Network Fault" 4 indicates "Rapid Fault/Restore Cycle"	Yes	Yes
5	Ring Fault Count	Number of ring faults since power up	Yes	No
6	Last Active node on port 1	Last active node at the end of chain through port 1 of active ring supervisor during ring fault	Yes	No
7	Last Active node on port 2	Last active node at the end of chain through port 2 of active ring supervisor during ring fault	Yes	No
8	Ring Participants Count	Number of devices in ring protocol participants list	Yes	No
9	Ring Protocol Participants List	List of devices participating in ring protocol	Yes	No
10	Active Supervisor Address	IP and/or MAC address of the active ring supervisor	Yes	Yes
11	Active Supervisor Precedence	Precedence value of the active ring supervisor	Yes	No
12	Capability Flags	Describes the DLR capabilities of the device	Yes	Yes
13	Redundant Gateway Config	Redundant Gateway configuration parameters	Yes	No
14	Redundant Gateway Status	0 – indicates the device is functioning as a non-gateway DLR node (gateway not enabled) 1 – indicates the device is functioning as a backup gateway 2 - indicates the device is functioning as the active gateway 3 – indicates gateway fault state due to loss of communication on uplink port	Yes	No

Attribute ID	Attribute Name	Description	Available from Ring Supervisor	Available from Ring Node
		4 – indicates the device cannot support the currently operating gateway parameters (Advertise Interval and/or Advertise Timeout) 5 – indicates gateway fault state due to partial network fault		
15	Active Gateway Address	IP and/or MAC address of the active gateway device	Yes	No
16	Active Gateway Precedence	Precedence value of the active gateway	Yes	No

Table 2 – DLR Attributes

3.3. DLR, High-Performance applications and IEEE-1588

Figure 2 depicts two different ring types. On the left, is the same simple, comparatively low performance ring depicted in figure 1. On the right, figure 2 shows a high performance ring requiring the fast recovery times provided by a beacon-based DLR network. The DLR protocol does not recognize the existence of multiple or overlapping rings. However, networks can contain more than one DLR ring so long as the rings are isolated such that DLR protocol messages from one ring are not present on another ring.

While the DLR protocol does not require support for IEEE-1588, high performance applications can benefit from both the time synchronization of IEEE-1588 and the fast fault detection and ring recovery inherent in the DLR protocol. Nodes which do not support IEEE-1588 time synchronization will introduce additional messaging jitter into the ring thus affecting the accuracy of time synchronization. It is therefore recommended that all nodes participating in such a high-performance ring support IEEE-1588 and the CIP time sync object. More specifically, while the device itself need not be time-aware (i.e. it need not process IEEE-1588 messages locally), it should include support for end-to-end (E2E) transparent clock.

Similarly, inserting Non-DLR nodes into a high performance ring will increase the time required for fault detection and ring recovery. For high performance applications such as CIP motion, it is preferable to connect a Non-DLR device through a DLR-compliant Tap as shown in figure 2. Note that figure 2 depicts both a switch which is compliant to the DLR protocol (S3) and a switch which does not (S2). Switch S2 is connected to the ring via a DLR tap. An infrastructure switch that does not specifically support DLR is prohibited from being directly connected to a DLR ring.




Figure 2 - DLR Network with two non-overlapping rings

3.4. DLR and QuickConnect.

QuickConnect is an ODVA technology intended to address specific applications in Automotive Manufacturing. Applications such as robots, tool changers and framers are required to quickly exchange tooling fixtures which contain a section or segment of an industrial network. This requires the network and nodes to be capable of quickly connecting and disconnecting, both mechanically, and logically. Consequently, when in QuickConnect mode, a port configured for forced speed and duplex mode, QuickConnect devices do not use Auto-MDIX (detection of the required cable connection type). Because Auto-MDIX is disabled for QuickConnect applications, extra care must be taken to ensure cabling connections are correct.

3.5. Inserting Non-DLR Devices in a DLR Ring

As shown in figure 3, the user may insert a Non-DLR device into a DLR ring. As Non-DLR devices are inserted into the ring, they have no adverse impact on worst case fault detection time (Note: a Non-DLR device is required to meet all the guidelines in sections 3.5.1 and 3.5.2; therefore MAC learning must be disabled). The supervisor will still detect a lack of beacon frames and initiate ring recovery. However, if a significant number of Non-DLR devices are included in the ring, isolating and diagnosing a fault becomes problematic if the fault occurs between two non-DLR nodes.




Figure 3 - DLR ring with a Non-DLR Device

For example, consider the topology in Figure 4. This simple ring consists of a supervisor (S1), four DLR nodes (D1 through D4) and a single Non-DLR device (N1). From a DLR perspective, N1 does not exist. It is simply part of the physical connection between D2 and D4. D4 considers D2 to be its DLR “neighbor and vice versa. Assume that a link failure occurs between N1 and D4 (as indicated by the red “X”). D4 will detect the failure and immediately inform the supervisor. The supervisor will proceed with the fault detection process and identify D2 as the “Last Active node on port 1” and D4 as the “Last Active node on port 2”. Further, the Network Administrator can query link status for D2 and D4 via the Ethernet Link Object. Therefore, the Network Administrator is aware that, from D4’s perspective, the link with D2 is down, while from D2’s perspective the link is still up. The supervisor will unblock traffic on its port 2, network traffic will resume and the Network Administrator can reasonably deduce that the fault lies either with D4’s lower port, N1’s right port or the cabling between these nodes.




Figure 4 - Simple ring with one Non-DLR node

Now consider the case depicted in Figure 5. In this case, Non-DLR devices N2 and N3 have been added to our simple ring. As before, D4 considers D2 to be its DLR “neighbor and vice versa. Assume a link failure between N2 and N3. In this case, because the link in question is not directly tied to a DLR-capable device, the link failure cannot be reported via the Ethernet Link Object. A beacon timeout must occur before the Supervisor becomes aware of the problem. Again, the supervisor will proceed with the fault detection process and identify D2 as the “Last Active node on port 1” and D4 as the “Last Active node on port 2”. The supervisor will unblock traffic on its port 2 and network traffic will resume. However, the Network Administrator will have more difficulty isolating the fault since it might exist anywhere between D2 and D4.




Figure 5 - Simple ring with three Non-DLR nodes

If multiple Non-DLR devices are to be inserted in the ring, it is preferable to insert DLR nodes between these devices for better isolation of faults (Figure 6).




Figure 6 - Simple ring with three Non-DLR nodes interspaced

3.5.1. Requirements for Devices not supporting the DLR Protocol

To prevent interference with the DLR protocol, a device which contains an embedded switch but does not support the DLR protocol must be carefully configured. In general, unicast MAC learning must be disabled on those ports connected to the DLR ring. Using Non-DLR devices can result in loss of unicast frames for some period of time following a ring fault or restoration. After a ring fault/restoration, the device's MAC learning tables may be invalid as devices are now reachable through different ports. Until the MAC learning tables are updated as a result of devices sending frames, unicast frames may not reach the target devices. These

configuration requirements mandate the use of devices supporting the following features in the DLR ring. The use of devices lacking these features in a DLR ring will interfere with the DLR protocol. These devices are considered non-compliant and therefore prohibited in the ring.

Required Features

- Disable unicast MAC learning on those ports connected to the DLR ring
- IEEE 802.3 operation:
 - Auto-negotiation, with 10/100Mbps, full/half duplex
 - Forced setting of speed/duplex
 - Auto MDIX (medium dependent interface crossover), in both auto-negotiate and forced speed/duplex modes. Note: This is a PHY and transformer issue, not an embedded switch issue.
- QoS:
 - 2 queues
 - High priority queue for DLR frames, with strict priority scheduling for the high priority queue
 - Prioritization via 802.1Q/D. Usage shall be consistent with the EtherNet/IP QoS scheme shown in Table 3. For non-IP frames the priority in the 802.1Q header should be used.

Recommended Features

- QoS:
 - 4 queues
 - Prioritization via DSCP Usage shall be consistent with the EtherNet/IP QoS scheme shown in Table 3. For IP frames the switch should use the DSCP value.

To ensure proper operation of devices not supporting the DLR protocol (as described in section 3.4.1) in a DLR ring, the user should configure the device as outlined in the following sections. Please note that the configuration steps required to support those networks using VLAN routing (VLAN-based ring network) and those network which do not use VLAN routing (non-VLAN based ring network) differ as noted below.

3.5.2. General Configuration Requirements for devices not supporting the DLR protocol

The use of devices not supporting the DLR protocol in a DLR ring may result in an undesired loss of unicast frames following fault detection and ring recovery. The user may avoid this situation by disabling unicast MAC learning on those ports connected to the DLR ring.

3.5.3. Configuration Requirements for a Ring Network

1. Quality of Service (QoS) based upon either the Differentiated Services Code Point (DSCP) contain in the Type of Service (TOS) field of the IPv4 header or upon the 3-bit priority field contained in the VLAN ID as specified in IEEE 802.1D/Q. Use of the DSCP for QoS is preferred. Default mapping of EtherNet/IP traffic to the DSCP and IEEE 802.1 D is shown in Table 3. While two priority queues are acceptable per the specification and likely acceptable for low performance applications, the use of four priority queues will ensure more deterministic delivery of high priority traffic. The user should carefully consider fault detection and ring recovery time requirements when implementing the QoS scheme for their application.
2. The switch ports connected directly to the ring must be configured to preserve IEEE 802.1Q tag priority of ring protocol frames when they pass through the ports.
3. Disable IP multicast filtering on the two ports of the switch connected to ring. This step will assure uninterrupted delivery of EtherNet/IP multicast connection data after a ring reconfiguration.

- Statically configure the three multicast addresses used by ring protocol to be forwarded only on two ports of switch connected to ring. This step must be done to prevent multicast ring protocol frames from being forwarded on other ports of switch. These addresses are:

Message Type	MAC Address
Beacon	01:21:6C:00:00:01
Neighbor Check Request/Response/Sign on	01:21:6C:00:00:02
Announce/Locate_Fault messages	01:21:6C:00:00:03

- Configure unicast MAC addresses of all configured ring supervisors statically into the MAC table of switch such that unicast traffic destined for ring supervisors will be forwarded through both ports of switch connected to ring. This step must be done to prevent switch from getting confused by bi-directional ring beacons from active ring supervisor.

Traffic Type	CIP Priority	DSCP	802.1D Priority ¹	CIP Traffic Usage (recommended)
PTP Event (IEEE 1588)	n/a	59 ('111011')	7	n/a
PTP General (IEEE 1588)	n/a	47 ('101111')	5	n/a
CIP class 0 / 1	Urgent (3)	55 ('110111')	6	CIP Motion
	Scheduled (2)	47 ('101111')	5	Safety I/O I/O
	High (1)	43 ('101011')	4	I/O
	Low (0)	39 ('100111')	3	No recommendation at present
CIP UCMM CIP class 3 All other EtherNet/IP encapsulation messages	All	35 ('100011')	3	CIP messaging
¹ Sending 802.1Q tagged frames is disabled by default				

Table 3 - Default DSCP and 802.1D Mapping for EtherNet/IP

3.6. DLR and Resiliency Protocols

Switch resiliency protocols such as MSTP have the ability to block one or more ports on the switch to prevent loops. In addition to resiliency protocols, some managed switches have advanced network protection

features that can also block or disable ports on the switch. These resiliency and management protocols often assume that the switches are direct neighbors to each other and may block or disable ports isolating DLR devices from the ring. It is important to disable any resiliency or management protocols that can disable or block the ports connected to the DLR ring. The Non-DLR device must be a passive device on the ring and rely on the DLR Beacon protocol to maintain a loop free topology on the ring.

For instance, as shown in Figure 2, DLR rings may be connected to networks supporting IEEE Spanning Tree protocols (RSTP, MSTP). Care must be taken to ensure that these protocols do not interfere with DLR and vice versa. Spanning tree protocols use special control frames called Bridge Protocol Data Units (BPDUs) to exchange information about network connections and their cost. The active Ring Supervisor will not forward multicast messages with the address 01:80:C2:00:00:00 (BPDU frames) from one ring port to the other, regardless of ring state. This feature ensures that the spanning tree protocol does not recognize and attempt to manage the DLR ring.

3.6.1. A non-DLR loop passing through a DLR ring

While DLR rings can coexist with spanning tree and other resiliency protocols, they must be properly isolated. Consider the example in Figure 7. Note that switches S2 and S3 each connect to the DLR ring through a DLR tap. Because the Ring Supervisor will block BPDUs packets sent by each switch, the RSTP algorithm will be unable to resolve the network structure potentially resulting in unaddressed nodes, unmanaged rings and network storms. For this reason, multiple connections into a DLR ring from a given network should be avoided.




Figure 7 - DLR ring with Multiple Connections to a Larger Network - DO NOT IMPLEMENT!

3.6.2. DLR Redundant Gateways

To avoid the situation shown in Figure 7, the protocol supports multiple connections to a DLR network through redundant gateway devices. As shown in Figure 8, a gateway has two DLR ports to connect to the DLR network and one or more uplink ports to connect to the network infrastructure outside of the DLR network. The gateway implements DLR protocol on its two DLR ports. It implements either IEEE 802.1D RSTP or IEEE 802.1Q MSTP on its uplink ports. Optionally other protocols such as STP may be implemented on an uplink port.

Only one gateway is active at any given time. In the case shown in figure 8, RDG1 is the active gateway while RDG2 is the backup gateway. Note that RDG1 also acts as the ring supervisor. The DLR protocol provides mechanisms for the automatically selecting the active gateway and for automatic switchover should the active gateway become unavailable. The remaining, or backup, gateway device(s) block traffic from being forwarded between the DLR and uplink ports. While blocking traffic in backup mode, a backup gateway only forwards DLR traffic between its two DLR ports thus preventing the creation of an unmanaged ring as described in Figure 7. Therefore, it is strongly recommended that multiple connections to the network infrastructure outside of the DLR network be made via these gateway devices.




Figure 8 - DLR ring with 2 Redundant Gateways