

# Programmer's Manual for Multi-Level Error Detection (MLED) Architecture

Prateek Jain

January 23, 2025



# Contents

<b>1</b>	<b>Introduction</b>	<b>5</b>
1.1	Overview of MLED . . . . .	5
1.2	File Transfer in MLED . . . . .	6
1.2.1	LS( $L_{n1}$ ) and LD( $L_{n1}$ ) . . . . .	6
1.2.2	LS( $L_{ij}$ ) and LD( $L_{ij}$ ), $i < n$ . . . . .	6
1.2.3	Relay Process and Virtual Links . . . . .	6
1.2.4	File Transfer Example . . . . .	7
<b>2</b>	<b>Configuration File</b>	<b>9</b>
2.1	Example Main Configuration File . . . . .	12



# Chapter 1

## Introduction

### 1.1 Overview of MLED

The MLED is structured with  $n$  levels, where  $n \geq 3$  and each level  $i$  can consist of  $j$  layers. The layer  $L_{ij}$  represents the layer located at level  $i$  with identifier  $j$ , referred to as the LayerID. Each layer  $L_{ij}$  is governed by a specific policy  $P_{ij}$ , which dictates the operations of the layer over its scope. The scope of a layer is defined as the portion(s) of the network over which it operates. Portion of a network refers to the successive part of the physical network between the source and the destination node having the same network characteristics, such as noise interference, that influences the Bit Error Rate (BER) of that part. Layers at level  $i$  have smaller scope than the layers at level  $i + 1$ . For example, in Figure 1.1, each physical link between two consecutive nodes represents a part of the network with a distinct BER. The scope of layer  $L_{21}$  spans from node 1 to node 3. This scope includes two portions of the network: the portion between node 1 and node 2, and the portion between node 2 and node 3.

In general, an MLED framework is denoted as  $\text{MLED}(n, P)$ , where  $n$  denotes the maximum number of levels and  $P$  denotes the set of policies such that  $P_{ij} \in P$ . The value of  $n$ ,  $j$  and the policy to be used at each layer  $L_{ij}$  in MLED is determined by the admissible UEP ( $\gamma$ ) specified by the user.

Now, consider a file transfer example shown in Figure 1.1 where the MLED framework with four levels,  $\text{MLED}(4, P)$ , is started by an application at the source to transfer a file to the destination. Both sending and receiving applications are marked with white squares. In the diagram, each horizontal band extending from the source node to the destination node represented by a different color is a different level of the MLED framework. The levels are numbered in a sequentially decreasing order starting from the top level that is represented by  $n$ . Each level has a distinct set of MLED processes shown by the colored squares.

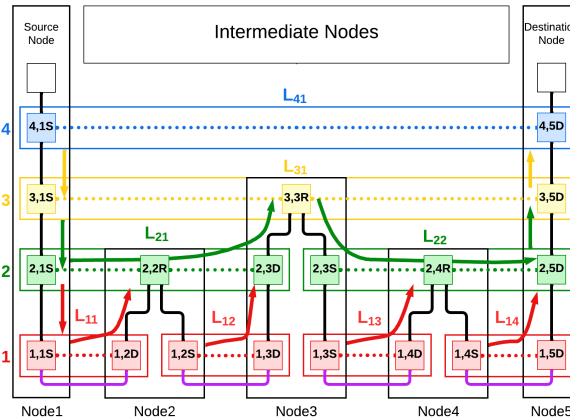


Figure 1.1:  $\text{MLED}(4, P)$  configured to run on 5 nodes.

Each process is named according to the level it belongs to, the node it operates on, and an identifier (S, D, or R) to specify whether the process functions as a Source, Destination, or Relay. For example, a process named “3,1S” indicates that it is a source process running on node 1 at level 3. These processes perform functions to realize the MLED. The vertical rectangles group the MLED processes at different

levels that run on the same node. The configuration shown in Figure 1.1 has 5 nodes. Each physical node is connected to its neighboring nodes via a physical link or an overlay link in unbalanced MLED.

Within each level, different layers are represented by colored horizontal rectangles. Each layer  $L_{ij}$  starts and ends with its own source (S) and destination (D) processes, referred to as  $\text{LayerSource}(L_{ij})$  and  $\text{LayerDestination}(L_{ij})$ , abbreviated as  $\text{LS}(L_{ij})$  and  $\text{LD}(L_{ij})$ , to differentiate them from the source and destination nodes between which the file transfer takes place. Depending on the level and the position of the layer, it may or may not contain a relay process. For instance, in Figure 1.1, the processes (3,1S), (3,3R), and (3,5D) represent the  $\text{LS}(L_{31})$ , the relay process, and the  $\text{LD}(L_{31})$ , respectively, for the layer  $L_{31}$ . Similarly, processes (4,1S) and (4,1D) are the  $\text{LS}(L_{41})$  and the  $\text{LD}(L_{41})$ , respectively, for the layer  $L_{41}$  but it has no relay process.

The layer specific policy  $P_{ij}$  defines various parameters and configurations that may include routing, addressing, congestion control, flow control or any other communication functionality. In this paper, we focus on the error detection policy.

Each MLED layer  $L_{ij}$  represents a logical connection between its  $\text{LS}(L_{ij})$  and  $\text{LD}(L_{ij})$  process. This logical connection is realised by one or more virtual links represented by horizontal dotted lines in Figure 1.1, connected by the relay process.

## 1.2 File Transfer in MLED

The file transfer process in MLED starts at level  $n$ . The fundamental operational unit of MLED is the MLED process. Various MLED processes operate to facilitate this transfer:

### 1.2.1 $\text{LS}(L_{n1})$ and $\text{LD}(L_{n1})$

At level  $n$  there is only one layer  $L_{n1}$ . The  $\text{LS}(L_{n1})$  and  $\text{LD}(L_{n1})$  enable inter-process communication (IPC) between application layer processes.

Given the likelihood of handling very large files that may not be stored locally,  $\text{LS}(L_{n1})$  reads files in chunks from a remote disk. It starts a server locally or remotely that streams the file in chunks. Based on its payload length, the  $\text{LS}(L_{n1})$  receives the file to be transferred in chunks, and forwards these chunks to the underlying layer at level  $(n-1)$ . Layer  $L_{n1}$  could implement its own error detection policy for the chunks it handles. The  $\text{LD}(L_{n1})$  collects the chunks from the lower layer at level  $(n-1)$ , reassembles them according to its payload length, and writes data to the disk or passes it to the destination application process.

If *verifyWholeFileHash* is enabled, then  $\text{LS}(L_{n1})$  conveys the SHA256 hash of the whole file to the  $\text{LD}(L_{n1})$ . Once the file is written to the disk successfully,  $\text{LD}(L_{n1})$  computes the hash of transferred file and compares it with the received hash value to verify the integrity of the file transfer. This check acts as an additional error detection mechanism over and above the layer specific error detection policies used in MLED.

### 1.2.2 $\text{LS}(L_{ij})$ and $\text{LD}(L_{ij})$ , $i < n$

Except for level  $n$ , the  $\text{LS}(L_{ij})$  receives the data from the layer at level  $(i+1)$ , splits the data according to its payload length, calculates a checksum according to the error detection policy  $P_{ij}$ , encapsulates the data and the checksum to create the *Protocol Data Unit (PDU)* for layer  $L_{ij}$  and pushes it down to the layer at level  $(i-1)$ , except at level 1. At level 1, the  $\text{LS}(L_{1j})$  transmits the data onto the underlying physical network.

At level 1, the  $\text{LD}(L_{1j})$  on the next node receives the data from the physical link. At level  $i > 1$ , the  $\text{LD}(L_{ij})$  receives the data from layer at level  $(i-1)$ , reassembles this data to form its PDU, decapsulates the checksum from the data, and verifies this checksum according to the error check policy  $P_{ij}$ . If no errors are detected in the data, it sends an *acknowledgement (ACK)*, otherwise it sends a *Negative Acknowledgement (NACK)* to its corresponding  $\text{LS}(L_{ij})$  and requests for a retransmission of the PDU.

### 1.2.3 Relay Process and Virtual Links

For a layer  $L_{ij}$ , any process between the  $\text{LS}(L_{ij})$  and  $\text{LD}(L_{ij})$  is a relay process (R). In Figures 1.1, the process (2,2R), (3,3R) and (2,4R) are relay processes for layers  $L_{21}$ ,  $L_{31}$  and  $L_{22}$  respectively. The relay process at level  $i$  connects two underlying layers by transferring the data from the  $\text{LD}(L_{(i-1)j})$  to the  $\text{LS}(L_{(i-1)(j+1)})$ . With this data transfer, it establishes the functionality of the layer at level  $i$  that

governs the policy over a larger part of the underlying physical network, encompassing multiple portions of the physical network.

Except for the layer  $L_{n1}$  and the layers at level 1, each layer contains at least two virtual links. The virtual link connects the  $LS(L_{ij})$  to a relay and this relay to the  $LD(L_{ij})$ . For example, in Figure 1.1 the layer  $L_{31}$  is a logical connection between  $LS(L_{31})$  and  $LD(L_{31})$  composed of two virtual links connecting  $LS(L_{31})$  to relay process (3,3R) and relay process (3,3R) to  $LD(L_{31})$ .

#### 1.2.4 File Transfer Example

To understand the MLED architecture, consider the file transfer from source node to the destination node in Figure 1.1. Level 4 has one layer  $L_{41}$  with  $LS(L_{41})$  and  $LD(L_{41})$ .  $LS(L_{41})$  starts a server and receives chunks according to its payload length. It forwards these chunks to layer  $L_{31}$  for delivery to the  $LD(L_{41})$  on Node 5. The  $LS(L_{31})$  splits the data according to the payload length of  $L_{31}$  and adds its own header creating its own PDU. This header includes a checksum, which is computed based on the error detection policy defined by  $P_{31}$  and the sequence number for this chunk. This PDU is passed to layer  $L_{21}$  to transfer it to  $LD(L_{31})$  on Node 5 via the relay process (3,3R). Similarly, layer  $L_{21}$  is responsible for transporting the PDUs received from layer  $L_{31}$ .  $LS(L_{21})$  encapsulates these PDUs within its own PDUs and forwards them to  $LD(L_{21})$  through the relay process (2,2R). Finally, the layer  $L_{11}$  works with the physical links represented by the purple line, moving the PDUs of layer  $L_{11}$  with their own unique checks from one node to the next one. After  $LD(L_{11})$  receives its PDU via the physical link, it decapsulates the PDU, validates the data with the attached checksum and sends an ACK or NACK to  $LS(L_{11})$ . Once verified as error-free, it discards the header of layer  $L_{11}$  from the PDU and passes the data to relay process (2,2R), which then forwards it to  $LS(L_{12})$ . Here relay process (2,2R) connects layer  $L_{11}$  and  $L_{12}$  to realize the functionality of layer  $L_{21}$ . This procedure repeats until  $LD(L_{21})$  reassembles the incoming data chunks into a PDU of layer  $L_{21}$ , conducts the integrity check specific to policy  $P_{21}$  and sends an ACK/NACK to  $LS(L_{21})$ .  $LD(L_{21})$  passes the verified data to the relay process (3,3R). Process (3,3R) forwards the data to  $LS(L_{22})$  which in turn forwards it to the  $LD(L_{22})$ . Process (3,3R) relays the data from layer  $L_{21}$  to layer  $L_{22}$  realizing the functionality of layer  $L_{31}$ . When  $LD(L_{31})$  forms its PDU by reassembling the data received from  $LD(L_{22})$ , it checks the PDU for errors in accordance with the error detection policy  $P_{31}$  and sends an ACK/NACK to  $LS(L_{31})$ . The data transfer sequence culminates at the destination node at level 4, where  $LD(L_{41})$  writes the data to disk storage.





## Chapter 2

# Configuration File

The MLED architecture relies on a configuration file to set up and customize the system for file transfer operations. This configuration file is a JSON document that specifies critical parameters, policies, and settings for each component of the architecture. The configuration file enables users to define the behavior of nodes, layer managers, and error detection policies, providing flexibility and control over the MLED's operation. Understanding these settings is essential for deploying and troubleshooting the MLED architecture effectively.

Below is an explanation of the keys used in the configuration file:

### Explanation of Configuration Keys

Key	Description
FileName	Name of the file to be transferred.
ReadChunkSizeInKB	Size of each chunk (in KB) when reading the file at level $n$ .
WriteChunkSizeInKB	Size of each chunk (in KB) when writing the file at level $n$ .
MemoryUsageLimitInGB	Maximum memory usage allowed (in GB) to buffer the chunks read from the file. The system must have at least this much free memory.
PrefetchInputFile	Boolean value indicating whether to prefetch the input file into memory before starting the transfer. The system must have free memory equal to the size of the file to be transferred. Otherwise, MLED automatically disables this option.
MLEDServerPort	Port used by the MLED central manager server to communicate with the workers.
HeartBeatTimeInterval	Time interval (in milliseconds) between heartbeat messages exchanged between the workers and the central manager server.
HeartBeatTimeOut	Timeout (in milliseconds) for heartbeat messages. In current implementation, no action is taken on timeout as we don't handle machine failures.
RetryTimeInterval	Time interval (in milliseconds) before retrying a failed connection request between various MLED servers and clients.
FacilityPortIP	IP address of the remote server where the file is stored. Use the local IP if the file is stored locally.
RemoteFilePath	Path to the directory where file to be transferred is stored. It could be a local file path or remote file path. This path must end with a '/' and not contain '.'.
RemoteSSHKey	Name of the SSH key used for accessing the remote file. If the file is stored locally then you need an SSH key to SSH into local machine as well. This key should be stored in <code>\$mled_dir_home/node_Node1/</code> and <code>\$mled_dir_home/scripts/</code> directory. The name should not contain any '.'.
RemoteUsername	Username for remote access or username on the local machine for local access.
FileServerPort	Port used for file reading operations remotely or locally.

Key	Description
MaxPortsToTry	In case the assigned <code>FileServerPort</code> is occupied, this parameter defines the maximum number of ports to attempt before failing.
MaxAttemptsPerPort	Maximum number of attempts to try to connect per port to read the file.
Network	Type of network protocol used (e.g., <code>tcp</code> ). Currently, only <code>tcp</code> is supported as a valid protocol.
Routing	Type of routing policy (e.g., <code>static</code> ). Currently, only <code>static</code> routing is supported.
MaxLayers	Maximum number of layers in the architecture i.e., the value of $n$
VerifyWholeFileHash	Boolean indicating whether to verify the entire end-to-end file hash for integrity. It uses <code>SHA256</code> as the hashing algorithm.
CreateFileToBypassCrcAndChecksum	Boolean indicating whether MLED is running to get the metadata to create an <i>adversarial test file</i> . User should not set this value to <code>True</code> . This parameter is set by MLED automatically depending on the value of <code>BypassBothCrcAndChecksum</code> parameter.
BypassBothCrcAndChecksum	Boolean indicating whether the errors introduced by the adversarial error model are undetected by both CRC and checksum policy employed at different levels. If <code>False</code> , then MLED introduces adversarial errors that are undetected only at the level of introduction.
Nodes	Array of objects representing each node in the system, including their configuration, processes, and forwarding rules.
NodeId	Unique identifier for the node.
NodeName	Name of the node (e.g., <code>Node1</code> , <code>Node2</code> ).
NodeIp	IP address of the data server running on node.
NodeType	Type of the node (e.g., <code>Source</code> , <code>Intermediate</code> , <code>Destination</code> ).
NodePort	Port used by the data server for communication.
forwardTo	Array of objects indicating the next MLED worker to which data will be forwarded.
processes	Array of objects specifying the MLED processes running on the node, including details like process ID and level depth for checks.
ProcessId	Unique ID assigned to the each process on a node.
LevelDepthUptoWhichChecksPerformed	Starting from level 1 what is the depth up to which this process will go in the MLED hierarchy.
ExecutorDetails	<p>This key represents the MLED processes that will run on this node. Each triplet separated by a hyphen (-) represents an MLED process. A triplet is structured as follows: the first part indicates the <b>node ID of the LS</b> for the process, the second part specifies whether the process is <b>enabled (E)</b> or <b>disabled (D)</b>, and the third part denotes the <b>node ID of the LD</b> of the process. It is important to note that all relay processes are always disabled, which means their status in the second part of the triplet will always be represented by <code>D</code>.</p> <p>For example, consider the sequence <code>1.D.2-2.E.3-1.D.3</code>. The first triplet, <code>1.D.2</code>, indicates that the process has <b>Node 1</b> as its LS, is <b>disabled</b>, and <b>Node 2</b> is its LD. The second triplet, <code>2.E.3</code>, signifies a process with <b>Node 2</b> as LS, is <b>enabled</b>, and <b>Node 3</b> is its LD. Finally, the third triplet, <code>1.D.3</code>, shows a process that has <b>Node 1</b> as LS, is <b>disabled</b>, and has <b>Node 3</b> as its LD.</p> <p>This structured representation provides a clear and concise way to describe the state and LS and LD of processes at various layers within the MLED system, eliminating the need for <i>destination resolution</i> in the current implementation.</p>

Key	Description
	The order in which these processes are defined is related to their visual position in MLED architecture. For each node, we define processes from level 1 to level $n$ and on each level, we define them from left to right. For example, in Fig. 1.1, the configuration of Node 2 would define its processes in the following order: process(1,2D), process(1,2S) and process(2,2R).
IsRelayProcess	A boolean option to indicate if the processes defined in <b>ExecutorDetails</b> has a relay process or not. If set, this means that the last process defined in the key <b>ExecutorDetails</b> is the relay process.
NodeAckIp	IP address used by the ACK/NACK server for acknowledgment messages.
NodeAckPort	Port used by the ACK/NACK server for acknowledgment messages.
forwardAckTo	Array of objects indicating the next MLED worker to which ACK/NACK will be forwarded.
LayerManagers	Array of objects representing layer managers, their configurations, such as the <b>NodeID</b> , <b>NodeName</b> , <b>NopeIp</b> and <b>NopePort</b> of the node on which a particular manager would run and the routing tables for the layers that are managed by this layer manager.
LayerManagerId	Unique identifier for the layer manager, representing the level it manages.
LayerStructure	Represents the structure of the level managed by the layer manager. It essentially consists of node IDs separated by a ‘.’ (dot), which indicate a virtual link between processes running on these nodes at the level managed by the layer manager.
SubLayerDetails	Array of objects describing the details of each layer at the level managed by the layer manager, such as window size, check policy, parameter, chunk size, PER, and cross-traffic generation. Cross-traffic generation parameters allow precise control over the generation and management of cross traffic in the system, enabling customization based on specific requirements. Cross traffic could be used for simulating multiple MLED flows through a particular source process to study the impact of increasing processing load on MLED.
NodeDetails	This array defines the LS and LD of the layer on which these particular parameters defined in <b>SubLayerDetails</b> would be applied.
GenerateCrossTraffic	A boolean parameter that determines whether cross traffic should be generated. If set to <b>true</b> , cross traffic will be enabled; if set to <b>false</b> , it will be disabled.
CrossTrafficGenerationRateInBps	Specifies the rate at which cross traffic is generated, measured in bytes per second (Bps).
NumberOfInstancesOfCrossTraffic	Indicates the number of threads or instances responsible for generating cross traffic. Each instance will generate traffic at the rate specified by <b>CrossTrafficGenerationRateInBps</b> .
routingTable	Array of objects defining the routing rules for the layer manager, including source, destination, and next-hop details. The <b>Priority</b> key in the <b>routingTable</b> is not utilized in the current implementation. However, it is reserved for potential future use, where it could be employed to prioritize one MLED path over another. At present, the implementation only supports a single path, so prioritization is not necessary or applicable.

## 2.1 Example Main Configuration File

Listing 2.1: Main Configuration File created by the user to customize MLED

---

```

1
2 {
3   "FileName": "100MB",
4   "ReadChunkSizeInKB": 1024,
5   "WriteChunkSizeInKB": 1024,
6   "MemoryUsageLimitInGB": 10,
7   "PrefetchInputFile": false,
8   "MLEDServerPort": 50000,
9   "HeartBeatTimeInterval": 1000,
10  "HeartBeatTimeOut": 10000,
11  "RetryTimeInterval": 1000,
12  "FacilityPortIP": "10.52.0.16",
13  "RemoteSSHKey": "mledcpp",
14  "RemoteUsername": "cc",
15  "FileServerPort": 51000,
16  "MaxPortsToTry": 10,
17  "MaxAttemptsPerPort": 5,
18  "Network": "tcp",
19  "Routing": "static",
20  "MaxLayers": "3",
21  "VerifyWholeFileHash": true,
22  "CreateFileToBypassCrcAndChecksum": false,
23  "BypassBothCrcAndChecksum": false,
24  "Nodes": [
25    {
26      "NodeId": 1,
27      "NodeName": "Node1",
28      "NodeIp": "10.52.0.16",
29      "NodeType": "Source",
30      "NodePort": 8000,
31      "forwardTo": [
32        {
33          "Ip": "10.52.0.16",
34          "Port": 8100,
35          "NodeName": "Node2",
36          "NodeId": 2
37        }
38      ],
39      "processes": [
40        {
41          "ProcessId": 1,
42          "LevelDepthUptoWhichCheckIsPerformed": 3,
43          "ExecutorDetails": "1.E.2-1.E.3-1.E.5",
44          "IsRelayProcess": false
45        }
46      ],
47      "NodeAckIp": "10.52.0.16",
48      "NodeAckPort": 9000,
49      "forwardAckTo": []
50    },
51    {
52      "NodeId": 2,
53      "NodeName": "Node2",
54      "NodeIp": "10.52.0.16",
55      "NodeType": "Intermediate",
56      "NodePort": 8100,
57      "forwardTo": [
58        {
59          "Ip": "10.52.0.16",
60          "Port": 8200,
61          "NodeName": "Node3",
62          "NodeId": 3
63        }
64      ],
65      "processes": [
66        {
67          "ProcessId": 1,
68          "LevelDepthUptoWhichCheckIsPerformed": 1,
69          "ExecutorDetails": "1.E.2-2.E.3-1.D.3",

```

```

70         "IsRelayProcess": true
71     }
72 ],
73     "NodeAckIp": "10.52.0.16",
74     "NodeAckPort": 9500,
75     "forwardAckTo": [
76     {
77         "Ip": "10.52.0.16",
78         "Port": 9000,
79         "NodeName": "Node1",
80         "NodeId": 1
81     }
82 ],
83 },
84 {
85     "NodeId": 3,
86     "NodeName": "Node3",
87     "NodeIp": "10.52.0.16",
88     "NodeType": "Intermediate",
89     "NodePort": 8200,
90     "forwardTo": [
91     {
92         "Ip": "10.52.0.16",
93         "Port": 8300,
94         "NodeName": "Node4",
95         "NodeId": 4
96     }
97 ],
98     "processes": [
99     {
100         "ProcessId": 1,
101         "LevelDepthUptoWhichCheckIsPerformed": 2,
102         "ExecutorDetails": "2.E.3-3.E.4-1.E.3-3.E.5-1.D.5",
103         "IsRelayProcess": true
104     }
105 ],
106     "NodeAckIp": "10.52.0.16",
107     "NodeAckPort": 9200,
108     "forwardAckTo": [
109     {
110         "Ip": "10.52.0.16",
111         "Port": 9500,
112         "NodeName": "Node2",
113         "NodeId": 2
114     }
115 ],
116 },
117 {
118     "NodeId": 4,
119     "NodeName": "Node4",
120     "NodeIp": "10.52.0.16",
121     "NodeType": "Intermediate",
122     "NodePort": 8300,
123     "forwardTo": [
124     {
125         "Ip": "10.52.0.16",
126         "Port": 8400,
127         "NodeName": "Node5",
128         "NodeId": 5
129     }
130 ],
131     "processes": [
132     {
133         "ProcessId": 1,
134         "LevelDepthUptoWhichCheckIsPerformed": 1,
135         "ExecutorDetails": "3.E.4-4.E.5-3.D.5",
136         "IsRelayProcess": true
137     }
138 ],
139     "NodeAckIp": "10.52.0.16",
140     "NodeAckPort": 9300,
141     "forwardAckTo": [
142     {

```

```

143         "Ip": "10.52.0.16",
144         "Port": 9200,
145         "NodeName": "Node3",
146         "NodeId": 3
147     }
148 ]
149 },
150 {
151     "NodeId": 5,
152     "NodeName": "Node5",
153     "NodeIp": "10.52.0.16",
154     "NodeType": "Destination",
155     "NodePort": 8400,
156     "forwardTo": [],
157     "processes": [
158         {
159             "ProcessId": 1,
160             "LevelDepthUptoWhichCheckIsPerformed": 3,
161             "ExecutorDetails": "4.E.5-3.E.5-1.E.5",
162             "IsRelayProcess": false
163         }
164     ],
165     "NodeAckIp": "10.52.0.16",
166     "NodeAckPort": 9400,
167     "forwardAckTo": [
168         {
169             "Ip": "10.52.0.16",
170             "Port": 9300,
171             "NodeName": "Node4",
172             "NodeId": 4
173         }
174     ]
175 }
176 ],
177 "LayerManagers": [
178     {
179         "LayerManagerId": 1,
180         "NodeId": 1,
181         "NodeName": "node1",
182         "NodeIp": "10.52.0.16",
183         "NodePort": 7000,
184         "LayerStructure": "1.2.3.4.5",
185         "SubLayerDetails": [
186             {
187                 "NodeDetail": [
188                     "1-S",
189                     "2-D"
190                 ],
191                 "WindowSize": 461,
192                 "CheckPolicy": "CRC",
193                 "CheckParameter": "0x04C11DB7",
194                 "ChunkSize": 65536,
195                 "PacketErrorRate": 0.1,
196                 "GenerateCrossTraffic": false,
197                 "CrossTrafficGenerationRateInBps": 9000,
198                 "NumberOfInstancesOfCrossTraffic": 30
199             },
200             {
201                 "NodeDetail": [
202                     "2-S",
203                     "3-D"
204                 ],
205                 "WindowSize": 446,
206                 "CheckPolicy": "CRC",
207                 "CheckParameter": "0x04C11DB7",
208                 "ChunkSize": 65536,
209                 "PacketErrorRate": 0,
210                 "GenerateCrossTraffic": false,
211                 "CrossTrafficGenerationRateInBps": 10000,
212                 "NumberOfInstancesOfCrossTraffic": 20
213             }
214         ],
215         "NodeDetail": [

```

```

216         "3-S",
217         "4-D"
218     ],
219     "WindowSize": 552,
220     "CheckPolicy": "CRC",
221     "CheckParameter": "0x04C11DB7",
222     "ChunkSize": 65536,
223     "PacketErrorRate": 0,
224     "GenerateCrossTraffic": false,
225     "CrossTrafficGenerationRateInBps": 11000,
226     "NumberOfInstancesOfCrossTraffic": 40
227 },
228 {
229     "NodeDetail": [
230         "4-S",
231         "5-D"
232     ],
233     "WindowSize": 475,
234     "CheckPolicy": "CRC",
235     "CheckParameter": "0x04C11DB7",
236     "ChunkSize": 65536,
237     "PacketErrorRate": 0,
238     "GenerateCrossTraffic": false,
239     "CrossTrafficGenerationRateInBps": 12000,
240     "NumberOfInstancesOfCrossTraffic": 60
241 }
242 ],
243 "routingTable": [
244     {
245         "Source": 1,
246         "Destination": 5,
247         "Priority": "1",
248         "NextHopNodeId": 2
249     },
250     {
251         "Source": 1,
252         "Destination": 4,
253         "Priority": "1",
254         "NextHopNodeId": 2
255     },
256     {
257         "Source": 1,
258         "Destination": 3,
259         "Priority": "1",
260         "NextHopNodeId": 2
261     },
262     {
263         "Source": 1,
264         "Destination": 2,
265         "Priority": "1",
266         "NextHopNodeId": 2
267     },
268     {
269         "Source": 2,
270         "Destination": 5,
271         "Priority": "1",
272         "NextHopNodeId": 3
273     },
274     {
275         "Source": 2,
276         "Destination": 4,
277         "Priority": "1",
278         "NextHopNodeId": 3
279     },
280     {
281         "Source": 2,
282         "Destination": 3,
283         "Priority": "1",
284         "NextHopNodeId": 3
285     },
286     {
287         "Source": 3,
288         "Destination": 5,

```

```

289         "Priority": "1",
290         "NextHopNodeId": 4
291     },
292     {
293         "Source": 3,
294         "Destination": 4,
295         "Priority": "1",
296         "NextHopNodeId": 4
297     },
298     {
299         "Source": 4,
300         "Destination": 5,
301         "Priority": "1",
302         "NextHopNodeId": 5
303     },
304     {
305         "Source": 2,
306         "Destination": 1,
307         "Priority": "1",
308         "NextHopNodeId": 1
309     },
310     {
311         "Source": 3,
312         "Destination": 2,
313         "Priority": "1",
314         "NextHopNodeId": 2
315     },
316     {
317         "Source": 3,
318         "Destination": 1,
319         "Priority": "1",
320         "NextHopNodeId": 2
321     },
322     {
323         "Source": 4,
324         "Destination": 3,
325         "Priority": "1",
326         "NextHopNodeId": 3
327     },
328     {
329         "Source": 4,
330         "Destination": 2,
331         "Priority": "1",
332         "NextHopNodeId": 3
333     },
334     {
335         "Source": 4,
336         "Destination": 1,
337         "Priority": "1",
338         "NextHopNodeId": 3
339     },
340     {
341         "Source": 5,
342         "Destination": 1,
343         "Priority": "1",
344         "NextHopNodeId": 4
345     },
346     {
347         "Source": 5,
348         "Destination": 2,
349         "Priority": "1",
350         "NextHopNodeId": 4
351     },
352     {
353         "Source": 5,
354         "Destination": 3,
355         "Priority": "1",
356         "NextHopNodeId": 4
357     },
358     {
359         "Source": 5,
360         "Destination": 4,
361         "Priority": "1",

```



```

362         "NextHopNodeId": 4
363     }
364 ]
365 },
366 {
367     "LayerManagerId": 2,
368     "NodeId": 1,
369     "NodeName": "node1",
370     "NodeIp": "10.52.0.16",
371     "NodePort": 7100,
372     "LayerStructure": "1.3.5",
373     "SubLayerDetails": [
374         {
375             "NodeDetail": [
376                 "1-S",
377                 "3-D"
378             ],
379             "WindowSize": 293,
380             "CheckPolicy": "CRC",
381             "CheckParameter": "0x9B",
382             "ChunkSize": 131072,
383             "PacketErrorRate": 0,
384             "GenerateCrossTraffic": false,
385             "CrossTrafficGenerationRateInBps": 17000,
386             "NumberOfInstancesOfCrossTraffic": 80
387         },
388         {
389             "NodeDetail": [
390                 "3-S",
391                 "5-D"
392             ],
393             "WindowSize": 419,
394             "CheckPolicy": "CRC",
395             "CheckParameter": "0x9B",
396             "ChunkSize": 131072,
397             "PacketErrorRate": 0,
398             "GenerateCrossTraffic": false,
399             "CrossTrafficGenerationRateInBps": 20000,
400             "NumberOfInstancesOfCrossTraffic": 90
401         }
402     ],
403     "routingTable": [
404         {
405             "Source": 1,
406             "Destination": 5,
407             "Priority": "1",
408             "NextHopNodeId": 2
409         },
410         {
411             "Source": 1,
412             "Destination": 4,
413             "Priority": "1",
414             "NextHopNodeId": 2
415         },
416         {
417             "Source": 1,
418             "Destination": 3,
419             "Priority": "1",
420             "NextHopNodeId": 2
421         },
422         {
423             "Source": 1,
424             "Destination": 2,
425             "Priority": "1",
426             "NextHopNodeId": 2
427         },
428         {
429             "Source": 2,
430             "Destination": 5,
431             "Priority": "1",
432             "NextHopNodeId": 3
433         },
434     ]

```

```

435         "Source": 2,
436         "Destination": 4,
437         "Priority": "1",
438         "NextHopNodeId": 3
439     },
440     {
441         "Source": 2,
442         "Destination": 3,
443         "Priority": "1",
444         "NextHopNodeId": 3
445     },
446     {
447         "Source": 3,
448         "Destination": 5,
449         "Priority": "1",
450         "NextHopNodeId": 4
451     },
452     {
453         "Source": 3,
454         "Destination": 4,
455         "Priority": "1",
456         "NextHopNodeId": 4
457     },
458     {
459         "Source": 4,
460         "Destination": 5,
461         "Priority": "1",
462         "NextHopNodeId": 5
463     },
464     {
465         "Source": 2,
466         "Destination": 1,
467         "Priority": "1",
468         "NextHopNodeId": 1
469     },
470     {
471         "Source": 3,
472         "Destination": 2,
473         "Priority": "1",
474         "NextHopNodeId": 2
475     },
476     {
477         "Source": 3,
478         "Destination": 1,
479         "Priority": "1",
480         "NextHopNodeId": 2
481     },
482     {
483         "Source": 4,
484         "Destination": 3,
485         "Priority": "1",
486         "NextHopNodeId": 3
487     },
488     {
489         "Source": 4,
490         "Destination": 2,
491         "Priority": "1",
492         "NextHopNodeId": 3
493     },
494     {
495         "Source": 4,
496         "Destination": 1,
497         "Priority": "1",
498         "NextHopNodeId": 3
499     },
500     {
501         "Source": 5,
502         "Destination": 1,
503         "Priority": "1",
504         "NextHopNodeId": 4
505     },
506     {
507         "Source": 5,

```

```

508         "Destination": 2,
509         "Priority": "1",
510         "NextHopNodeId": 4
511     },
512     {
513         "Source": 5,
514         "Destination": 3,
515         "Priority": "1",
516         "NextHopNodeId": 4
517     },
518     {
519         "Source": 5,
520         "Destination": 4,
521         "Priority": "1",
522         "NextHopNodeId": 4
523     }
524 ]
525 },
526 {
527     "LayerManagerId": 3,
528     "NodeId": 1,
529     "NodeName": "node1",
530     "NodeIp": "10.52.0.16",
531     "NodePort": 7300,
532     "LayerStructure": "1.5",
533     "SubLayerDetails": [
534         {
535             "NodeDetail": [
536                 "1-S",
537                 "5-D"
538             ],
539             "WindowSize": 325,
540             "CheckPolicy": "checksum",
541             "CheckParameter": "2",
542             "ChunkSize": 262144,
543             "PacketErrorRate": 0,
544             "GenerateCrossTraffic": false,
545             "CrossTrafficGenerationRateInBps": 22000,
546             "NumberOfInstancesOfCrossTraffic": 15
547         }
548     ],
549     "routingTable": [
550         {
551             "Source": 1,
552             "Destination": 5,
553             "Priority": "1",
554             "NextHopNodeId": 3
555         },
556         {
557             "Source": 1,
558             "Destination": 3,
559             "Priority": "1",
560             "NextHopNodeId": 3
561         },
562         {
563             "Source": 3,
564             "Destination": 5,
565             "Priority": "1",
566             "NextHopNodeId": 5
567         },
568         {
569             "Source": 3,
570             "Destination": 1,
571             "Priority": "1",
572             "NextHopNodeId": 1
573         },
574         {
575             "Source": 5,
576             "Destination": 1,
577             "Priority": "1",
578             "NextHopNodeId": 3
579         },
580     ]

```

```
581         "Source": 5,  
582         "Destination": 3,  
583         "Priority": "1",  
584         "NextHopNodeId": 3  
585     }  
586 ]  
587 }  
588 ]  
589 }
```

---

Due to the complex structure of the configuration file required for customizing the MLED framework, users may find it challenging to manually create or modify the file accurately. To simplify this process, a graphical user interface (GUI) is available at <https://mled.bu.edu> to create configuration files for both balanced and unbalanced MLED architectures. This GUI streamlines the process by providing an interface to generate the configuration file with minimal effort. By leveraging the GUI, users can avoid potential errors and ensure the configuration aligns with their specific network requirements.