

Operational impact of maximum LRIT periodic information exchange under the coastal state entitlement

TECHNICAL REPORT

Contents

List of Acronyms	2
1 Executive Summary	3
2 Background and Regulatory Context	3
2.1 <i>LRIT Entitlement Under SOLAS Regulation V/19-1</i>	3
2.2 <i>The LRIT Data Distribution Plan (DDP)</i>	4
2.3 <i>Current operational status of the LRIT system</i>	4
2.4 <i>LRIT Capacity Implications and Mitigation Mechanisms for Increased Traffic</i>	5
Pool of Connections and Processing Capacity in Distributed Communication Systems..	5
Session Lifecycle at the DC-IDE Interface	5
LRIT Governance	6
3 Questions to be answered	6
4 Methodology	6
4.1 <i>Data Sample</i>	6
4.2 <i>Limitations and Caveats</i>	7
4.3 <i>Treatment of non-standard-rate transmissions</i>	7
4.4 <i>Analysis Procedure</i>	7
4.5 <i>Operational and theoretical load under full coastal entitlement exercise</i>	8
5 Aggregated data overview	8
6 Results	10
6.1 <i>Average Load Estimation</i>	10
6.2 <i>Temporal Analysis — Transmission Peak Times</i>	10
6.3 <i>Increase in storage for IDE</i>	13
6.4 <i>Operational impact for Data Centres</i>	14
7 Conclusions.....	14

List of Acronyms

DC	LRIT Data Centre
DDP	Data Distribution Plan
IDE	International Data Exchange
IMO	International Maritime Organization
LRIT	Long-Range Identification and Tracking
NM	Nautical Mile
SOLAS	International Convention for the Safety of Life at Sea
UTC	Coordinated Universal Time

1 Executive Summary

1.1 This report presents the findings of an analysis conducted by the LRIT Coordinator to estimate the volume and geographic distribution of LRIT messages that would flow through the International LRIT Data Exchange (IDE) if all Contracting Governments were to exercise, in full, their entitlement to receive LRIT information under SOLAS regulation V/19-1.

1.2 The analysis is based on sample data drawn from the 2025 audit schedule, covering 71 Data Centres (DCs). For each LRIT position report received in the sample, the geographic coordinates were cross-referenced against the LRIT Data Distribution Plan (DDP) in force at the moment of Timestamp 4, to determine all Contracting Governments (CG) within whose 1,000-nautical-mile coastal zone the transmitting vessel was located at that time. A database record was created for each such entitled CG, representing one message that, under a fully exercised entitlement regime, would be routed through the IDE to the LRIT Data Centre (DC) associated with the CG. The key findings are displayed in table 1:

Table 1 - Analysis key findings

Metric	Value
Total LRIT positions analysed (all DCs)	3,042,131
Total database records generated	17,604,806
Average records per position	5.7
Estimated annual IDE message volume (extrapolated)	214,191,806
Peak transmission hour (message/sec)	292 message/sec
Number of Contracting Governments as data recipients	124 SOLAS Contracting Governments ¹
Data Centres covered in the sample	71

1.3 The results offer a quantitative foundation for a comprehensive evaluation of the operational impact, concerning IDE throughput, DC ingestion load, and routing infrastructure, of a scenario wherein all Contracting Governments actively receive the LRIT information to which they are entitled under SOLAS.

2 Background and Regulatory Context

2.1 LRIT Entitlement Under SOLAS Regulation V/19-1

2.1.1 SOLAS regulation V/19-1 establishes the obligations of ships to transmit LRIT information and the rights of CGs to receive it. CGs are entitled to receive LRIT information in three distinct capacities: as a flag State, in respect of ships entitled to fly their flag; as a port State, in respect of ships proceeding to a port or place under their jurisdiction; and as a coastal State, in respect of ships navigating within 1,000 nautical miles of their coast, regardless of their flag.

2.1.2 Coastal State entitlement is the primary factor driving message volume in a fully operational LRIT system. Because these entitlement zones extend up to 1,000 NM from the coast of each CG, they overlap extensively across the world's oceans. As a result, a single ship's position can simultaneously fall within the zones of multiple CGs, each independently entitled to receive that LRIT information, directly increasing the volume of messages flowing through the system.

¹ 10 non-metropolitan territories and two special administrative regions are part of the LRIT, and they may be associated with LRIT DCs other than those used by their Contracting Governments.

2.1.3 To receive LRIT information in a given area, a CG must establish a Coastal State Standing Order (CSSO) in the IMO GISIS system for its LRIT polygons representing the geographic area it wishes to monitor. When processing position reports from their associated ships, DCs check each ship's location against the existing CSSOs to determine which LRIT Data Users (LDUs) are entitled to receive that information. This entire process runs automatically, without human intervention, on the DC servers.

2.2 *The LRIT Data Distribution Plan (DDP)*

2.2.1 CGs are required to upload information on their Internal Waters, Territorial Waters, and 1000 NM zones to the IMO GISIS database. This information is represented as polygons in GML format and distributed to all DCs via the DDP file and its subsequent updates. For the purposes of this analysis, the version of the DDP in force at the time of each Timestamp 4 event was used to determine whether a given ship position fell within a particular government's entitlement zone. This approach ensures that the analysis accurately reflects the state of the DDP at the moment each message was received by the DC, thereby accounting for any updates made to the plan throughout the sample period.

2.3 *Current operational status of the LRIT system*

2.3.1 In practice, only a very small number of CGs have configured CSSOs covering the full extent of their 1,000 NM entitlement zone. Most CGs instead limit data receipt to ships flying their own flag, making no use of their coastal entitlement. As a result, the LRIT traffic volume currently observed on the IDE represents only a fraction of what would be generated if all CGs fully exercised their entitlement.

2.3.2 The purpose of this analysis is to quantify that gap, providing the LRIT Coordinator and Contracting Governments with a data-driven estimate of the operational implications of full entitlement exercise.

2.3.3 As a reference, in 2025, the IDE received and processed 13,479,083 periodic messages under CSSO, including those addressing the IDF. For this study, the IDF was included because it significantly affects message flow.

2.3.4 The flow of messages in 2025 indicates that 88.2% of all messages in the system are LRIT information provided under CSSO, 0.7% are provided in response to LRIT requests, such as port requests or SAR requests, and 10.1% are other messages, such as LRIT Requests, system status, and receipt messages.

2.3.5 The requirements for the IDE related to timely message processing are: to use a store-and-forward buffer to ensure LRIT information is received (MSC.263(84)/Rev.1, paragraph 10.3.3); to process and handle any input within 30 seconds of receipt and to provide the appropriate output; and to be capable of receiving and processing at least 100 reports per second (MSC.1/Circ.1259/Rev.9, paragraph 4.1.1). In addition, MSC.1/Circ.1259/Rev.9, paragraph 3.3.1.2, states that designers and implementers are advised to adopt asynchronous patterns and reliable, non-blocking mechanisms to receive the HTTP response message, to ensure optimal performance.

2.3.6 The requirement to process and handle any input within 30 seconds of receipt and to provide the appropriate output is verified in every IDE audit, with no findings in the last audits.

2.3.7 The IDE's capability to receive and process at least 100 reports per second was reportedly tested twice in the Developmental Testing environment in 2013 and 2016, during testing of IDE software versions 2.2.0 and 3.0.0, respectively. A new test is being conducted by

the IDE at the time this report is being written.

2.4 LRIT Capacity Implications and Mitigation Mechanisms for Increased Traffic

2.4.1 The study aimed to assess the impact of the increased message traffic on the LRIT system. To perform the analysis, it was necessary to verify the system's current capacity against the existing requirements. The risk to consider is that surges in transmission could reach a point where the IDE would be unable to accept new connections while still processing existing ones. Thus, it is not sufficient that the average number of incoming messages is below 100 messages per second, because the average metric does not reflect surges or idle periods.

2.4.2 The LRIT architecture is a resilient distributed system based on an asynchronous communications protocol using SOAP over TLS 1.2 that provides the data interface between components. Every component offers and consumes services from other components to operate within the system. Each new communication creates a session that should end with the message being sent and received. To receive these incoming communications, each component provides a pool of connections to consumers. The pool of communications is standard in network systems; every server would have a configured pool to handle incoming traffic.

Pool of Connections and Processing Capacity in Distributed Communication Systems

2.4.3 The capacity to maintain a pool of incoming communication sessions and the capacity to process the messages delivered within those sessions are independent parameters in a server. For example, in theory, the IDE may maintain a pool of open transport connections from transmitting LRIT Data Centres, but be unable to process the messages carried over those connections due to resource saturation or internal queue congestion. Conversely, the IDE may have available processing capacity while being unable to accept new connections because the connection pool has reached its configured limit. The latter scenario can occur during a surge of incoming communications or as a result of a malicious denial-of-service (DoS) attack.

Session Lifecycle at the DC-IDE Interface

2.4.4 In SOAP over TLS 1.2 (typically over HTTP), a communication exchange is bound to the lifecycle of the underlying transport connection and the request–response pattern. A connection may be closed upon: completion of the HTTP/SOAP request–response exchange; expiry of the transmitting DC's application-level response timeout without receiving a reply; network- or transport-layer closure by either party; or release by the IDE of an idle connection exceeding its configured idle timeout.

2.4.5 When a timeout occurs at the transmitting DC, the connection may be closed or abandoned without notification to the IDE at the application layer. In such cases, any processing completed subsequently by the IDE cannot be delivered to the originating Data Centre over that connection.

2.4.6 Timeout parameters exist independently at each communication layer, from the TCP transport layer through the TLS and HTTP layers to the application layer. The connection pool is an application-level resource management mechanism that governs the number of simultaneous transport connections that the receiving component maintains. These parameters are independent of the component's processing capacity and must be considered together to ensure reliable handling of incoming communications.

LRIT Governance

2.4.7 Timeout configuration and connection pool limits operate independently of the processing capacity of a component; however, all three parameters must be considered together to ensure that incoming connections and messages are handled reliably under varying load conditions. The current LRIT technical documentation establishes requirements for processing capacity and processing time, but provides no guidance on timeout configuration or connection pool sizing.

2.4.8 A further mitigation mechanism widely implemented in distributed systems is message redelivery upon failure. In the LRIT system, the Technical Specifications (MSC.1/Circ.1259/Rev.9, paragraph 3.3.3.2) require the IDE to attempt redelivery of a message up to three times within a 12-minute window when connectivity issues occur with a destination LRIT Data Centre. No equivalent redelivery requirement is currently prescribed for LRIT Data Centres themselves.

2.4.9 Earlier versions of the technical specifications included redelivery requirements applicable to the International LRIT Data Centre (IDC). However, when the IDC provisions were subsequently removed, this requirement was not carried forward. As a result, it is reasonable to assume that some LRIT Data Centre implementations may still incorporate redelivery mechanisms based on those earlier specifications, despite the absence of a current formal requirement; however, this is a significant issue that should be confirmed, since it would provide a resilience layer to ensure message delivery from DCs to IDE.

3 Questions to be answered

3.1 The analysis aims to answer the following questions:

- .1 What would be the increase in the message flow through the IDE if all Contracting Governments decide to exercise their coastal State entitlements?
- .2 Under the existing requirements, would the IDE be able to accommodate the increased message flow?
- .3 What would be the estimated number of incoming and outgoing messages DCs should expect if all Contracting Governments decide to exercise their coastal State entitlements?

4 Methodology

4.1 Data Sample

4.1.1 The analysis was conducted on sample data extracted from the 2025 LRIT audit schedule. Therefore, the experimental sample is formed by aggregating the audit sample data, and it covers:

- All 71 Data Centres;
- All LRIT information transmitted by ships indicated by the associated Administrations to transmit to each DC during its 30-day audit sample period;
- 39,517 ships registered to the DCs in their audit sample period;
- A total of 3,042,131 individual LRIT position standard reports (6/6 hrs) across

all DCs;

- The data sample period spans from July 2024 to August 2025.

4.1.2 The data sample was assessed for its suitability for generalisation over a one-year period, taking into account each LRIT DC, its respective sample period, and the number of registered ships. Several statistical non-parametric tests were performed to verify that the data are representative and that no bias was inadvertently introduced.

4.2 *Limitations and Caveats*

4.2.1 The following limitations should be noted when interpreting the results of this analysis:

- .1 The DDP as used in this analysis reflects the geographic entitlement zones as submitted by Contracting Governments, and incorrections in the polygons may affect the results; and
- .2 The analysis does not account for on-demand (polled) transmissions requested by individual governments or overreporting from individual ships. Only periodic messages in the standard rate of 6/6 hours were included.

4.3 *Treatment of non-standard-rate transmissions*

4.3.1 The analysis considered only the 6/6 hour regular transmission of LRIT information, since this is the concern this study aims to address.

4.3.2 Transmissions in rates faster than 6/6 hours or responses to polling requests have been excluded from the sample programmatically at runtime while loading the samples.

4.4 *Analysis Procedure*

4.4.1 Each of the 71 DCs 30-day samples was initially processed individually, and the results were logged in a database.

4.4.2 For each LRIT position in the 30-day sample data of a DC, the following procedure was applied:

- .1 Ship identification and associated Administration, DC, geographic coordinates (latitude, longitude) and Timestamp 4 were extracted from the record.
- .2 The version of the DDP in force at the moment of Timestamp 4 was retrieved from the DDP archive.
- .3 A point-in-polygon test was performed to determine all CG in whose 1,000 NM coastal zone polygon the ship's position fell at that instant.
- .4 CG associated with the same DC of the sample data are excluded since this information is routed internally, not using the IDE²
- .5 One database record was created for each such Contracting Government, representing one LRIT message that the DC would be expected to route to that

² Cooperative DC, Regional DC, and National DC providing services to other Administrations may route LRIT information internally, and only the journal is periodically sent to the IDE.

government's Data Centre.

- .6 The new registry contains originator ship IMO number, data user (Administration), originator DC, destination data user (coastal CG), destination DC, position, timestamp4;
- .7 Timestamp 4 was recorded for each generated record to enable temporal analysis of transmission peaks.

4.4.3 After all records were processed, they were adjusted for the same 30-day period, forming an aggregated 30-day sample that represents a period in which all ships are transmitting. Such data provide the basis for the yearly generalisation and facilitate the calculations performed during the data analysis.

4.4.4 The annualization further accounts for the current number of ships actively transmitting to each DC, as recorded in the DC's ship registry at the time of the audit, to allow scaling of results to the full active fleet.

4.5 *Operational and theoretical load under full coastal entitlement exercise*

4.5.1 While assessing the sample data, it was noted that 39,517 ships in the sample data, accounting for 3,042,131 pieces of standard 6/6 hours LRIT information sent to the DCs indicated by their Administrations. This is the operational data used in the analysis.

4.5.2 However, it was observed that for such a number of ships, a total of 4,742,040 individual pieces of LRIT information should be generated in the same period, if all ships produced one message every 6 hours as per the requirement. The sample analysis report describes the causes for this Fleet-wide Transmission Completeness Rate of 64.2%, noting this is a known phenomenon in the system.

4.5.3 In order to assess the load under full coastal entitlement exercise, this matter was considered, and the results are discussed for the operational scenario, but also for the theoretical one, where the Fleet-wide Transmission Completeness Rate is 100%.

5 **Aggregated data overview**

5.1 The total estimated messages under the full CSSO entitlement exercise for a 30-day period, based on the sample analysis and the current operational status of the LRIT, is 17,604,806. If considering the full theoretical capacity, with all registered ships transmitting 4 messages per day, the expected load is 27,467,703 messages per month.

5.2 To produce an estimate of the year-long IDE message flow, the total number of database records generated for the sample period was scaled to a full calendar year using the following formula:

$$\text{Annual Estimate}(\text{operational}) = (\text{Total Records}/\text{Sample Days}) \times 365 = 214,191,806 \text{ MT1}$$

5.3 Considering a scenario where all ships transmit every 6 hours and do not miss transmissions (theoretical capacity):

$$\text{Annual Estimate}(\text{max_theoretical}) = (\text{Total Records}/\text{Sample Days}) \times 365 = 334,190,386 \text{ MT1}$$


5.4 **Figure 1** presents two graphs, the first with the distribution of Timestamp 4 of periodic LRIT position report to be sent every second, for all DCs combined in a unique 30-period.

5.5 The second graph shows the moving average for 25 seconds in future, which gives an insight into the consequences of a 25-second timeout in the application layer. The moving average is calculated as follows:

$$MA(i) = (1/25) * \sum_{t=i}^{i+24} x(t)$$

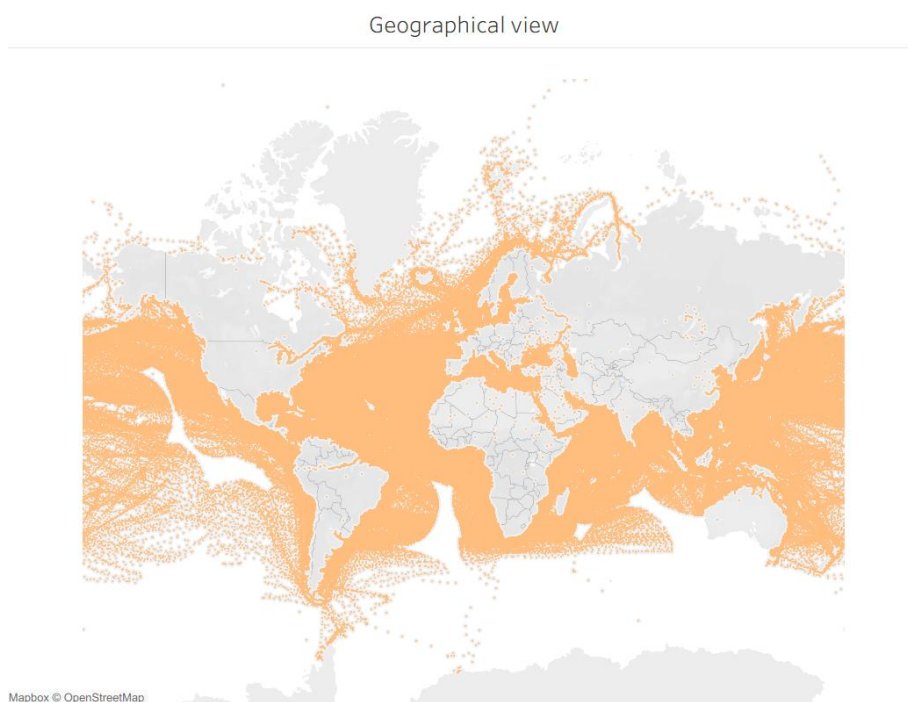
5.6 The peak observed in the moving average (nearly 200 messages) is not attributed to an outline in the ship's transmission. This phenomenon is historically associated with connection issues in the LRIT or CSP infrastructure (e.g., a CSP issue or a DC outage). When an issue occurs, the participant components may buffer the outgoing traffic and release it when possible, causing communication peaks.

Figure 1 - Estimated message flow through IDE per second.



5.7 A geo-referenced visualisation of estimated annual message traffic (**Figure 2**) was generated by aggregating LRIT data received by DCs from ships instructed by their Administrations to report to them. The resulting plot illustrates the volume and distribution of SOLAS ships' transmissions to the LRIT system over a full year, allowing coastal States to assess the information they would receive if they were to exercise their coastal entitlement.

Figure 2 - Geo-referenced LRIT positions from sample data.



6 Results

6.1 Average Load Estimation

6.1.1 The analysis of the operational data indicates that 214,191,806 messages are expected to be sent through the IDE in a one-year period if the CG exercises their coastal entitlement in full. This total represents an average of 6.8 messages per second. If the maximum theoretical capacity is reached, the average messages per second is estimated in 10.6. Considering the distributed load over time, in both scenarios the IDE will be operating well below its processing capacity of 100 messages per second.

6.1.2 The average increase in outgoing traffic of messages for the IDE is the same as the income, since the IDE receives the LRIT position reports and forwards them to the destination DCs.

6.2 Temporal Analysis — Transmission Peak Times

6.2.1 **Figure 3** presents the distribution of LRIT transmissions across one month, aggregated across all Data Centres in the sample. This analysis identifies peak and off-peak periods of IDE activity, which are relevant to dimensioning routing capacity and scheduling maintenance windows. A peak of messages is a second in which the IDE receives more than its processing capacity per second, 100 messages per second.

6.2.2 Considering distributed systems which communicate with asynchronous messages, such as the LRIT system, it is expected that in some situations communications will arrive faster than the system can process them. However, the system should be able to keep such communications, or sessions, alive while they have not been processed. As mentioned before, communication resilience may benefit from tuning the server connection pool and timeout mechanisms, and from implementing redelivery mechanisms.

6.2.3 Considering redelivering mechanism is not a requirement for DCs and should be further discussed, the analysis here intends to estimate the minimum connection pool and session timeout for the IDE to be able to couple with the peaks indicated by the sample data, considering its nominal capacity of processing 100 messages per second.

6.2.4 The analysis considers the composition of the DCs' data sample and the messages that would be generated by each piece of LRIT information. The 30-day aggregated sample data contain 3,042,131 periodic LRIT information sent by the ships to the DC indicated by their Administration. These messages cause 17,604,806 periodic messages to be sent to coastal states registered in other DCs, based on their 1000 NM coastal area. In short, due to the overlap of 1000 NM polygons, each piece of LRIT information from a ship would, on average, be sent to 5.7 other CG under their CSSO through the IDE if they exercise their coastal entitlement in full. Inter DC traffic is excluded because it does not affect the IDE.

6.2.5 Over the 30-day aggregated sample period, the message rate received by the IDE shows a clear central tendency of around 6.8 messages per second, representing the typical level of activity. Values up to 10 messages per second were registered in 75% of the time, indicating that the system will operate within this range under normal conditions. This suggests a stable and consistent workload for most of the time.

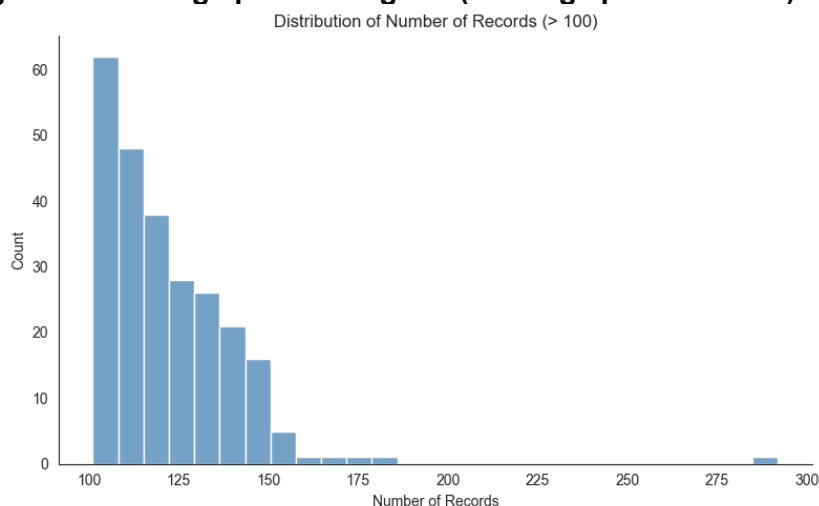
6.2.6 Statistically, outliers are identified using the interquartile range (IQR) method ($IQR = Q3 - Q1$; upper threshold = $Q3 + 1.5 \times IQR$). For the operational scenario, any value above 25 messages per second is classified as an outlier in the aggregated sample data. When extrapolated to the theoretical scenario, outliers correspond to values exceeding 40 messages per second. In the 30-day aggregated dataset, a total of 249 seconds recorded traffic above 100 messages per second, representing approximately 0.01% of the total operating time.

6.2.7 For the aggregate data sample, the maximum number of messages observed in a single second was 292, with the distribution showing a clear positive skew, as illustrated in the histogram. Only 10 seconds of recorded traffic exceeded 150 messages per second (Figure 3).

6.2.8 In a theoretical scenario, assuming continuous transmissions from all ships on a strict 6/6-hour cycle without interruption, traffic levels would increase by approximately 50%. However, analysis of the data does not identify any temporal pattern indicating concentration of messages at specific time intervals, nor any alignment with the 6/6-hour transmission cycle. No cyclic behaviour was observed.

6.2.9 Instead, the data exhibits short-term clustering. Further analysis indicates that the principal burst of messages within the sample originated from a single DC within a short interval of a few seconds. This behaviour is consistent with historical observations by the LRIT Coordinator during audits, which indicate that significant disturbances in message flow are typically associated with issues or outages affecting CSPs or LRIT system components, such as the IDE or DCs.

Figure 3 - Message peak histogram (Message per sec > 100).



6.2.10 To assess the nominal ability of the IDE to handle these peak times, the data was further analysed, calculating the minimum connection pool required for every second in the monthly aggregated data:

$$CPmin(t_0) = M(t_0)$$


$$CPmin(t) = \max(0, CPmin(t-1) - 100) + M(t)$$

M(t): number of messages received at time t (per second)
CPmin(t): minimum required connection pool size

6.2.11 This metric indicates the size of the connection pool required to handle the incoming traffic, considering the IDE capacity of handling 100 messages per second.

6.2.12 **Figure 4** shows that during a burst period, considered an outlier and non-related to the ship's behaviour, the number of connections needed in the pool reaches 292 on day 6.

Figure 4 - Message peak and connection pool per second (aggregated sample data)




6.2.13 The data was further analysed to identify the maximum time interval between seconds that did not require an extra buffer (i.e. when the sum of the remaining connections and the incoming ones is lower than 100 messages). The resulting maximum time interval was 19 seconds, providing indication of the minimum timeout value that should be used when configuring the LRIT system components.

6.2.14 Special consideration was given to the full theoretical scenario where all ships are able to transmit their periodic information with no gaps. Given that it is unrealistic to assume these new transmissions would occur at the exact seconds we observe communications in the actual sample, a model was developed to distribute the new load evenly across the data.

6.2.15 **Figure 5** presents the simulation of such traffic. Considering the 100 messages per second capacity of the IDE, the pending connections reached 329, indicating the minimum pool size and the biggest gap between the two values of 100 messages in the pool was 20 seconds, suggesting a potential number for the timeout at the application level.

Figure 5 - Message peak and connection pool per second (theoretical scenario).



6.2.16 The foregoing analysis is based on the assumption that DCs do not implement re-delivery mechanisms, which could otherwise mitigate pressure on the IDE. From an architectural standpoint, the implementation of such mechanisms is recommended, even in the current scenario where only limited periodic reports are exchanged via the IDE, as it would enhance resilience to short-duration surges and temporary outages lasting a few minutes.

6.3 Increase in storage for IDE

6.3.1 Taking into account a binary XML representation of the periodic message, removing parts not stored by the IDE, such as latitude and longitude, a single message reduces to nearly 1.3KB uncompressed.

6.3.2 The Message flow is expected to increase from 15M to 214M messages. Thus, in theory, the raw XML data received will increase from approximately 33GB to 500 GB. When representing XML in a relational database, there is overhead due to indexes and internal database metadata. This overhead will vary depending on the implementation approach and the database management system. On the other hand, if the data is parsed and only the values of the files are stored, the storage size can be a fraction of the original data.

6.3.3 As presented, it is not feasible to calculate the actual increment in disk space for the IDE databases to couple with the new influx of information. However, it is a fair estimation to consider that the annual increment in size for the data will increase 15 times when compared to the current scenario, since LRIT position reports correspond to 90% of the current traffic

6.3.4 If considering the theoretical maximum capacity, where all ships transmit every 6 hours with no gap, the increase in messages will be from 14M to 334M, and the data may increase 23 times from the current annual increment.

6.4 Operational impact for Data Centres

Outgoing traffic:

6.4.1 The aggregated sample data show the receipt of 3,042,131 position reports by CG as flag States, and that these reports may imply 17,604,806 messages being sent to CG in other DCs if they exercise full coastal State entitlement. This suggests that, on average, every message received by the DC will cause 5.7 messages to be sent to other DCs, on average, and these messages must be stored in the DC.

6.4.2 As an example, if a DC has 100 ships, this implies 146,000 annual standard reports (100x4x365) that the DC should be receiving. In a scenario with full exercise of coastal entitlement, these DC, on average, will send 832,200 (146,000 x 5.7) messages to other CGs.

6.4.3 Considering that the full MT1 has all timestamps and the position of the ship. We may consider 1.6KB per message for estimation. One year will require around 1.5GB of raw data space for 100 ships, and this may be reflected in a database volume several times larger (e.g., 3x), due to metadata and DBMS overhead if stored as XML, or it can be compacted if saved as values in database tables.

6.4.4 For instance, a DC with 5,000 ships may need to provision an extra 75GB of raw data and consider how much this will impact the database.

Incoming traffic

6.4.5 It is not possible to provide a generalised rule to estimate incoming traffic, because this is highly dependent on the traffic density in the 1000NM from the coastline of the CG associated with the DCs.

6.4.6 However, the aggregated simulation developed in this experiment allows the LRIT Coordinator to inform each DC individually of their estimated traffic if they exercise their full 1000 NM coastal entitlement. Since the size of MT1 is the same as described in the paragraph above, these DCs can estimate incoming traffic and the need for storage resources.

6.4.7 The LRIT Coordinator notes that there is no requirement enforcing CG to receive information in the 1000 NM region; thus, it is sensible to assume that CG will only activate their polygons when their DCs are able to confirm they have the capacity to handle the incoming traffic.

7 Conclusions

7.1 The present study used a statistical experiment based on sample data that considered one month of traffic from each DC and aggregated these months to estimate annual traffic on the LRIT in a scenario in which all CGs exercise the full 1000 NM coastal entitlement. The suitability of the sample for a 12-month generalisation was assured through tests in a pre-processing phase.

7.2 The statistical analysis outcome, considered together with software engineering concepts, standards and specifications, allowed for the estimation of the impact on the IDE and the DCs in terms of message flow and storage needs.

7.3 The LRIT Coordinator notes that the information provided is generic and does not consider the specifics of each implementation, which may cause variation in the results. The

aim was not to provide precise data, which would be impossible by the very nature of a statistical experiment, but to look at the dimensions of potential impacts in all CG requires information in their 1000 NM, assisting the ongoing discussion on this matter.

7.4 Based on the outcome of the study, the existing technical requirements and the current ship reporting performance, it is concluded that:

- .1 with regard to message processing capacity, the existing requirements for the IDE are sufficient to accommodate the increased message flow if all Contracting Governments decide to exercise their coastal State entitlements.
- .2 the peak of transmissions, which are usually caused by a shoreside component such as DC or CSPs, are short and infrequent, and can be managed through appropriate configuration of system parameters.
- .3 the projected increase in data volume may require appropriate scaling of storage capacity at both IDE and Data Centre level if needed. The study suggests that DCs would require a storage capacity of 1.5GB per 100 ships per year for the outgoing messages and the IDE around 500GB per year (when the LRIT messages are saved in raw XML format). These data storage capacity may already be available considering that it was assumed when the LRIT system was designed that Contracting Governments would actively monitor their coastal waters.

7.5 The study outcome includes data on the estimated increase in incoming and outgoing message flow for individual DCs, which is not published in this document. However, the DCs and Contracting Governments may consult the LRIT Coordinator for detailed information regarding their DCs.

7.6 The study confirms the importance of a redelivery mechanism for the IDE and the DCs, especially when there is a peak in the number of messages sent or received.
