# Thread Network Performance

## Test Report Version 1.0

ESPRESSIF

# About This Document

This report aims to evaluate the performance of the Thread network. The tests are conducted on Espressif's Thread SoCs and SDK. The test results highlight the network key metrics, including latency, throughput, packet loss rate, as well as large-scale capability and stability. The findings demonstrate that the Thread network consistently delivers stable performance across all scenarios, making it a reliable solution for IoT deployments.

## Release notes

| Date | Version | Release Notes |
| --- | --- | --- |
| Jun 2025 | V1.0 | Initial release |

## Documentation Change Notification

Espressif provides email notifications to keep customers updated on changes to technical documentation. Please subscribe at https://www.espressif.com/en/subscribe.

## Certification

Download certificates for Espressif products from https://www.espressif.com/en/certificates.

# Table of Contents

# 1. Overview

To evaluate the connectivity performance and scalability of the Thread network, we conducted tests under two different topologies:

- Large-scale mesh topology with 300 nodes: used to assess the network's capacity and stability under dense deployment condition.
- 10-hop linear topology: used to measure Thread network throughput, latency, and packet loss rate.

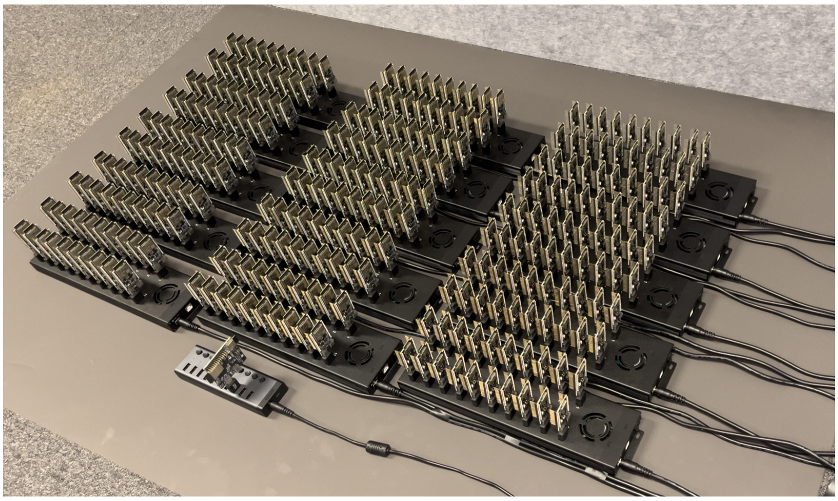

Figure 1-1. Large-scale Mesh Topology

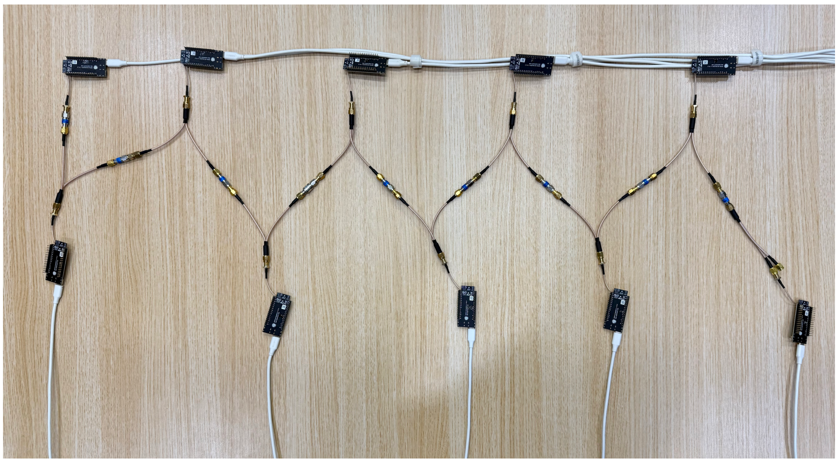

Figure 1-2. 10-hop Linear Topology

**Note:**

*To ensure a stable and reproducible testing environment, the test nodes in the 10-hop linear topology are connected via attenuator, with MAC filtering enabled to control network topology. This configuration ensures that each test node receives only one-hop or two-hop messages, as illustrated in figure 1-2.*

*All the tests are conducted in a shielding box.*

# 2. Performance Summary

This section provides a concise summary of key test results, offering a quick overview of the Thread network performance.

## 2.1. Network Capability

In our evaluation, a 40-node network consisting of 5 FTDs and 35 MTDs formed and stabilized within 30 seconds. In a larger 300-node mesh network with 20 FTDs and 280 MTDs, stabilization was achieved within 2 minutes. Both networks maintained reliable connectivity and communication throughout the evaluation period.

The Thread network demonstrates strong stability, even under large-scale deployment.

## 2.2. Connectivity Performance

In a 10-hop linear Thread network topology, performance tests were conducted using iPerf to measure TCP and UDP throughput. Additionally, the Ping command with varying payload lengths was used to evaluate latency and packet loss for both unicast and multicast traffic.

The result summary is provided below, with additional details available in Chapter 3.

Table 2-1. Test Results Summary

| Type | | 1 Hop | 2 Hop | 3 Hop | 4 Hop | 5 Hop | 6 Hop | 7 Hop | 8 Hop | 9 Hop |
|---|---|---|---|---|---|---|---|---|---|---|
| TCP Throughput (Kbps) | | 101 | 39 | 25 | 23 | 22 | 22 | 20 | 19 | 19 |
| UDP Throughput (Kbps) | | 103 | 41 | 27 | 27 | 26 | 26 | 26 | 26 | 26 |
| Type | Size (byte) | Round-Trip Latency (ms) | | | | | | | | |
| Unicast | 10 | 14 | 26 | 38 | 50 | 62 | 74 | 86 | 98 | 110 |
| | 70 | 18 | 35 | 51 | 68 | 84 | 100 | 117 | 133 | 149 |
| | 650 | 101 | 250 | 349 | 399 | 419 | 438 | 458 | 476 | 494 |
| | 1200 | 181 | 443 | 634 | 715 | 736 | 757 | 776 | 798 | 818 |
| Multicast | 10 | 15 | 60 | 107 | 154 | 200 | 247 | 293 | 340 | 387 |
| | 70 | 22 | 76 | 133 | 190 | 247 | 302 | 357 | 413 | 469 |
| | 650 | 137 | 352 | 631 | 951 | 1249 | 1548 | 1844 | 2145 | 2442 |
| | 1000 | 206 | 525 | 941 | 1423 | 1869 | 2325 | 2775 | 3231 | 3686 |
| Type | Size (byte) | Loss Rate (%) | | | | | | | | |
| Unicast | 10 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | 70 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | 650 | 0 | 0 | 0 | 0.04 | 0.06 | 0.12 | 0.13 | 0.15 | 0.17 |
| | 1200 | 0 | 0 | 0 | 0.09 | 0.12 | 0.22 | 0.31 | 0.35 | 0.44 |
| Multicast | 10 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | 70 | 0 | 0 | 0 | 0 | 0 | 0 | 0.01 | 0.02 | 0.12 |
| | 650 | 0 | 0.05 | 0.11 | 0.13 | 0.25 | 0.24 | 0.8 | 1.14 | 2.5 |
| | 1000 | 0 | 0.05 | 0.12 | 0.19 | 0.33 | 0.47 | 1.36 | 2.04 | 4.12 |

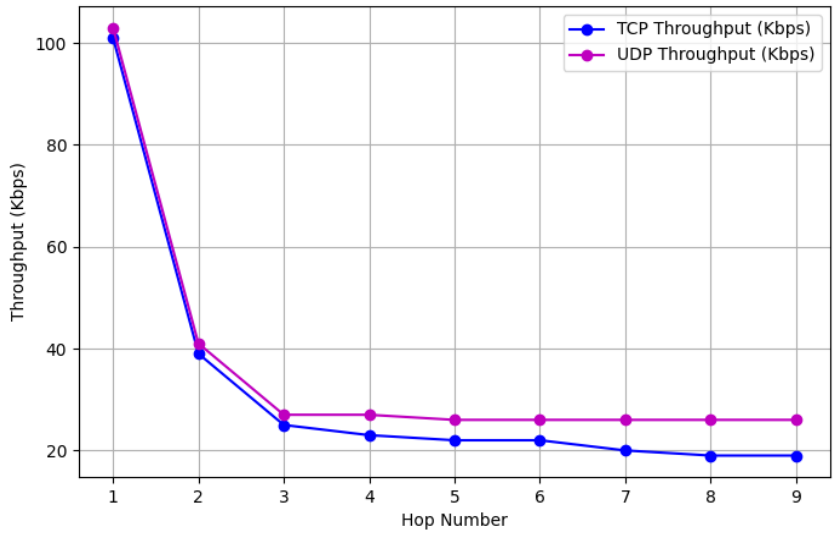


Figure 2-1. TCP & UDP Throughput vs. Hop Number

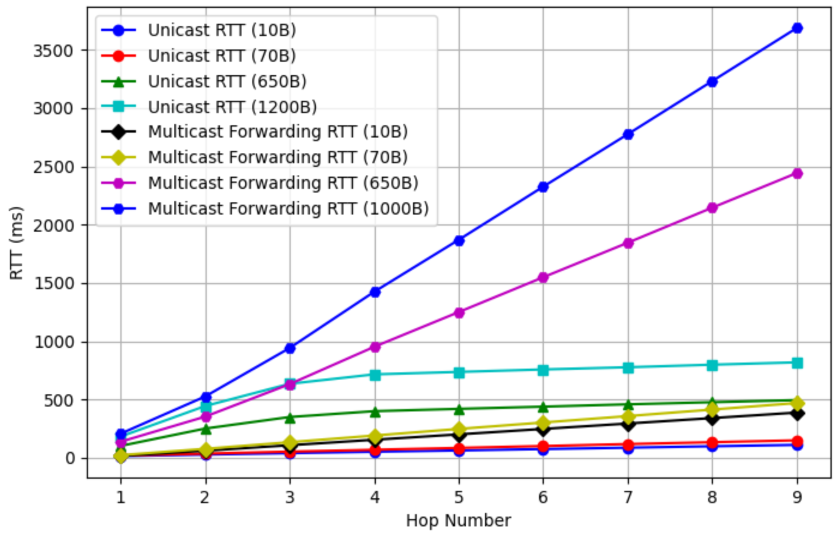


Figure 2-2. RTT vs. Hop Number

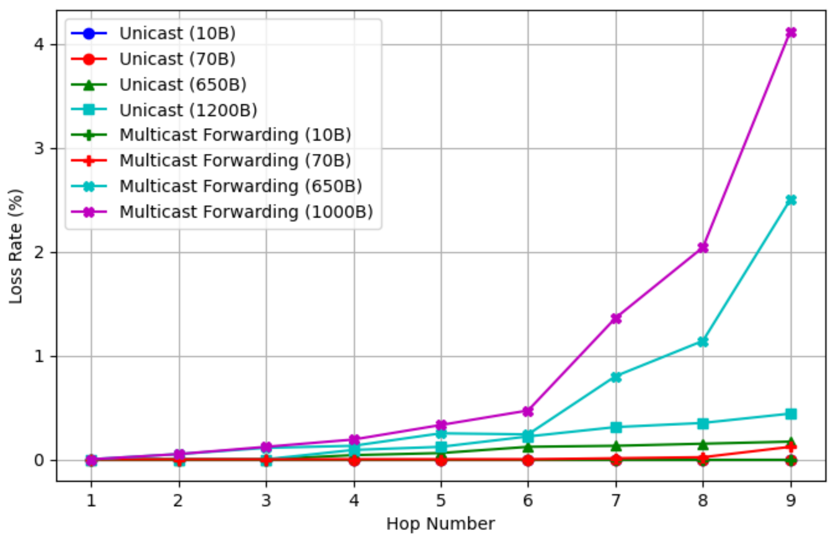


Figure 2-3. Packet Loss Rate vs. Hop Number

# 3. Detailed Test Results

## 3.1. Network Formation Time and Stability

A comprehensive analysis was conducted to evaluate the network's formation time and stability under varying conditions.

In a 40-node mesh network consisting of 5 FTDs and 35 MTDs, after power-up, devices quickly established the network, reaching a stable state within 30 seconds.

In a large-scale 300-node mesh network with 20 FTDs and 280 MTDs, the entire network achieved full operational status within 2 minutes, despite the significantly higher node density.

To validate inter-node communication, random devices were selected for ping tests. These selected devices consistently exhibited seamless communication, underscoring the robustness and stability of the network under all tested conditions. Continuous monitoring throughout the test confirmed that all nodes successfully joined the network and maintained stable, reliable connections, with no observed disruptions.

These findings demonstrate that the network can integrate rapidly and maintain reliable operation, making it well-suited for professional applications which require stable and scalable connectivity.

The demo video is available on YouTube at the following link:

Espressif Large-Scale Thread Network Performance Demonstration

## 3.2. Thread Border Router Capability

The test was conducted using Espressif's Thread Border Router (BR) solution, based on a hardware platform combining the ESP32-S3 and ESP32-H2. The ESP32-S3 is integrated with 2 MB of external PSRAM to accommodate larger volumes of network data. It's tested in a 300-node Thread mesh network, and evaluate the bi-directional connectivity, service discovery, and NAT64 performance.

### 3.2.1. Service Registration and Discovery Performance

In the 300-node Thread network, each node registered a DNS service with the BR, with each service approximately 1 KB in size. The BR successfully handled all registrations without failure. All services were discovered from backbone network as expected, demonstrating stable performance under high load.

These results confirm that the Thread network can efficiently support large-scale service registration and discovery. The test also validates the BR's capability to reliably manage simultaneous service registrations from a 300-node network.

### 3.2.2. NAT64 Performance

The BR's NAT64 (IPv6-to-IPv4 translation) feature enables Thread devices in the network to communicate with IPv4-based devices on the internet, particularly cloud services.

During testing, the NAT64 mechanism enabled seamless IPv6-to-IPv4 communication across all nodes. Each node in the 300-node Thread network successfully established a TLS session with the cloud. On the BR, each session consumed only 44 bytes of memory, and the results showed stable network performance, with zero application packet loss, consistent throughput, and minimal memory overhead.

These results highlight the efficiency of NAT64 in supporting large-scale deployments within a Thread network.

It is worth noting that the overall data rate through the BR is constrained by the IEEE 802.15.4 physical layer, which has a maximum throughput of 250 Kbps. As a result, the BR’s concurrent data throughput is inherently limited by this specification.

## 3.3. Unicast Performance

This section evaluates the unicast performance of the Thread network under varying payload sizes, focusing on throughput, latency, and packet loss rate.

### 3.3.1. Throughput

TCP and UDP throughput were evaluated across multiple hops using iPerf, revealing a significant performance drop after the first hop, followed by a more gradual decline as the hop count increased.

For TCP, the highest throughput was observed at 1-hop (101 Kbps), but it dropped sharply to 39 Kbps at 2 hops. Beyond 3 hops, the throughput stabilized at a lower rate, gradually declining to 22 Kbps at 5 hops. These results indicate that multi-hop transmission significantly impacts TCP performance, with most of the degradation occurring within the first few hops.

For UDP, a similar trend was observed. Throughput started at 103 Kbps at 1-hop and dropped significantly to 41 Kbps at 2 hops. It continued to decline beyond 3 hops, stabilizing around 26 Kbps from 5 hops onward. This suggests that UDP performance also degrades primarily in the early hops and remains relatively stable over longer distances.

The throughput degradation observed in the initial hops is primarily due to increased forwarding delays. Intermediate nodes must receive and retransmit each frame, incurring additional time and bandwidth overhead.

The results are summarized in table 3-1.

Table 3-1. Throughput Test Results (Unit: Kbps)

| Type | 1 Hop | 2 Hop | 3 Hop | 4 Hop | 5 Hop | 6 Hop | 7 Hop | 8 Hop | 9 Hop |
|---|---|---|---|---|---|---|---|---|---|
| TCP | 101 | 39 | 25 | 23 | 22 | 22 | 20 | 19 | 19 |
| UDP | 103 | 41 | 27 | 27 | 26 | 26 | 26 | 26 | 26 |

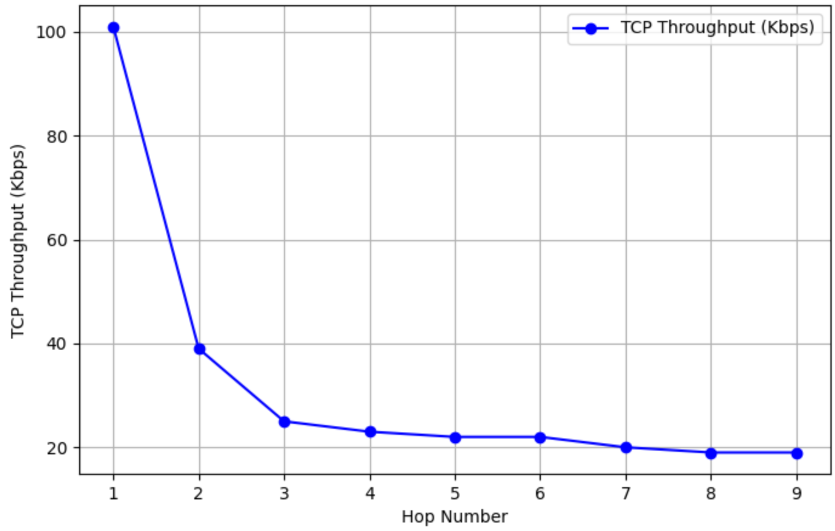


Figure 3-1. TCP Throughput vs. Hop Number

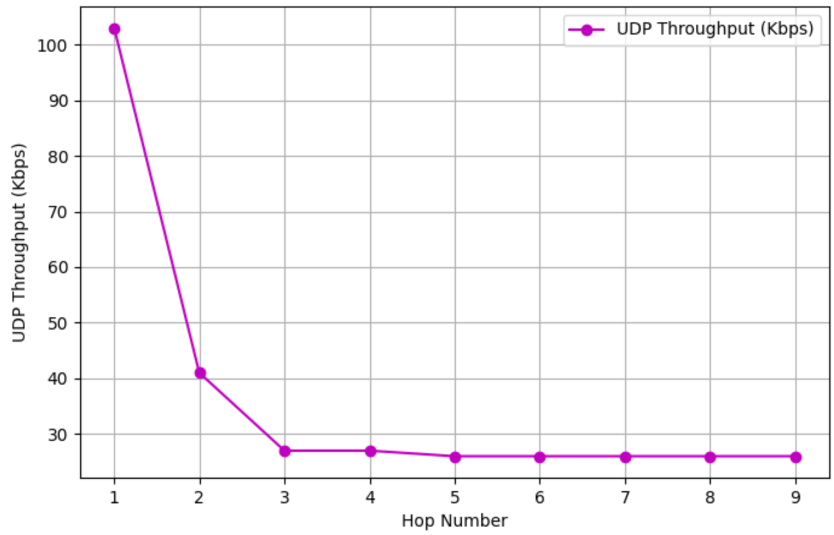


Figure 3-2. UDP Throughput vs. Hop Number

### 3.3.2. Round-Trip Latency

The Ping round-trip time (RTT) test was conducted using payload sizes of 10, 70, 650, and 1200 bytes across multiple hops. The test measured the minimum (Min), average (Avg), and maximum (Max) RTT for each hop.

#### Observations

- Small Payload Sizes (10 & 70 bytes)
  - For 10-byte payload, RTT increased gradually with hop count, with average RTT ranging from 14 ms (1 hop) to 110 ms (9 hops).

  - For 70-byte payload, RTT values were slightly higher, with an average RTT increasing from 18 ms (1 hop) to 149 ms (9 hops).
- Medium Payload Size (650 bytes)
  - A notable increase in RTT was observed, with minimum RTT ranging from 90 ms (1 hop) to 448 ms (9 hops).
  - The average RTT rose from 101 ms (1 hop) to 494ms (9 hops), indicating a steady increase due to larger packet transmission time.
- Large Payload Size (1200 bytes)
  - This packet size yielded the highest RTT values, with minimum RTT starting at 166 ms (1 hop) and reaching 760 ms (9 hops).
  - Average RTT ranged from 181 ms (1 hop) to 818 ms (9 hops).

### Conclusions

- RTT increases predictably with hop count due to cumulative forwarding delays.
- Larger payloads introduce significantly higher RTT.
- RTT variation (Max–Min) is more prominent with larger packet sizes, indicating possible network congestion or queuing effects at intermediate hops.

Table 3-2. Unicast Round-Trip Latency Test Results (Unit: ms)

| Size (byte) | Type | 1 Hop | 2 Hop | 3 Hop | 4 Hop | 5 Hop | 6 Hop | 7 Hop | 8 Hop | 9 Hop |
|---|---|---|---|---|---|---|---|---|---|---|
| 10 | Min | 11 | 21 | 32 | 41 | 53 | 64 | 75 | 86 | 95 |
| | Avg | 14 | 26 | 38 | 50 | 62 | 74 | 86 | 98 | 110 |
| | Max | 28 | 40 | 101 | 70 | 97 | 116 | 108 | 124 | 142 |
| 70 | Min | 15 | 30 | 45 | 60 | 76 | 91 | 105 | 122 | 135 |
| | Avg | 18 | 35 | 51 | 68 | 84 | 100 | 117 | 133 | 149 |
| | Max | 35 | 66 | 68 | 88 | 150 | 124 | 202 | 160 | 228 |
| 650 | Min | 90 | 235 | 299 | 365 | 382 | 398 | 418 | 426 | 448 |
| | Avg | 101 | 250 | 349 | 399 | 419 | 438 | 458 | 476 | 494 |
| | Max | 144 | 295 | 415 | 616 | 655 | 685 | 662 | 699 | 718 |
| 1200 | Min | 166 | 423 | 564 | 666 | 687 | 695 | 722 | 741 | 760 |
| | Avg | 181 | 443 | 634 | 715 | 736 | 757 | 776 | 798 | 818 |
| | Max | 205 | 490 | 849 | 877 | 946 | 1051 | 991 | 1017 | 1069 |

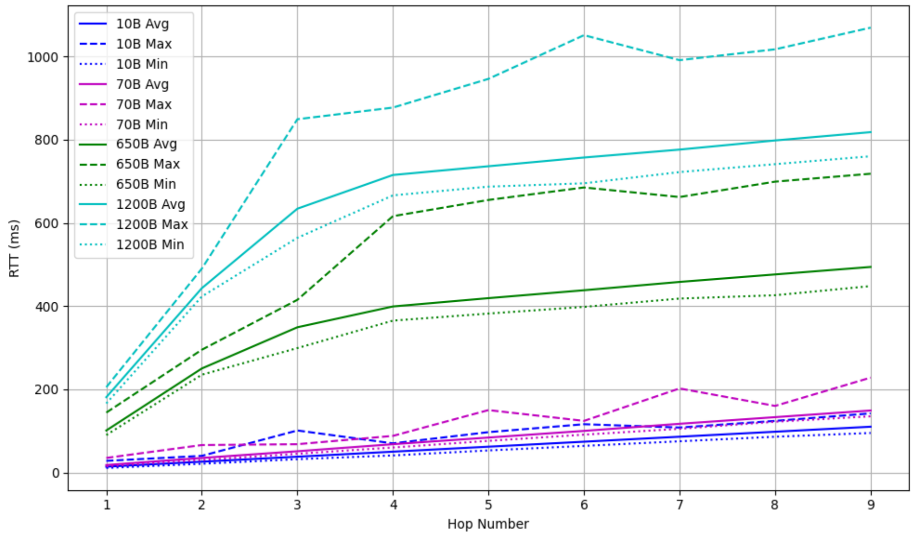


Figure 3-3. Unicast RTT (Min/Avg/Max) vs. Hop Number

### 3.3.3. Packet Loss Rate

To evaluate network reliability across different payload sizes, a multi-hop ping test was conducted. The test measured packet loss for four packet sizes: 10 bytes, 70 bytes, 650 bytes, and 1200 bytes.

- For 10-byte and 70-byte payloads, the ping interval was set to 1 second. Results showed 0% packet loss across all hops, indicating stable communication and efficient packet transmission.
- For 650-byte and 1200-byte payloads, the ping interval was increased to 3 seconds to accommodate the larger payloads.
  - For 650-byte payload, no loss occurred within the first three hops. Minor loss appeared beyond that, peaking at 0.17% at hop 9.
  - For 1200-byte payload, the first three hops also showed no loss. Packet loss began at hop 4 (0.09%) and increased with distance, reaching a peak of 0.44% at hop 9.

These results confirm that the network maintains high reliability, with only minimal packet loss observed at larger packet sizes and over longer transmission distances.

Table 3-3. Unicast Packet Loss Rate Test Results (Unit: %)

| Size (byte) | 1 Hop | 2 Hop | 3 Hop | 4 Hop | 5 Hop | 6 Hop | 7 Hop | 8 Hop | 9 Hop |
|---|---|---|---|---|---|---|---|---|---|
| 10 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 70 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 650 | 0 | 0 | 0 | 0.04 | 0.06 | 0.12 | 0.13 | 0.15 | 0.17 |
| 1200 | 0 | 0 | 0 | 0.09 | 0.12 | 0.22 | 0.31 | 0.35 | 0.44 |

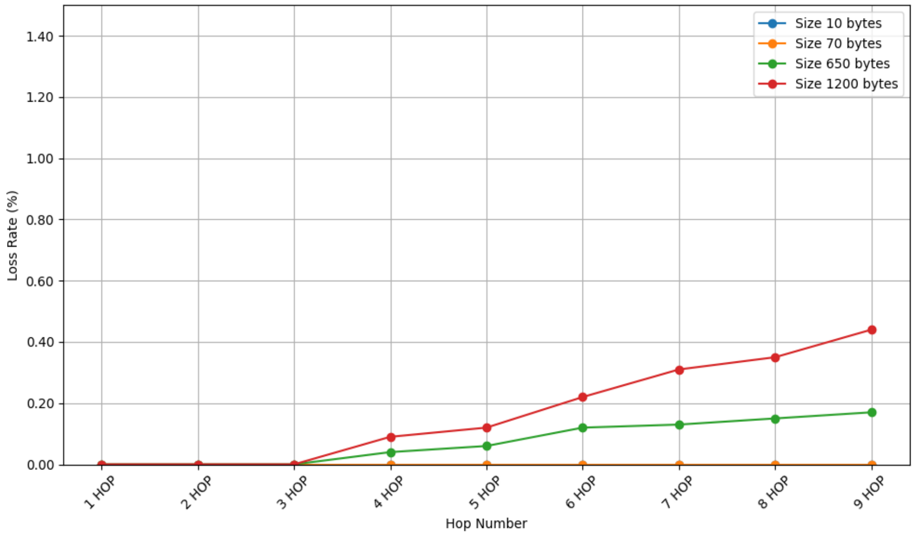


Figure 3-4. Unicast Loss Rate vs. Hop Number

## 3.4. Multicast Performance

This section evaluates the multicast performance of the Thread network using the Ping command, focusing on the latency and packet loss rate for multicast traffic.

All nodes were added to a multicast group, and the multicast address was pinged from the first-hop node. Upon receiving the multicast request, each node—regardless of its hop distance—sent a unicast response back to the first-hop node. The response time of each node was then measured accordingly.

### 3.4.1. Round-Trip Latency

To assess the transmission latency of multicast messages, RTT measurements were taken for four packet sizes (10 bytes, 70 bytes, 650 bytes, and 1000 bytes) across up to 9 hops.

#### Observations

- Small Payload Sizes (10 & 70 bytes)
  - For 10-byte payload, RTT increased steadily with hop count, from an average of 15 ms (1 hop) to 387ms (9 hops).
  - For 70-byte payload, average RTT ranged from 22 ms (1 hop) to 469 ms (9 hops), slightly higher than 10-byte packets.
  - At 9 hops, the maximum RTT reached 643 ms for 10-byte and 667 ms for 70-byte packets, reflecting moderate variance due to network conditions.
- Medium Payload Size (650 bytes)
  - RTT increased significantly with hop count, with minimum values rising from 98 ms (1 hop) to 1979 ms (9 hops).
  - Average RTT rose from 137 ms (1 hop) to 2442 ms (9 hops), showing the pronounced impact of larger payload sizes.

- Maximum RTT also rose sharply, with the 9-hop case peaking at 2742 ms.
- Large Payload Size (1000 bytes)
  - RTT values for 1000-byte payload were the highest among all sizes. Minimum RTT grew from 167 ms (1 hop) to 3363 ms (9 hops).
  - Average RTT increased from 206 ms to 3686 ms, reflecting considerable latency introduced by both payload size and hop count.
  - Maximum RTT fluctuated significantly, with the highest recorded value reaching 4209 ms at 9 hops, suggesting potential queuing delays or retransmission penalties in deeper hop ranges.

### Conclusions

- RTT increases proportionally with hop count in multicast scenarios, due to cumulative forwarding and processing delays.
- Larger packet sizes lead to significantly higher RTT values, where delays grow sharply due to increased transmission time and buffering.
- RTT variation (Max–Min) is more prominent with larger, indicating potential network congestion, queuing, or retransmission effects at intermediate hops.

Table 3-4. Multicast Round-Trip Latency Test Results (Unit: ms)

| Size (byte) | Type | 1 Hop | 2 Hop | 3 Hop | 4 Hop | 5 Hop | 6 Hop | 7 Hop | 8 Hop | 9 Hop |
|---|---|---|---|---|---|---|---|---|---|---|
| 10 | Min | 12 | 28 | 55 | 83 | 113 | 141 | 173 | 195 | 243 |
| | Avg | 15 | 60 | 107 | 154 | 200 | 247 | 293 | 340 | 387 |
| | Max | 70 | 154 | 220 | 285 | 341 | 402 | 449 | 501 | 643 |
| 70 | Min | 18 | 43 | 78 | 118 | 155 | 195 | 240 | 280 | 323 |
| | Avg | 22 | 76 | 133 | 190 | 247 | 302 | 357 | 413 | 469 |
| | Max | 71 | 164 | 252 | 314 | 395 | 461 | 522 | 602 | 667 |
| 650 | Min | 98 | 289 | 529 | 719 | 1002 | 1247 | 1555 | 1844 | 1979 |
| | Avg | 137 | 352 | 631 | 951 | 1249 | 1548 | 844 | 2145 | 2442 |
| | Max | 226 | 650 | 875 | 1258 | 1571 | 1952 | 2322 | 2596 | 2742 |
| 1000 | Min | 167 | 434 | 818 | 1245 | 1669 | 2086 | 2532 | 2842 | 3363 |
| | Avg | 206 | 525 | 941 | 1423 | 1869 | 2325 | 2775 | 3231 | 3686 |
| | Max | 313 | 712 | 1179 | 1811 | 2175 | 2627 | 3118 | 3750 | 4209 |

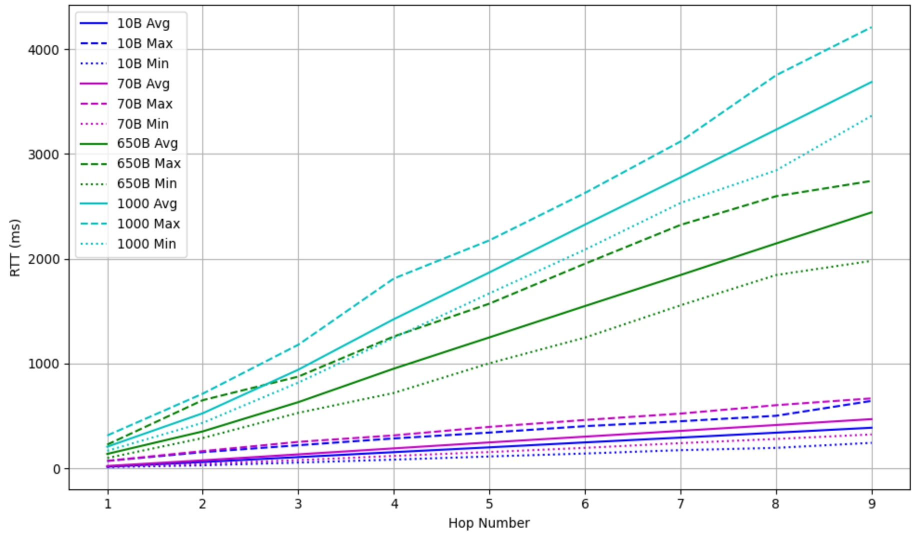


Figure 3-5. Multicast RTT (Min/Avg/Max) vs. Hop Number

### 3.4.2. Packet Loss Rate

To evaluate the reliability of multicast across multiple hops, a ping loss rate test was conducted using the same 10-hop linear topology as in the unicast tests. Four packet sizes were tested: 10 bytes, 70 bytes, 650 bytes, and 1000 bytes—with fixed ping intervals adjusted for each payload size.

- For 10-byte and 70-byte payloads, multicast delivery remained robust across all hops, with no packet loss observed for 10-byte packets and only minimal loss for 70-byte packets.
- For 650-byte packets, packet loss appeared at hop 2 (0.05%) and gradually increased, reaching 2.5% at hop 9.
- For 1000-byte packets, packet loss was first observed at hop 2 (0.05%), increasing to a peak of 4.12% at hop 9.

Overall, the multicast forwarding mechanism exhibited high reliability for small packets and maintained acceptable loss levels for larger payloads across extended hop counts.

Table 3-5. Multicast Packet Loss Rate Test Results (Unit: %)

| Size (byte) | 1 Hop | 2 Hop | 3 Hop | 4 Hop | 5 Hop | 6 Hop | 7 Hop | 8 Hop | 9 Hop |
|---|---|---|---|---|---|---|---|---|---|
| 10 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 70 | 0 | 0 | 0 | 0 | 0 | 0 | 0.01 | 0.02 | 0.12 |
| 650 | 0 | 0.05 | 0.11 | 0.13 | 0.25 | 0.24 | 0.8 | 1.14 | 2.5 |
| 1000 | 0 | 0.05 | 0.12 | 0.19 | 0.33 | 0.47 | 1.36 | 2.04 | 4.12 |

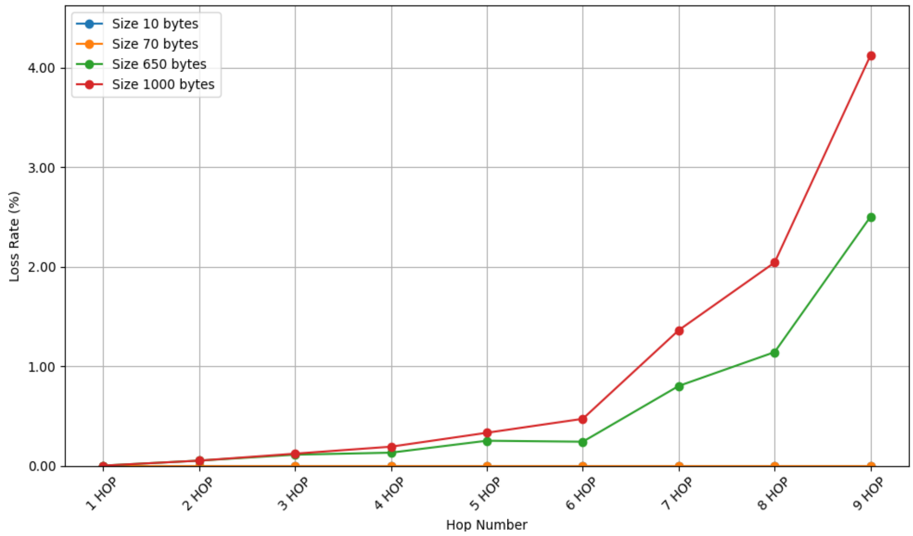


Figure 3-6. Packet Loss Rate vs. Hop Number

## 3.5. Communication Range

To assess the communication range of Thread devices, the tests were conducted in an open park area.

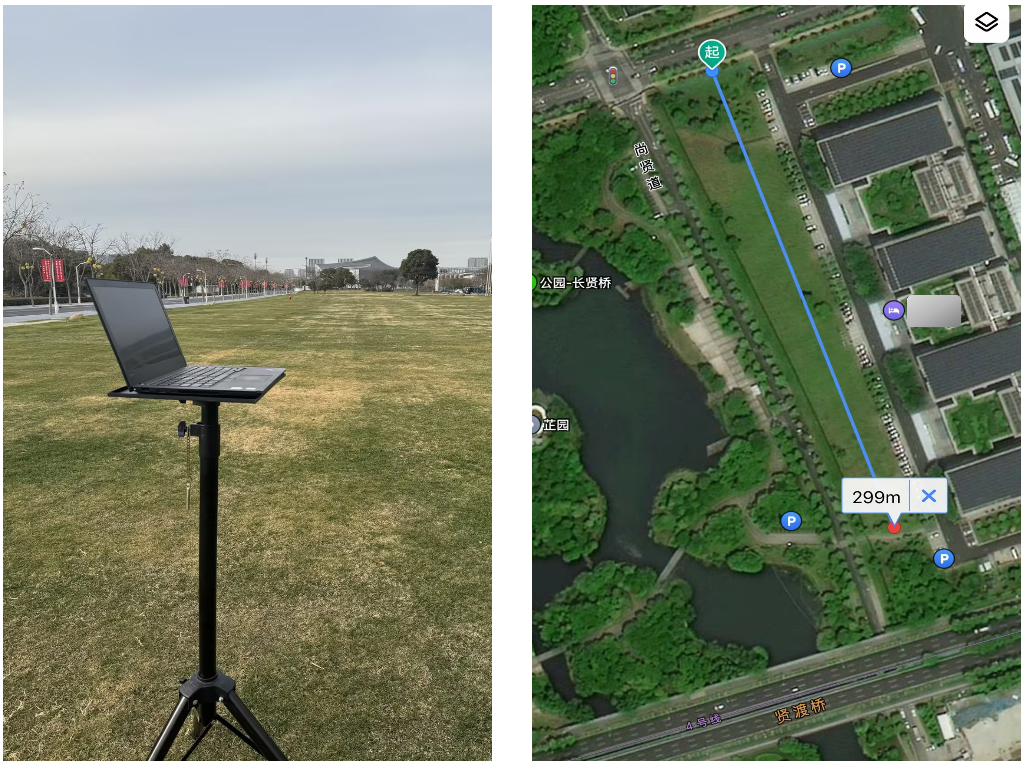


Figure 3-7. Communication Range Assessment

With both devices positioned 1.5 meters above the ground and transmitting at +20 dBm, stable communication was maintained over a distance of 250 meters.

Table 3-6. Test Results

| Range (m) | Join Network | Ping Loss Rate (%) |
| --- | --- | --- |
| 50 | Succeed | 0 |
| 100 | Succeed | 0 |
| 150 | Succeed | 0 |
| 200 | Succeed | 0 |
| 250 | Succeed | 0 |
| 300 | Succeed | 28 |

# Appendix A.

The Hardware and Software platforms used in the tests:

Hardware

- Espressif Thread Border Router
- ESP32-H2
- ESP32-C6

Software

- ESP-IDF : v5.2.5
- Thread FTD and MTD: ot_cli
- Thread Border Router: esp-thread-br
- iPerf Tool: espressif/iperf
- Some key software configurations see table A-1.

Table A-1. Test Results

| Test Configuration | Value |
|---|---|
| FREERTOS_HZ | 1000 |
| IEEE802154_TIMING_OPTIMIZATION | Y |
| LWIP_IRAM_OPTIMIZATION | Y |
| LWIP_EXTRA_IRAM_OPTIMIZATION | Y |
| OPENTHREAD_NUM_MESSAGE_BUFFERS | 1024 |
| OPENTHREAD_SPINEL_RX_FRAME_BUFFER_SIZE | 8192 |
| OPENTHREAD_MLE_MAX_CHILDREN | 30 |
| OPENTHREAD_CONFIG_MLE_ATTACH_BACKOFF_MAXIMUM_INTERVAL | 5000 |
| MDNS_MAX_SERVICES | 500 |