

TCP Connection Metrics in Metric Browser

On this page:

- Viewing TCP Flow Metrics
- TCP Flow Metric Descriptions
 - # Client Limited
 - # Client Zero Window
 - # Connection Errors
 - # Connection Requests
 - # Current Established Connections
 - # Data Retransmits
 - # Delayed Acks (Data Piggy Back)
 - # Delayed Acks (Timeouts)
 - # Errors
 - # Fin Errors
 - # Flows (<1KB)
 - # Flows (1k - 10k)
 - # Flows (10k - 100k)
 - # Flows (100k-1MB)
 - # Flows (1MB - 10MB)
 - # Flows (>10MB)
 - # Flows - Handshake (1SD)
 - # Flows - Handshake (2SD)
 - # Flows - Handshake (3SD)
 - # Flows - Lifetime (1SD)
 - # Flows - Lifetime (2SD)
 - # Flows - Lifetime (3SD)
 - # Flows - RTT (1SD)
 - # Flows - RTT (2SD)
 - # Flows - RTT (3SD)
 - # Flows - TCP Data Rxmt
 - # Flows - TCP Resets
 - # Flows - w/o TCP SACK
 - # Flows - w/o TCP Timestamp
 - # IP Fragment Count
 - # Loss Pkts
 - # Nagle Delays
 - # PIE Events
 - # RetransmissionTimeouts
 - # SACK Retransmits
 - # Server Limited
 - # Server Zero Window
 - # Syn Blackholes
 - # Syn Resets
 - # TCP Resets - Established
 - # TCP Resets - Fin
 - # TTL Changes (1 - 2 hops)
 - # TTL Changes (3 - 4 hops)
 - # TTL Changes (>=5 hops)
 - Avg # TTL Hops
 - Delayed Acks Lag (usec)
 - Initial RTT (usec)

- Latency - RTT (usec)
- Lifetime (usec)
- Loss (per mille)
- Nagle Delays Lag (usec)
- Pkts per Sec
- Rx Pkts per Sec
- Rx Throughput (BPS)
- TCP Handshake (usec)
- Throughput (BPS)
- Tx Pkts per Sec
- Tx Throughput (BPS)

Related pages:

- [Network Visibility Metrics](#)
- [KPI Metrics in Network Dashboard and Application Flow Map](#)
- [KPI Metrics in Right-Click Dashboards](#)
- [TCP Connection Metrics in Metric Browser](#)
- [Node Metrics in Network Dashboard and Metric Browser](#)
- [Dynamic Monitoring Mode and Network Visibility](#)

Network Visibility agents can calculate an extensive set of metrics based on the TCP flows observed by the network agents. In addition to KPI, PIE, and Troubleshooting metrics, you can view advanced metrics for network elements of interest (tiers, nodes, links, and Connections) in the Metric Browser. Using time-based charts, you can detect and analyze behavior such as

- TCP flow setup/teardown times and errors
- Distribution of TCP flows for an element by throughput, setup time, lifetime, and round-trip time
- **Flows closed by TCP resets**
- Distribution of flows based on TCP configuration options (Selective Acknowledgement, Timestamp)

Note the following:

- For these metrics, the Metric Browser uses the terms *Connection* and *Flow* interchangeably. See [Flows, Links, and Connections](#).
- Network Agents do not collect metrics for individual Connections by default. The recommended workflow is to enable TCP Flow metric collection only when you need to diagnose a performance issue on an associated node or Connection. See [Dynamic Monitoring Mode and Network Visibility](#).

Viewing TCP Flow Metrics

To view advanced network metrics, open the Metric Browser and navigate to Application Infrastructure Performance. The Metric Browser shows connection/flow metrics for the following aggregations:

- Application Infrastructure Performance >
 - Advanced Network >
 - Flows >
 - **All Connections/Flows for Application**
 - *tier-name*
 - Advanced Network >
 - Flows >
 - Call from <tier> to <service>
 - Call from <node> to <node-or-service>
 - **Individual Connection/Flow (<node> to <service>)**
 - **All Connections/Flows for Network Link (<tier> to <service>)**
 - **All Connections/Flows for tier**

TCP Flow Metric Descriptions

Metric Name	Description / Notes	Default Monitoring Mode
# Client Limited	The number of TCP window updates sent by a client node indicating that it is receiving data at a faster rate than the application can process.	KPI
# Client Zero Window	<p>If a client's TCP buffer is full, it sends Zero as the TCP receive window size to indicate that it cannot receive more data. The data transfer then stops until the client can process the data in its buffer.</p> <p>A high rate of Client TCP Zero Window messages indicates a problem with either</p> <ul style="list-style-type: none"> • The TCP configuration on the node, or • The client-side application using that flow. 	KPI
# Connection Errors	The sum of Syn Resets + Syn Blackholes + TCP Resets - Established .	KPI
# Connection Requests	The total number of connection requests (successful and unsuccessful) sent during the selected time window.	KPI
# Current Established Connections	The number of established (set-up) flows/connections.	KPI
# Data Retransmits	Number of TCP data packets that were retransmitted for all TCP Connections.	KPI
# Delayed Acks (Data Piggy Back)	The average number of times a receiving node sent an ACK (acknowledgement) by "piggybacking" the ACK onto another message.	Advanced
# Delayed Acks (Timeouts)	<p>The average number of times a receiving node sent an ACK (acknowledgement) because its Delayed ACK timer expired.</p> <p>This is a "worst-case" delay for the Delayed ACK algorithm, and occurs most often when Nagle's algorithm and Delayed ACK are enabled on the sending and receiving node (respectively). A high rate of delayed ACKs can contribute significantly to the average Latency on a Connection.</p>	Advanced
# Errors	The number of TCP messages sent indicating an error in setting up the connection (SYN errors) or tearing down the connection (FIN errors).	KPI
# Fin Errors	The number of errors seen while tearing down the connections (TCP FIN errors). A lot of connections in FIN wait states can cause delays in creating new connections.	KPI
# Flows (<1KB) # Flows (1k - 10k) # Flows (10k - 100k) # Flows (100k-1MB) # Flows (1MB - 10MB) # Flows (>10MB)	Use these metrics to analyze the distribution of flows by throughput.	Advanced

<p># Flows - Handshake (1SD)</p> <p># Flows - Handshake (2SD)</p> <p># Flows - Handshake (3SD)</p>	<p>The number of TCP flows with connection-setup times that are 1, 2, or 3 Standard Deviations outside (higher or lower than) the average for that connection group.</p> <p>Under ideal conditions, all flows should be within 1SD of the average.</p>	Advanced
<p># Flows - Lifetime (1SD)</p> <p># Flows - Lifetime (2SD)</p> <p># Flows - Lifetime (3SD)</p>	<p>The number of TCP flows with lifetimes that are 1, 2, or 3 Standard Deviations outside (higher or lower than) the average.</p> <p>Under ideal conditions, all flows should be within 1SD of the average. 2SD or 3SD flows indicate inconsistent flow treatment along the network path. These inconsistencies can occur when a network service decreases the available bandwidth intentionally (<i>bandwidth throttling</i>) or prioritizes some types of traffic over others (<i>traffic shaping</i>).</p> <p>Short-lived connections, even if they are intermittent, can indicate an issue worth investigating. The more short-lived connections get generated, the more resources are spent setting up and tearing down these connections.</p>	Advanced
<p># Flows - RTT (1SD)</p> <p># Flows - RTT (2SD)</p> <p># Flows - RTT (3SD)</p>	<p>Number of flows whose Round Trip Times are 1, 2, or 3 Standard Deviations higher than the average for all flows in the parent group. Under ideal conditions, all flows should be within 1SD of the average.</p> <p>High-RTT connections, even if they are intermittent, can indicate an issue worth investigating. If the average response time for an ecommerce web app is 2 seconds (acceptable), but 5% of transactions are 20 seconds or higher (not acceptable), this can result in a significant number of unhappy customers and lost revenue.</p>	Advanced
# Flows - TCP Data Rxmt	Number of flows within which any packet was retransmitted. Retransmissions are an indication of packet loss.	Advanced
# Flows - TCP Resets	The average number of flows that were closed by a TCP Reset.	KPI
# Flows - w/o TCP SACK	Number of flows with the TCP Selective Acknowledgement (SACK) option disabled. With SACK enabled, a receiver can send SACK packets to acknowledge receipt of multiple data packets in the case of lost segments. This improves network performance by reducing the number of retransmissions. With SACK disabled, the receiver must resend all the packets after the last lost segment even if they were received by the peer.	Advanced
# Flows - w/o TCP Timestamp	Number of flows with the TCP Timestamp option disabled. This option is used for calculating more accurate Round Trip Times.	Advanced
# IP Fragment Count	The total number of IP fragments attributes to this TCP connection group. A high level of fragmentation is an indication of network issues and can severely affect application performance.	KPI
# Loss Pkts	The total number of packets that were lost (sent but never received).	Advanced
# Nagle Delays	The number of times a message-send event was delayed because the sending node had to wait for previously-sent data to be acknowledged (ACK'd). This occurs most often when Nagle's algorithm and Delayed ACK are enabled on the sending and receiving node (respectively).	Advanced

# PIE Events	<p>Performance Impacting Elements (PIE) are useful for identifying the location of actual or potential bottlenecks:</p> <ul style="list-style-type: none"> • Client Limited and Client Zero Window events indicate a possible problem on the client node. • RTOs indicate a possible problem on the network path between two nodes. • Server Limited and Server Zero Window events indicate a possible problem on the server node. 	KPI
# RetransmissionTimeouts	<p>Retransmission Timeouts (RTOs) are a sign of network packet loss, which results in retransmission of data when the TCP retransmission timer expires (Timer to make sure data is ACK'ed).</p> <ul style="list-style-type: none"> • Typically this timer varies from 200ms-3 sec by default (different for each operating systems and their versions). • It causes severe performance degradations as considerable amount of time gets wasted to resend the lost data. • TCP falls back to the "Slow Start" phase impacting the performance even more. 	KPI
# SACK Retransmits	Average number of packets that were retransmitted due to a SACK message that indicated an unreceived packet.	KPI
# Server Limited	The number of TCP window updates sent by a server node indicating that it is receiving data at a faster rate than the application can process.	KPI
# Server Zero Window	<p>If a server's TCP buffer is full, it sends Zero as the TCP receive window size to indicate that it cannot receive more data. The data transfer then stops until the server can process the data in its buffer.</p> <p>A high rate of Server TCP Zero Window messages indicates a problem with either</p> <ul style="list-style-type: none"> • The TCP configuration on the node, or • The client-side application using that flow. 	KPI
# Syn Blackholes	The number of connection attempts that went unanswered and resulted in a failure. Syn blackholes can severely impact application performance.	KPI
# Syn Resets	The number of connection attempts that were explicitly refused by the other host. Syn resets can severely impact application performance.	KPI
# TCP Resets - Established	The average number of times an established TCP flow was reset.	KPI
# TCP Resets - Fin	The average number of times a TCP flow was reset while it was in the process of getting closed.	KPI
# TTL Changes (1 - 2 hops) # TTL Changes (3 - 4 hops) # TTL Changes (>=5 hops)	<p>The number of routing-hop changes experienced by this connection. Frequent variations in routing-hop changes indicate routing problems in the network and can severely affect app performance.</p> <p>Under ideal conditions, the number of hops should be consistent.</p>	Advanced
Avg # TTL Hops	The average number of routing-hop changes experienced by this connection.	Advanced
Delayed Acks Lag (usec)	The average amount of lag that Delayed ACKs are adding to the overall Latency of the Connection.	Advanced

Initial RTT (usec)	Round trip time for the initial two SYN packets (between SYN and SYN-ACK or SYN-ACK and ACK depending upon whether the agent is running on client or server).	KPI
Latency - RTT (usec)	Average round-trip latency (from packet transmission to acknowledgement) for all packets.	KPI
Lifetime (usec)	Average lifetime (from initial setup to final teardown) for TCP sessions in the Connection group.	KPI
Loss (per mille)	The number of packets lost per 1000 packets sent. "Per mille" is a percentage with one additional digit of precision.	KPI
Nagle Delays Lag (usec)	The average amount of lag that Nagle delays are adding to the overall Latency of the Connection.	Advanced
Pkts per Sec	Average rate of packets sent and received	KPI
Rx Pkts per Sec	Average rate of packets received	Advanced
Rx Throughput (BPS)	Average rate of bytes received	Advanced
TCP Handshake (usec)	Round trip time for the initial three-way connection setup, for all flows in the parent group: <ol style="list-style-type: none"> 1. SYN (client --> server) 2. SYN-ACK (client <-- server) 3. ACK (client --> server) 	KPI
Throughput (BPS)	Throughput (bytes per second) for the application of interest on all TCP sessions.	KPI
Tx Pkts per Sec	Rate of packets sent	Advanced
Tx Throughput (BPS)	Rate of traffic received	Advanced