


Choosing a Transport Protocol for Real-time Data across complex networks



Introduction

Choosing the right protocol for data transport depends on the needs of the application and the nature of the data. Most applications use a standard network transport protocol: either TCP or UDP. However these default choices do not satisfy strategic applications that use time-critical data. This guide lays out the important factors in choosing a protocol alternative. It's not for everyone; the high priority needs of strategic applications and real-time data are as follows.

Needs

Web browsing and email are well served by standard TCP. Forgiving streaming applications like audio are typically satisfied by UDP. More demanding applications have the following needs:

Application Needs

Mission critical
Uses a complex network not optimized for it
Spans long distances
Application owner believes it is more important than most network traffic

Data Needs

Near real-time
Streams, not files
Latency and jitter sensitive
Packet loss sensitive
High throughput (> 40 Mbps)

There are 6 main criteria to consider in choosing a protocol that meets these needs...

1 **Reliable:** Important data must arrive at its intended destination. IP networks divide data into packets for transmission, and those packets must arrive in the correct order. A reliable protocol assures delivery in the correct order. TCP is considered a reliable protocol, UDP is not. If we stopped here then TCP would be fine, but this criteria isn't the only one. TCP fails to meet others discussed below.

Real-time data is time sensitive. As a result, some applications use UDP because it is latency sensitive. However UDP isn't reliable. UDP caters to the time sensitivity but has no reliability features. Choose a protocol that includes guaranteed delivery.



2 High Throughput: Long fat networks are a challenge for transport protocols. A long fat network is one with a high throughput capability, but with a relatively long transmission time usually due to distance. Protocols like TCP that emphasize congestion avoidance are adversely affected by long transmission times, and thus cannot supply high throughput on long networks, however fat the pipe.

Age is also a factor in throughput. Protocols that were designed over 10 years ago did not anticipate the need for high data rates that modern applications require. For example, a recently launched satellite sensor supplies earth observation data at 3.8 Gbps. This rate was practically unimaginable in the late 70's and 80's when TCP and UDP were developed.



3 Flexibility: On just the first two criteria, the need for tradeoffs becomes clear. The best reliability leads directly to reduced throughput. A good protocol provides ways to make tradeoffs in key properties. These include latency, overhead, and error correction. Balancing these to fit mission requirements is the property of flexibility. Most users will be well served by choosing a protocol with good flexibility since it is unlikely that the designer had your application and network in mind.



4 Latency Sensitive: Certain applications are sensitive to latency and jitter. Most people are familiar with pauses in video stream playback. In small doses, these are annoying but not unbearable. However a live interactive video conference becomes impossible when data isn't received in a timely way.

A latency sensitive protocol prioritizes the turn around time from request to answer and is aware of the importance in minimizing delays. It minimizes transmission time beyond the physical delays caused by the length of the network path. By nature, a protocol that waits for a response before sending the next packet takes longer than one that does not wait. Real-time data requires a protocol that is latency sensitive.



5 Performant: The word performant is an engineering term for something that meets the performance expectations for which it was created. It's as fast and efficient as you would expect it to be, while still remaining within other constraints. There is no implication that its the optimal or best solution.

For real-time transport protocols, a performant solution operates as expected on networks with packet loss, reordering, duplication, jitter, and drop outs. In the face of those challenges, it meets the goals for reliability, latency, and throughput.



6 **Visibility:** Network conditions vary over time. Even a network with no impairments can develop them as traffic grows or the configuration is modified. It can be useful for the transport protocol to provide insights into network impairments and impacts. Reporting these conditions will allow parameters to be adjusted as conditions vary. It's a chance to future proof the solution, and the emerging view is that resilient networks require situational awareness of threats.



There are a couple of other things to worry about...

Support

If you have issues deploying the protocol, or if it isn't tuned to your specific situation, you will need support. Open source protocols offer only the support of the open-source community. While some open source projects have a large and robust community, others are maintained as a hobby by a single individual, or have been completely abandoned. A commercial product will naturally have a company that is dedicated to helping you. What level of support will you need?

Implementation Quality

A protocol may use sound design principles and have the attributes you need, but how solid is the software implementation that is available? Some protocols were developed as a research project, but the implementation was never validated or characterized. Some implementations even have known shortcomings and defects. It is difficult to judge, but it is important to consider the quality of the implementation in addition to the design of the protocol.



When looking for a protocol to transport real-time data, these are some common choices:

TCP Variants

TCP-Cubic
 TCP-Illinois
 SCPS-TP
 SCTP

Reliable UDP

RUDP
 UDT
 ENET
 NORM
 SMPTE 2022-1

Commercial

DataDefender
 FASP
 MPT
 SpeedyPackets

TCP Variants

As a group, these protocols preserve the main focus of TCP: reliability and congestion avoidance. They attempt to improve on TCP throughput and decrease the impact of network congestion on long, fat networks. The variants differ in the way they manage congestion, control data rate, and react to the lack of acknowledgments. Some differ on handling packet loss due to congestion or corruption. They generally don't prioritize low latency delivery, but there are substantial differences among them for networks that include a long transmission path such as satellite link.





Reliable UDP Variants

Underlying UDP isn't reliable. As a group, Reliable UDP variants provide better latency sensitivity and throughput than the TCP variants while improving on UDP for reliability. As a category, they are better than TCP variants for real-time data.

There are differences on the individual factors between the variants, so one isn't universally best.

Commercial Products

When shopping for a commercial protocols, you'll find that they satisfy the needs of the application areas that were targeted by their developers. Its probably not a surprise, but as a group they perform better than the previous two categories.

All offer superior reliability, but differ on latency sensitivity and throughput. Flexibility and visibility are also very different. Differences in application need will guide the priority of these properties. What types of data and applications are the products optimized for?



Conclusion

Generally, TCP variants will not provide the low latency and high throughput required for real-time data and strategic applications operating over complex networks.

There are choices amongst the Reliable UDP protocols and Commercial products that make sense. In the next piece we assess the protocols on each of the 6 factors and identify the best choices.

For more information,
visit: Kratos RT Logic
www.rtlogic.com

Look for the next piece in the series:
The Best Protocol for Real-time Data Transport from
www.rtlogic.com/products/datadefender

