

Guidelines for Internet Congestion Control at Endpoints

Gorry Fairhurst, University of Aberdeen, UK
draft-fairhurst-ccwg-cc-00
replaced draft-fairhurst-tsvwg-cc

BCP41: Congestion Control Principles

Outrageous opinion(s):

BCP41: Congestion Control Principles

Outrageous opinion(s):

- 1. Much of what was learned in the 90s is still relevant!**

BCP41: Congestion Control Principles

Outrageous opinion(s):

- 1. Much of what was learned in the 90s is still relevant!**
- 2. Links, hosts, and routers are very different!**

BCP41: Congestion Control Principles

Outrageous opinion(s):

- 1. Much of what was learned in the 90s is still relevant!**
- 2. Links, hosts, and routers are very different!**
- 3. TCP provided safe operation for decades, what comes next?**

BCP41: Congestion Control Principles

Outrageous opinion(s):

- 1. Much of what was learned in the 90s is still relevant!**
- 2. Links, hosts, and routers are very different!**
- 3. TCP provided safe operation for decades, what comes next?**

The target is to replace or update RFC 2914 (Sally Floyd, 2000)

Requirements from the RFC Series

2000

1 Gbps fibre had just emerged;
100 BT was quite new.



2023

.... things have moved on.

Requirements from the RFC Series

2000

1 Gbps fibre had just emerged;
100 BT was quite new.



2023

.... things have moved on.

Changes in the network: At one time, it was common that the serialisation delay of a packet at the bottleneck formed a large proportion of the round time of a path, motivating a need for conservative loss recovery. This is not often the case for today's higher capacity links.

Requirements from the RFC Series

2000

1 Gbps fibre had just emerged;
100 BT was quite new.



2023

.... things have moved on.

Changes in the network: At one time, it was common that the serialisation delay of a packet at the bottleneck formed a large proportion of the round time of a path, motivating a need for conservative loss recovery. This is not often the case for today's higher capacity links.

Changes in practice: Changes over time in the way that protocol mechanisms are deployed in Internet endpoints

Requirements from the RFC Series

2000

1 Gbps fibre had just emerged;
100 BT was quite new.



2023

.... things have moved on.



Changes in the network: At one time, it was common that the serialisation delay of a packet at the bottleneck formed a large proportion of the round time of a path, motivating a need for conservative loss recovery. This is not often the case for today's higher capacity links.

Changes in practice: Changes over time in the way that protocol mechanisms are deployed in Internet endpoints

Still true: The Internet continues to be heterogenous!

Requirements from the RFC Series

Requirements

- 3.1 The need to React to Congestion
- 3.2. Tolerance to a Diversity of Path Characteristics
- 3.3. Protection of Protocol Mechanisms

Incipient v persistent congestion

Congestion

Incipient congestion: a consequential side effect of the statistical multiplexing of packet flows the pattern of arriving traffic results in over consumption of the path resources.

Persistent congestion: could adversely impact other flows, e.g., starving them of resources; or reduce the efficiency of the path (e.g., congestion collapse).

Principles of Congestion Control

4.1. Robustness: Timers and Retransmission

Principles of Congestion Control

- 4.1. Robustness: Timers and Retransmission
- 4.2. Detecting and Reacting to Incipient Congestion

Principles of Congestion Control

4.1. Robustness: Timers and Retransmission

4.2. Detecting and Reacting to Incipient Congestion

4.3. Avoiding Persistent Congestion

(i.e. avoiding Congestion Collapse and Flow Starvation)

Next Steps

If there is interest, can we adopt this? and will we finish it?

I said that I would give-up if nobody joined me - someone has offered, will you come and help also?

*In a 1948 speech to the House of Commons, Churchill paraphrased
'Those who fail to learn from history are condemned to repeat it.'*