Responsiveness under Working Conditions

draft-ietf-ippm-responsiveness-01

Christoph Paasch, Randall Meyer, **Stuart Cheshire**, Omer Shapira, Matt Mathis

Upcoming update

Major changes

- Issue #62 (Flaw in “Working Conditions” algorithm)
  - Algorithm should terminate when delay stops increasing, not when throughput stops increasing
- Issue #17 (Use well-known URIs)
  - Default well-known URI that all web servers can supply if they wish

Issue #62: Problem statement

1 Gb/s
256 ms RTT
Bandwidth-Delay Product (BDP) = 32 MB

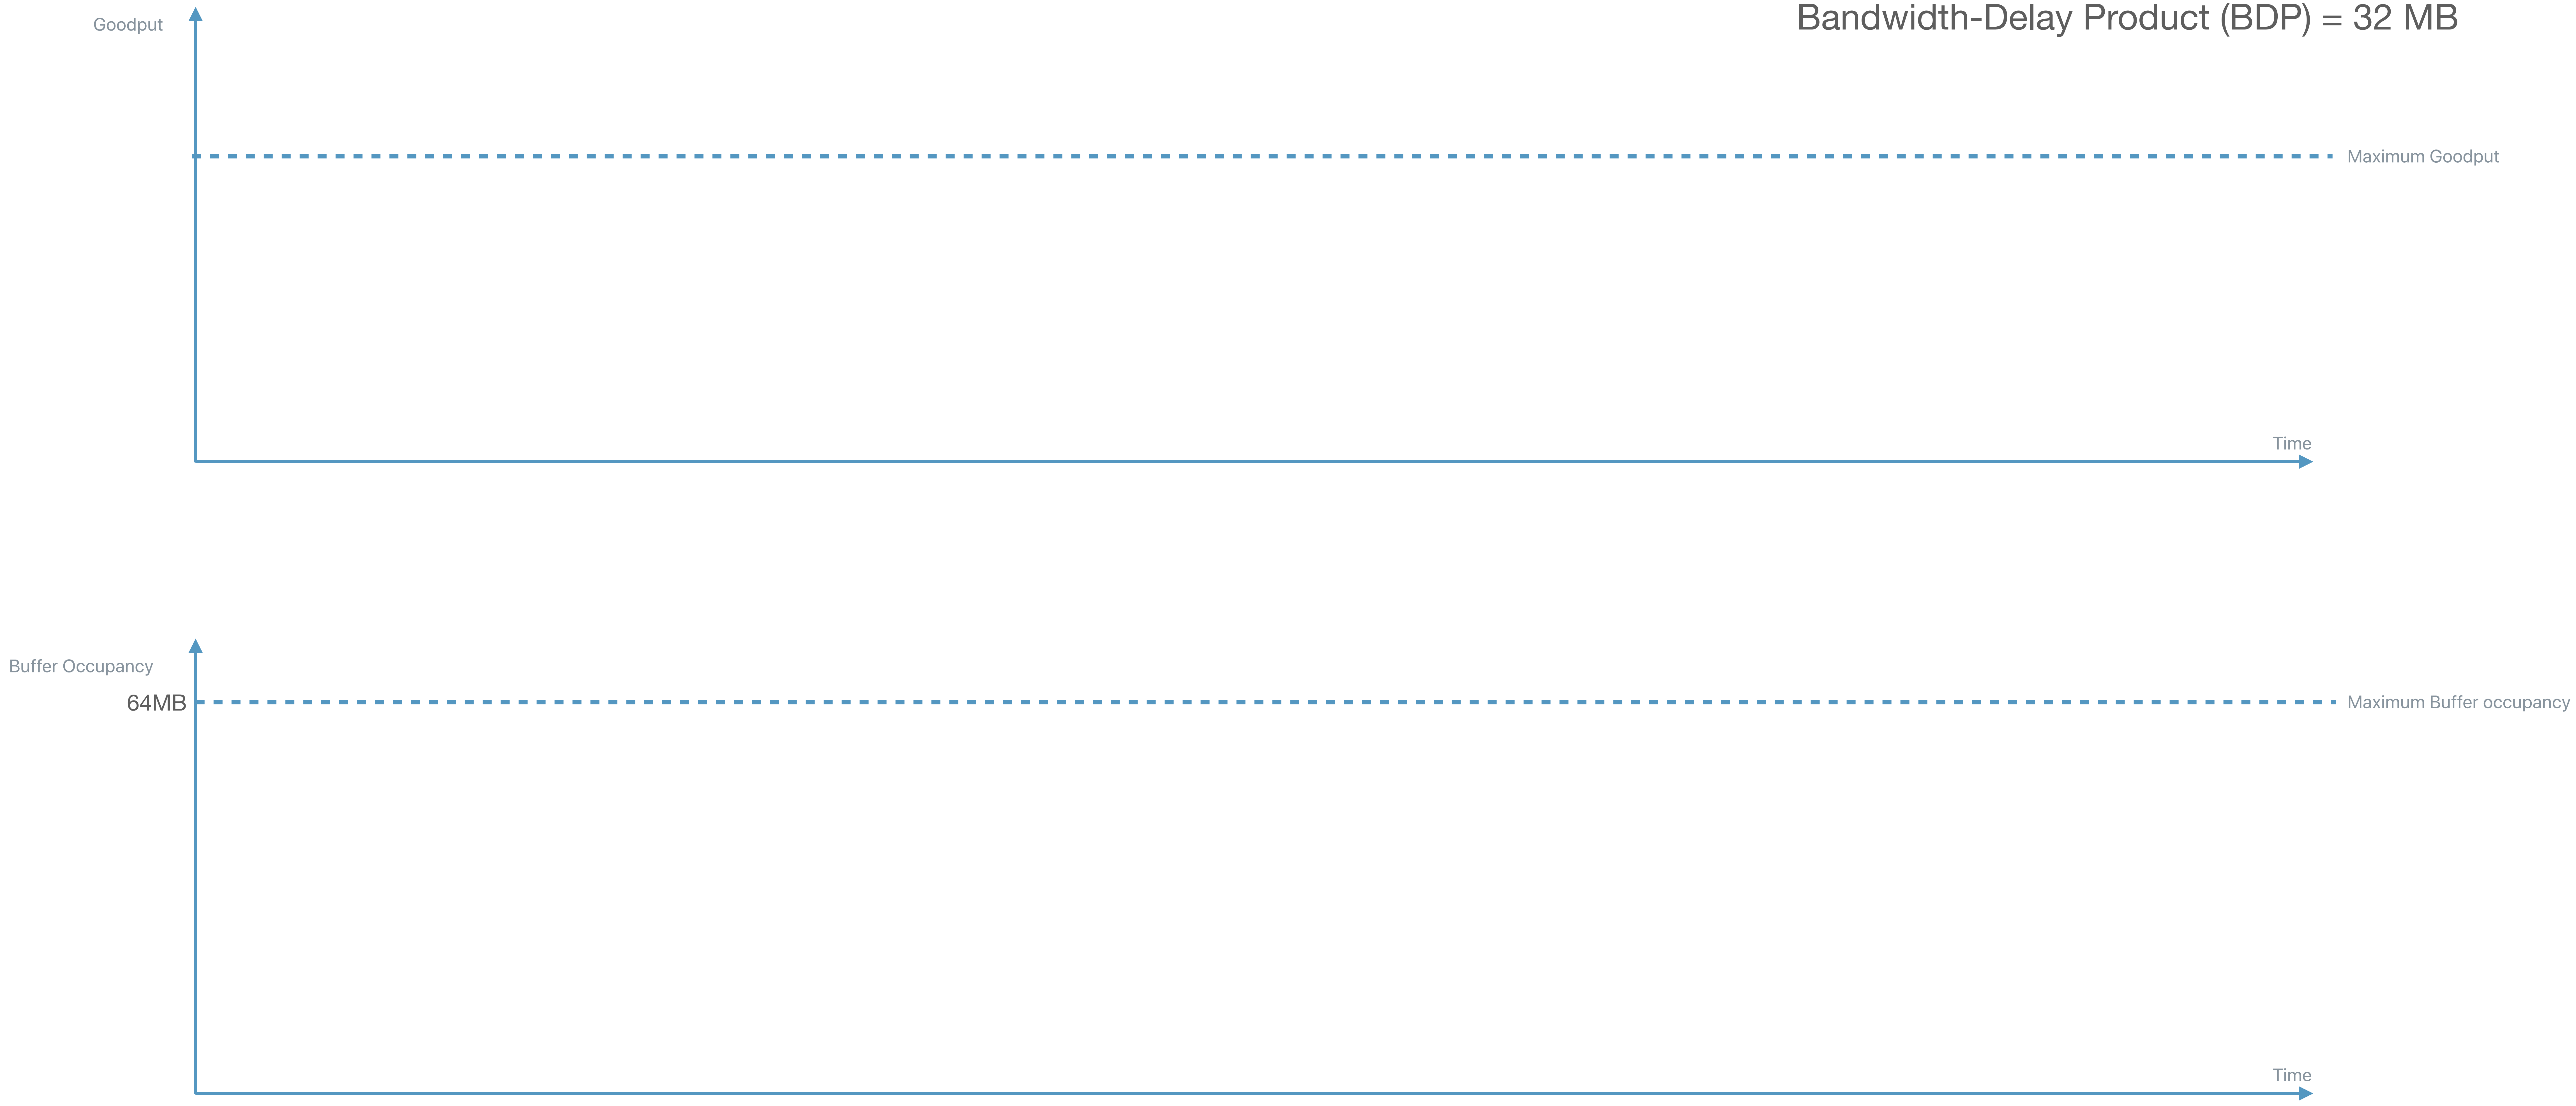

Issue #62: Problem statement

1 Gb/s
256 ms RTT
Bandwidth-Delay Product (BDP) = 32 MB

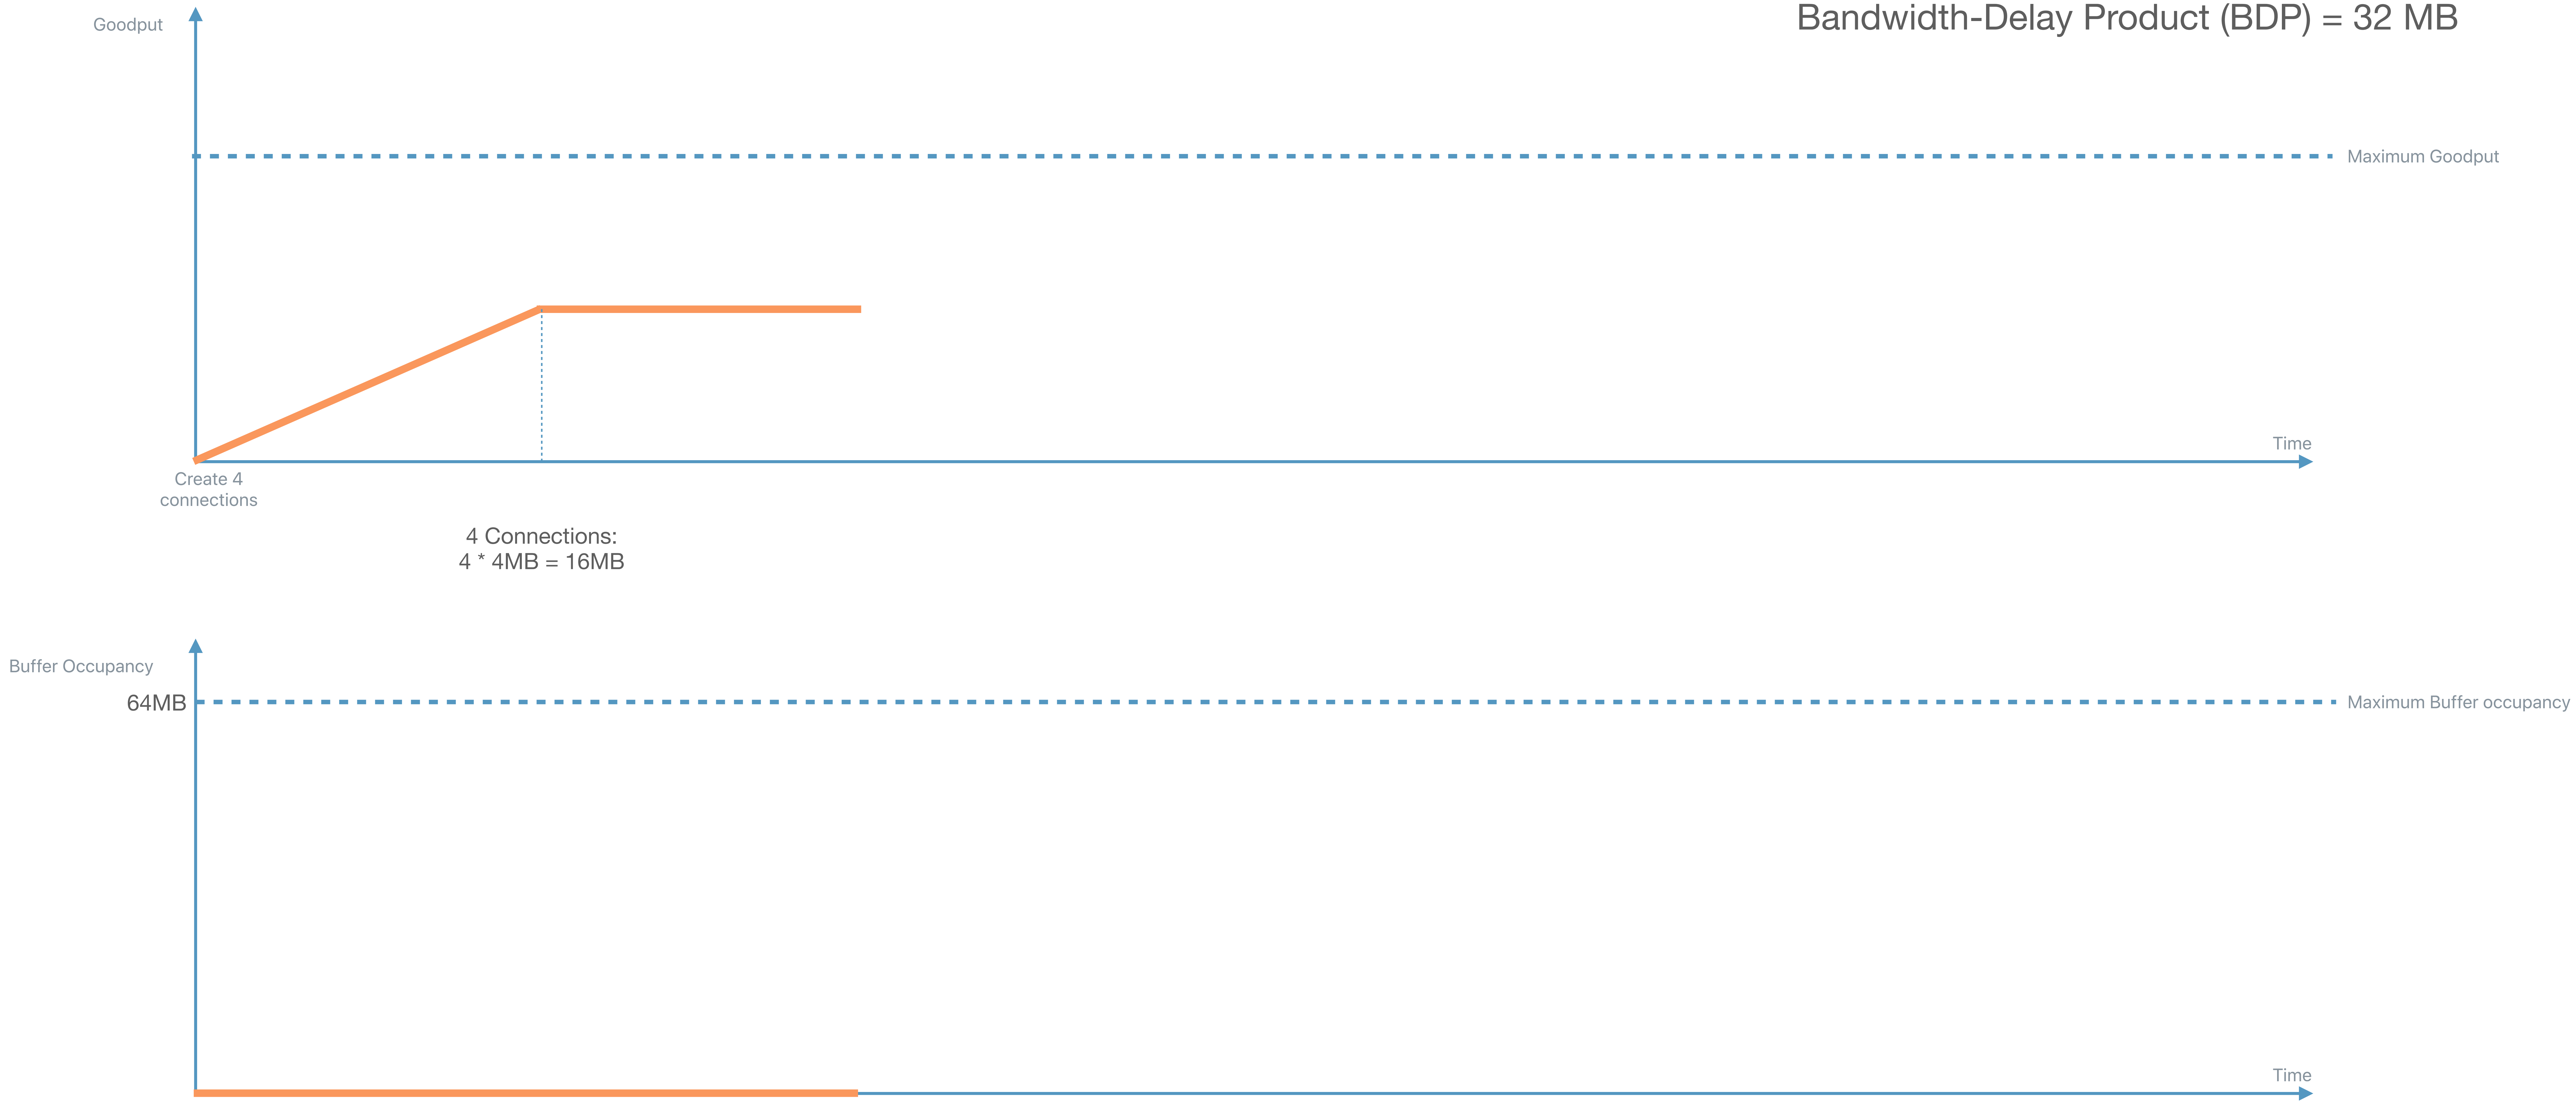

Issue #62: Problem statement

1 Gb/s
256 ms RTT
Bandwidth-Delay Product (BDP) = 32 MB

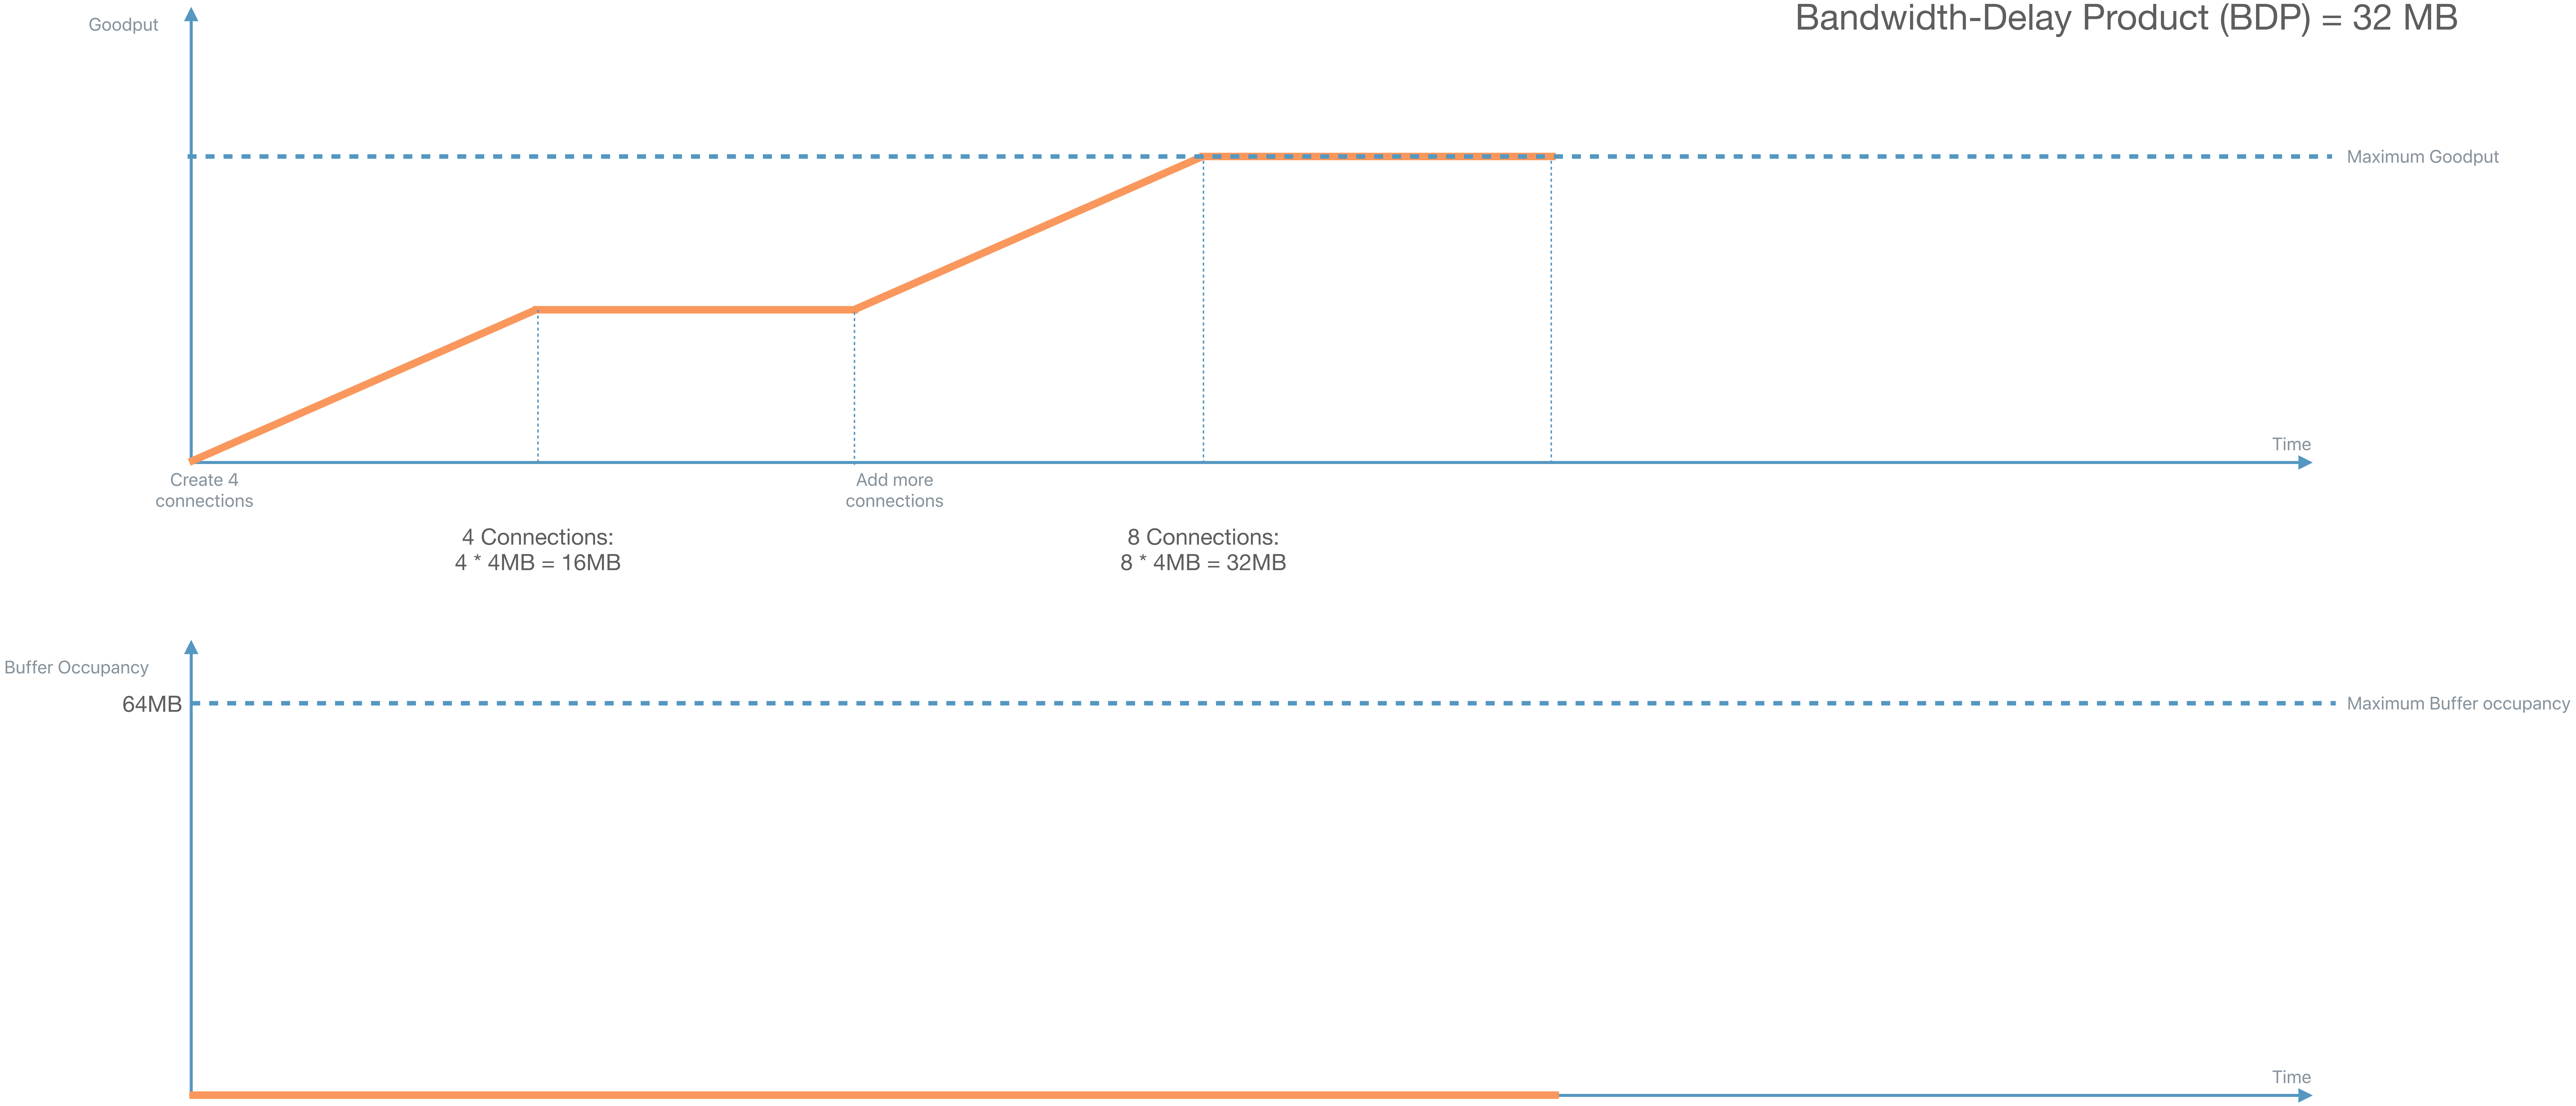

Issue #62: Problem statement

1 Gb/s
256 ms RTT
Bandwidth-Delay Product (BDP) = 32 MB

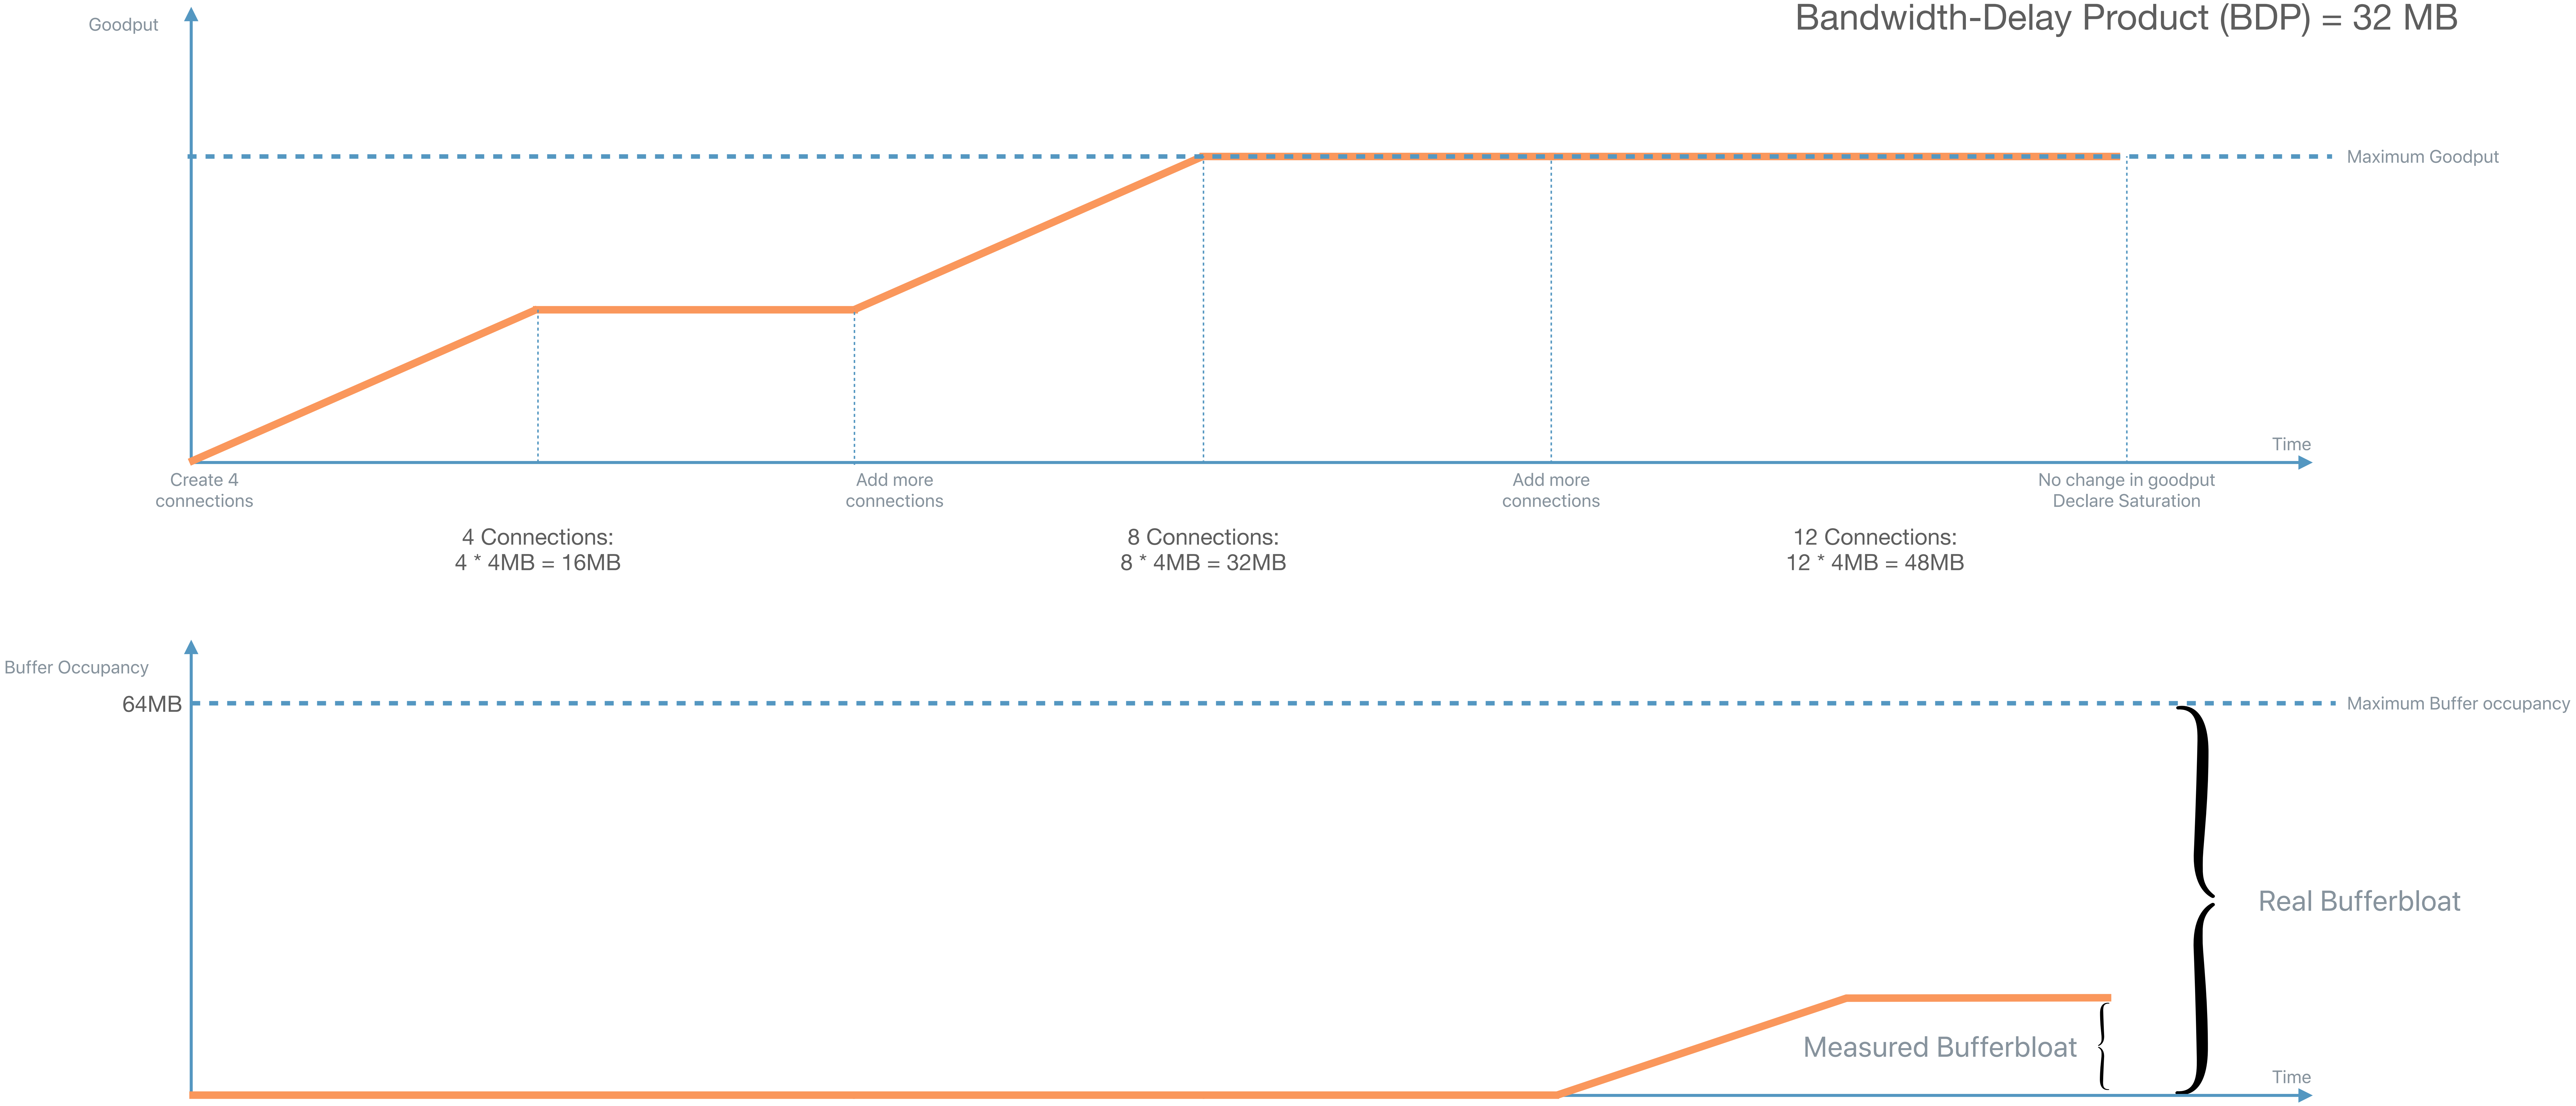

Issue #62: Solution

- Continuously add 1 connection to the pool every second
- Continuously probe responsiveness every 100ms
- Monitor goodput and responsiveness evolution over time
- Once both stabilize, declare saturation and use final responsiveness result

Issue #17: Use well-known URIs

- Json-configuration allows to bootstrap a responsiveness test
- Servers may implement a default responsiveness measurement endpoint
- Well-known URI to host the json-configuration

```
.well-known/responsiveness.json
```

Remaining issues

https://github.com/network-quality/draft-ietf-ippm-responsiveness/issues

- Issue #55: Specify confidence in the results
  - If stability is not reached within a certain timeframe, responsiveness can produce an “early” result with a lower confidence.
- Issue #66: Allow non-TLS measurements
  - Low-end devices don’t have sufficient CPU to load the link with TLS traffic

Responsiveness under Working Conditions

draft-ietf-ippm-responsiveness-01

Christoph Paasch, Randall Meyer, **Stuart Cheshire**, Omer Shapira, Matt Mathis