

# MPTCP – Multipath TCP

WG Meeting

Berlin, IETF-87, 30<sup>th</sup> July 2013

Philip Eardley

Yoshifumi Nishida

- Note taker
- Jabber [IMPORTANT]
- Please include “-mptcp-” in your draft names
- Please say your name at the mike

# Note Well

**Any submission to the IETF intended by the Contributor for publication as all or part of an IETF Internet-Draft or RFC and any statement made within the context of an IETF activity is considered an "IETF Contribution". Such statements include oral statements in IETF sessions, as well as written and electronic communications made at any time or place, which are addressed to:**

- the IETF plenary session,
- any IETF working group or portion thereof,
- the IESG, or any member thereof on behalf of the IESG,
- the IAB or any member thereof on behalf of the IAB,
- any IETF mailing list, including the IETF list itself, any working group or design team list, or any other list functioning under IETF auspices,
- the RFC Editor or the Internet-Drafts function

**All IETF Contributions are subject to the rules of RFC 3978 (updated by RFC 4748) and RFC 3979 (updated by RFC 4879). Statements made outside of an IETF session, mailing list or other function, that are clearly not intended to be input to an IETF activity, group or function, are not IETF Contributions in the context of this notice.**

**Please consult RFC 3978 (and RFC 4748) for details.**

**A participant in any IETF activity is deemed to accept all IETF rules of process, as documented in Best Current Practices RFCs and IESG Statements.**

**A participant in any IETF activity acknowledges that written, audio and video records of meetings may be made and may be available to the public.**

# WG Item Status

- MPTCP Application Interface Considerations  
(draft-ietf-mptcp-api)
  - RFC6897
  - Thanks to Michael & Alan

# Milestones

- Dec 2012: Consensus on what high-level changes are needed to the current MPTCP Experimental document in order to progress it on the standards track
  - Apr 2013: Implementation advice (Informational) to IESG
  - Aug 2013: Use-cases and operational experiences (Informational) to IESG
  - Dec 2013: MPTCP-enabled middleboxes (Informational) to IESG
  - Dec 2013: MPTCP standards track protocol to IESG
- 
- These need revising
  - We're progressing on most of these, except for the middlebox one?

# Other mptcp stuff in Berlin

- Sunday: tutorial – Olivier Bonaventure
- Wed 3pm: mptcp interop (Dahlem)
- Thurs 9am: mptcp update in tsvarea

# Agenda

1. Chairs update (Chairs, 15 mins)
  - a. Status update
  - b. Implementation survey
2. FreeBSD MPTCP Implementation Update (Nigel Williams, 15 mins)
3. Linux MPTCP Implementation Update (Christoph Paasch, 15 mins)
4. Conformance Tests for Multipath TCP (Yvan Coene, 10 mins)
5. MPTCP Use Case (Costin Raiciu, 10 mins)
6. Residual threats analysis for MPTCP (Marcelo Bagnulo, 15 mins)
7. Discussion on next steps (Chairs, 20 mins)
8. The NorNet Testbed: A Platform for Evaluating Multi-Path Transport in the Real-World Internet (Thomas Dreibholz, 10min)
9. Why protocol stacks should be in user-space? (Michio Honda, 10 mins)

# Status - Implementations

- We have 5 independent implementations!
  - Linux, UCLouvain
  - FreeBSD, Swinburne
  - Commercial OS, Anon \*
  - NetScaler, Citrix
  - User-space \*\*
- RFC6824 is well implemented and understood
- Interoperate with Linux 'reference'

\* Not publicly available

\*\* RFC compliant, but no longer maintained



# Implementations survey

## (1) signalling

```

|
| Question 3: Support for MPTCP's signalling functionality
| MPTCP's signalling messages are: MP_CAPABLE, MP_JOIN, Data transfer
| (DSS), ADD_ADDR, REMOVE_ADDR, MP_FASTCLOSE. There are sub-questions
| for MP_JOIN and DSS.
|
|           | UCLouvain | Swinburne | Anon      | Citrix    |
|MP_CAPABLE |Yes        |Yes        |Yes        |Yes        |
|MP_JOIN    |Yes        |Yes        |Yes        |Yes        |
|initiated by|first end  |either end |first end  |first end  |
|#subflows  |32         |8          |no limit   |6          |
|DSS        |Yes        |Yes        |Yes        |Yes        |
|DATA ACK   |4 bytes   |4 or 8 byte|4 or 8 byte|4 or 8 byte|
|Data seq num|4 bytes   |4 or 8 byte|4 or 8 byte|4 or 8 byte|
|DATA_FIN   |Yes        |Yes        |Yes        |Yes        |
|Checksum   |Yes        |No         |Yes        |Yes        |
|ADD_ADDR   |Yes        |No         |No (never) |No (never?)|
|REMOVE_ADDR|Yes        |No         |Partly     |Yes        |
|FAST_CLOSE |Yes        |No         |Yes        |Yes        |
|

```

- Signalling works well
- ADD\_ADDR needs more discussion
- Details in draft-eardley-mptcp-implementations-survey

# Implementations survey

## (2) fallback

```
|
| Question 4 asks about fallback from MPTCP: if a middlebox mangles
| MPTCP's signalling by removing MP_CAPABLE, MP_JOIN, DSS or DATA_ACK;
| if data is protected with Checksum in DSS option; if fallback to TCP
| uses an infinite mapping; and if any corner cases have been found.
|
|      | UCLouvain | Swinburne | Anon   | Citrix |
|MP_CAPABLE | Yes      | Yes      | Yes    | Yes    |
|MP_JOIN    | Yes      | Yes      | Yes    | Yes    |
|DSS        | Yes      | No       | Yes    | Yes    |
|DATA_ACK   | Yes      | No       | No     |        |
|Checksum   | Yes      | No       | Yes    | Yes    |
|infinite map | Yes     | Yes      | Yes    | Yes    |
|corner cases| No      |          | Yes    | Yes    |
|
```

- Fall-back to TCP works well
- A few clarifications are needed
- Details in [draft-eardley-mptcp-implementations-survey](#)

# Implementations survey

## (3) congestion control

```
|
| Question 8 asks about congestion control and related issues: how
| traffic is shared across multiple subflows; support for 'handover';
| and support of RFC6356 (or other) coupled congestion control.
|
|           | UCLouvain | Swinburne |   Anon   | Citrix   |
|sharing    |shared, RTT|shared     |active/back|active/back|
|handover   |Yes        |           |Yes        |Yes       |
|coupled cc |Yes        |No         |No         |No        |
|other ccc  |Yes, OLIA  |No         |No         |No        |
|MP-PRIO & B|Yes        |No         |Yes        |Yes       |
|
```

- Use of mptcp for 'active standby'
- OLIA is proposed improvement to RFC6356, draft-khalili-mptcp-congestion-control
- Several other multipath CC algorithms in the literature
- Details in draft-eardley-mptcp-implementations-survey

# Implementations survey

## (4) API

```
|
| Question 9 is about the API: how legacy applications interact with
| the MPTCP stack, and if implemented the RFC6897 API for MPTCP-aware
| applications.
|
|           | UCLouvain | Swinburne | Anon   | Citrix |
| legacy apps | default   | sysctl    | private API | configured |
| MPTCP API   | No        | No        | No     | No     |
| advanced API | No        | No        | No     | No     |
|
```

- API not really been explored yet
- Details in draft-eardley-mptcp-implementations-survey

# Next steps

- Moving RFC6824 to Standards track
  1. ADD\_ADDR needs more discussion
  2. Fall-back needs a bit more clarification
  3. ‘Better’ security may be needed
    - Now: during initial handshake exchange keys in clear, then use keyed HMAC
  4. More operational experience of different use cases, scenarios...
    - “particularly looking for cases where MPTCP could be detrimental in some way”
- Implementation advice (heuristics)
- MPTCP-aware middlebox (where at least one end host is MPTCP-enabled)
- How do we progress on these?