

LE codepoint: preliminary results and ongoing work in the IETF

MAP Research Group, March 2017

Gorry Fairhurst

Ana Custura

Andre Venne

Tom Jones (edgetrace)

(University of Aberdeen, UK)

neat



Which DSCP for LE?

- RFC 3662 designates CS1 for 'low priority' data
- RFC 4594 talks about using CS1 (DSCP 8) for LE
- Proposed LE PHB: 000010 (DSCP 2)
 - draft-ietf-tsvwg-le-phb-01
- Measurements over 2 campaigns using *PATHscope* and *PATHspider* from the core to web servers.

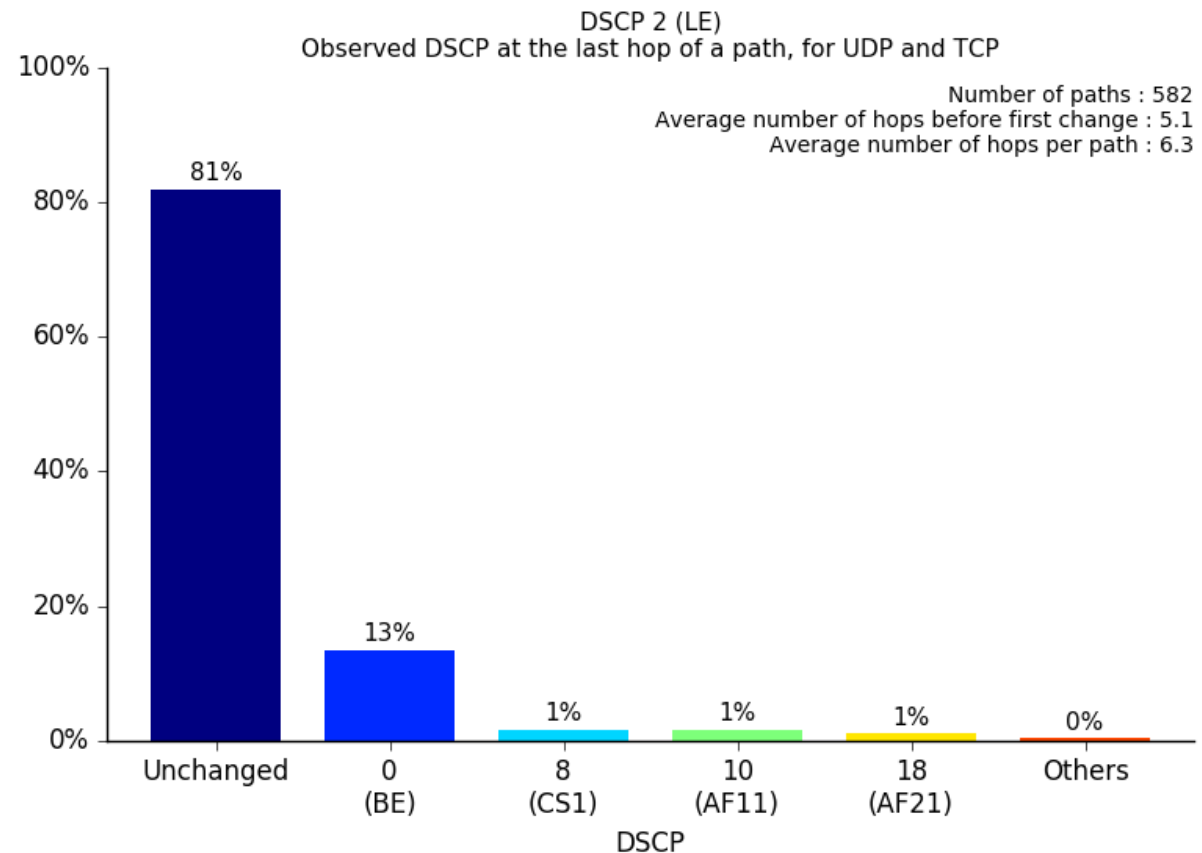
Preliminary results - PATHspider.net

- Entire Alexa Top 1 million website list resolved and IP addresses A/B tested for DSCP 0 & 2
 - Tested 4.2 Million unique paths
 - 3-way handshake completed for *both* DSCPs for 99.998% of paths
 - 7200 failures attributed to packet loss when testing
- Similar results for DSCP 46 (EF) (99.97% of paths)
- *Setting a DSCP did not blackhole packets*

(Details at end of slides)

Preservation of DSCP 2 (000010) Results for *582 unique paths*


- DSCP traverses 81% of paths unchanged
- 19% remarked, but 13% remarked to 0 (BE)
- No difference between TCP & UDP



Pathologies

%age *routers* changing DSCP

- DSCPs 1 to 7 had lowest modification rate, changed for <3% of routers; (unaffected by routers modifying upper 3 bits)
- A DCSP that sets the upper 3 bits is modified by a router using RFC791 ToS Semantics
- CS1 (DSCP 8) changed by 5% more routers than DCSP 2



Problems with routers using ToS Semantics

- In a study of 524 routers, 88 (16.79 %) bleached the upper three bits of the DS field:

10, 16, 26, 34 (AFx1) & 000111 = 2

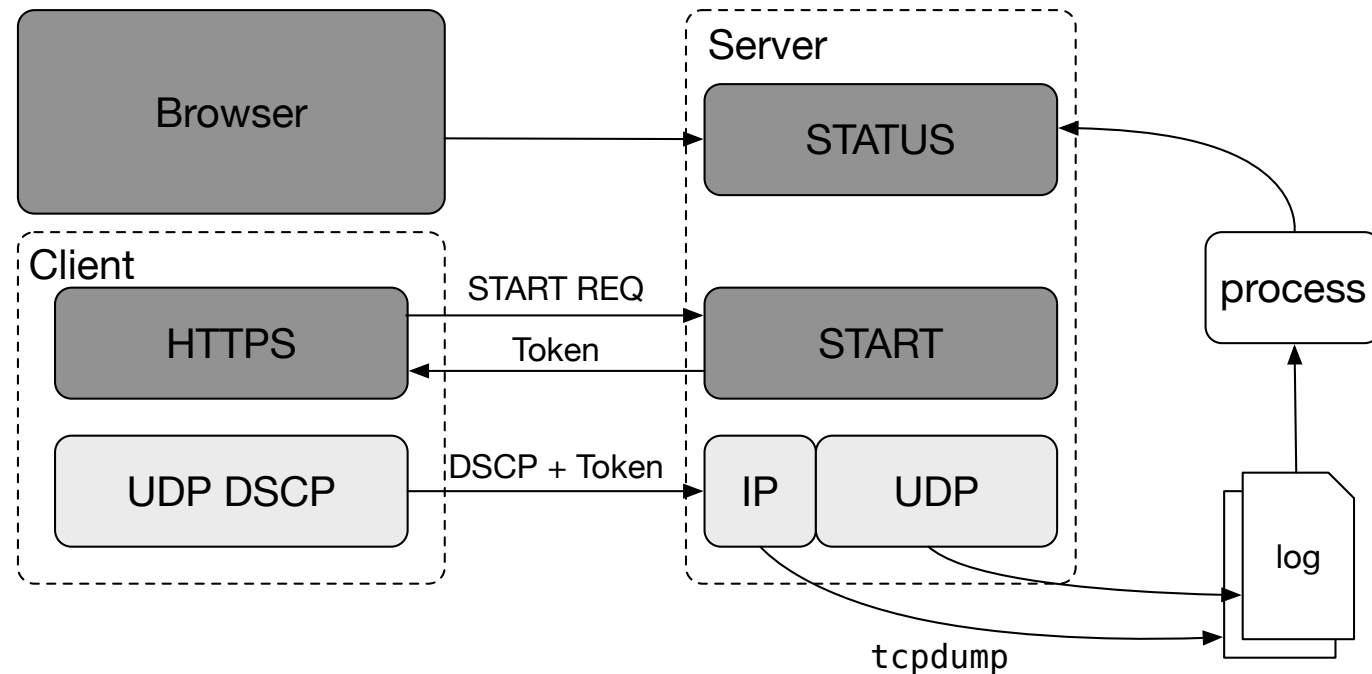
12, 18, 28, 26 (AFx2), 44 (Voice Admit) & 000111 = 4

14, 22, 30, 38 (AFx3), 46 (EF) & 000111 = 6

- AFx1, etc can be remarked as LE, resulting in priority inversion
- CS1 priority already ambiguous using ToS Semantics

Next steps: Edgetrace

- We need to find out about DSCPs in the Edge
- We made an easy to install go command line tool
- This tests DSCP values from your current location
- Please download at trace.erg.abdn.ac.uk



Conclusion

– Is DSCP2 good for LE?


- **DSCP 2 has good survivability**
 - Good compatibility with 802.11 (can be updated)
 - Not affected by pathologies resulting from ToS semantics
- **Routers/Hosts still using ToS semantics problematic**
 - Routers can cause priority inversion for higher DSCPs
 - ToS applications (Portable OpenSSH sets DSCPs 2 & 4)
- **DSCP 2 is better than CS1**
 - CS1 was more often remarked
 - CS1 was more prone to priority inversion

Questions & Answers

- Please use *Edgetrace* to measure edge networks
- try it and help us understand more of what really works



Observed DSCP at end of path



Detailed plots

Windows



MAC



Linux



BSD



trace.erg.abdn.ac.uk

Extra slides - for info

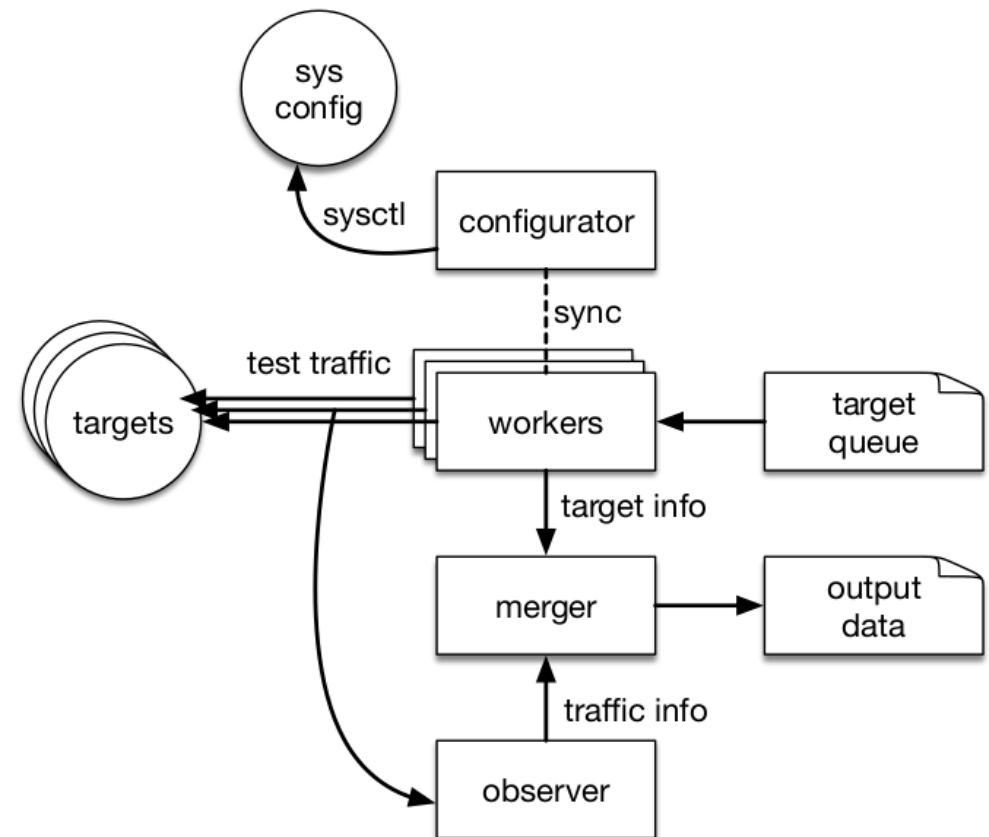
Acknowledgements

- **This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 688421, and was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 15.0268. The opinions expressed and arguments employed reflect only the authors' views. The European Commission is not responsible for any use that may be made of that information. Further, the opinions expressed and arguments employed herein do not necessarily reflect the official views of the Swiss Government.**
- **This work is funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 644399 (MONROE) through the open call project PREC. The views expressed are solely those of the author(s).**
- **This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 644334 (NEAT). The views expressed are solely those of the author(s)**



PATHspider Tool (EU MAMI Project)

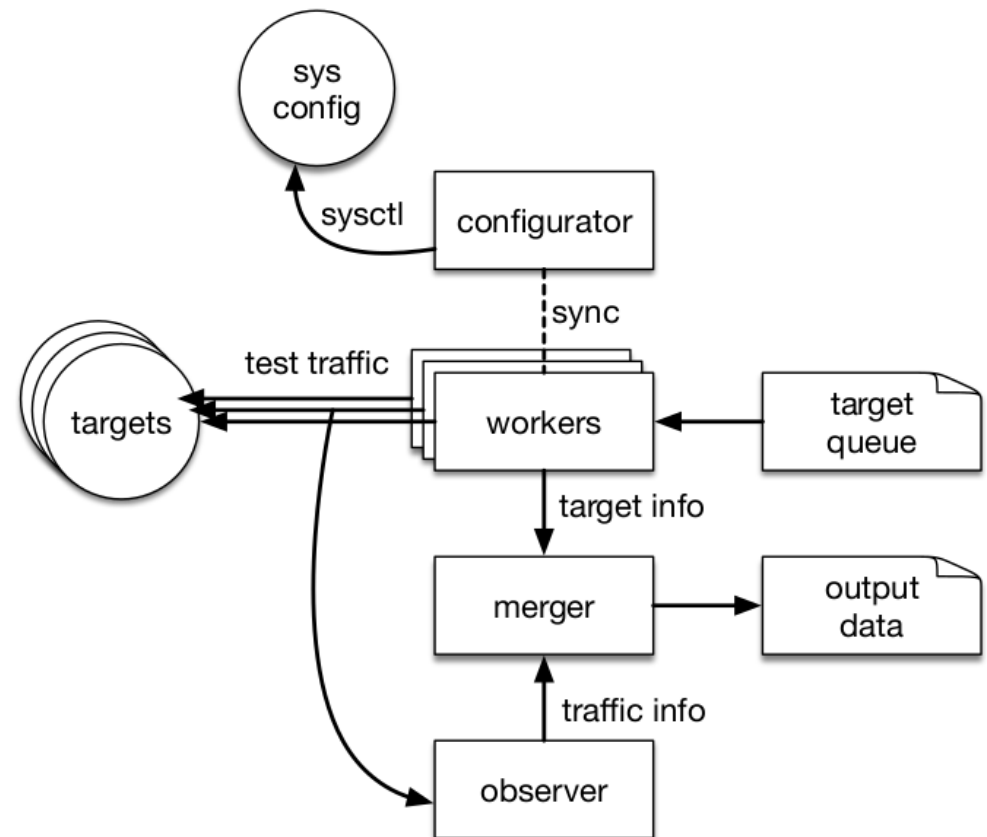
- Tool to perform path transparency measurements
- Entire Alexa Top 1 million website list resolved to IP address
- IP addresses tested for A/E connectivity



<https://pathspider.net/>

PATHspider Measurements

- Tested for A/B connectivity using DSCP 2
- A connection is considered successful if the TCP 3-way handshake completed for *both* DSCP 0 and DSCP 2
- In total, over 4.2 million paths tested



<https://pathspider.net/>

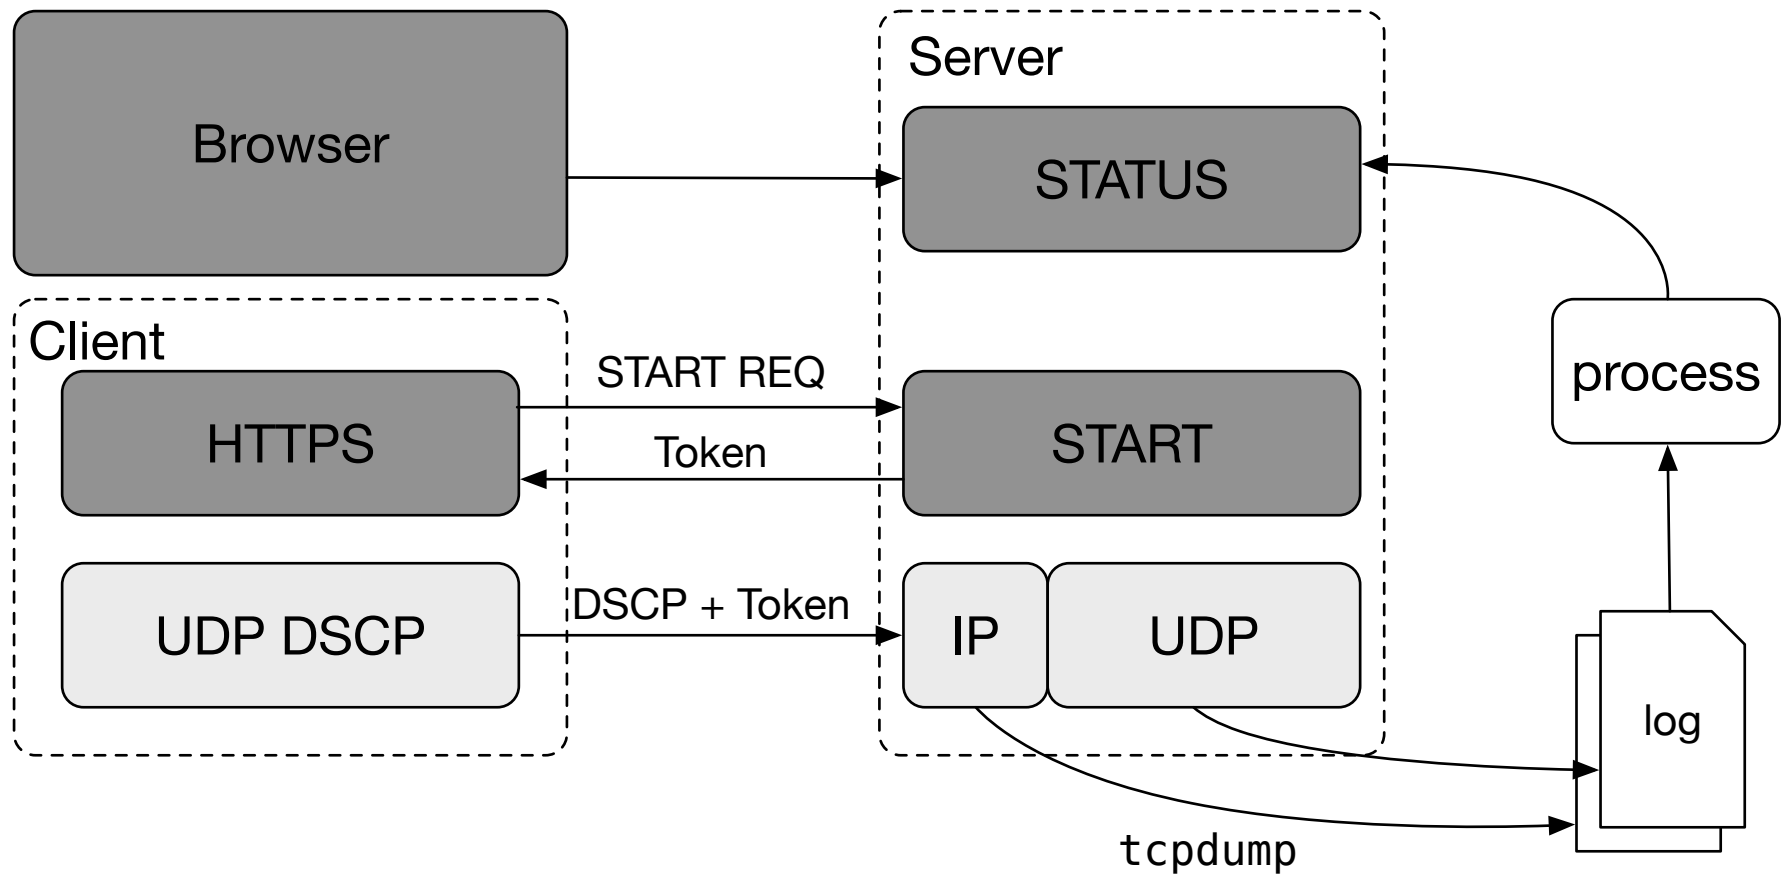


Other known usage (backwards compatibility)

- SSH code base still uses ToS semantics:
 - `#define IPTOS_LOWDELAY 0x10`
 - `#define IPTOS_THROUGHPUT 0x08`
 - `if (options->ip_qos_interactive == -1)`
 - `options->ip_qos_interactive = IPTOS_LOWDELAY;`
 - `if (options->ip_qos_bulk == -1)`
 - `options->ip_qos_bulk = IPTOS_THROUGHPUT;`
- Current DSCP used interactive and bulk data SSH sessions are 2 (000010) and 4 (000110).
 - This could also result in priority inversion
 - Portable OpenSSH bug number 1856: https://bugzilla.mindrot.org/show_bug.cgi?id=1856
 - Bug recommending replacing ToS options for SSH interactive sessions and SCP with AF21 and AF11

Edge trace

- An easy to install that tests which DSCPs are usable from your current location



Edgetrace Software Stack

- **Client**
 - go command line tool
 - runs everywhere™
- **Server**
 - nginx front end
 - go server



Edgetrace Operation


- Client requests session token from server
- Adds to Token:
 - Send DSCP
 - OS
- Sends
 - 10 datagrams per DSCP send
 - Paced 5 packets per second



neat



Observed DSCP at end of path



[Detailed plots](#)

Windows



Binary1

Binary2

Binary3

MAC



Binary1

Binary2

Binary3

Linux



Binary1

Binary2

Binary3

BSD



Binary1

Binary2

Binary3

DSCPs tested by edgetrace

DF EF (LBE)	CS0	AF11	AF31
	CS1	AF12	AF32
	CS2	AF13	AF33
	CS3		
	CS4	AF21	AF41
	CS5	AF22	AF42
	CS6	AF23	AF43



JSON blob returned from the request to the edgetrace server

```
{  
  "host": "139.133.204.55",  
  "time": "20170321090918",  
  "token": "475fc335577904615f3f2299e40c8c84763279fe",  
  "dscp": 0,  
  "description": "ERG TestRun",  
  "os": "darwin"  
}
```

