

# Protocol Identification in Zeek

Jan Grashöfer

Zeek 7.2  
Version



## radius.log | RADIUS authentication attempts

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp for when event happened
uid & id	string	Underlying connection info - See conn.log
username	string	Username, if present
mac	string	MAC address, if present
framed_addr	addr	Address given to network access server, if present
tunnel_client	string	Address (IPv4, IPv6, or FQDN) of initiator end of tunnel, if present
connect_info	string	Connect info, if present
reply_msg	string	Reply message from server challenge
result	string	Successful or failed authentication
ttl	interval	Duration between first request and either Access-Accept message or an error

## ssl.log

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp for when event happened
uid & id	string	Underlying connection info - See conn.log
version	string	SSL/TLS version
cipher	string	Cipher suite
server_name	string	Server name
curve	string	Curve
server_name	string	Server name
last_alert	string	Last alert
next_protocol	string	Next protocol
established	bool	Established
cert_chain_f	string	Certificate chain

## sip.log | SIP analysis

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp when request happened
uid & id	string	Underlying connection info - See conn.log
trans_depth	count	Pipelined depth into request/response transaction
method	string	Verb used in SIP request (INVITE, etc)
uri	string	URI used in request
date	string	Contents of Date: header from client
request_from	string	Contents of request from: header
response_to	string	Contents of response from: header
response_to	string	Contents of response to: header
reply_to	string	Contents of Reply-To: header from client
call_id	string	Contents of Call-ID: header from client
seq	string	Contents of CSeq: header from client
subject	string	Contents of Subject: header from client
request_path	vector	Client message transmission path, extracted from headers
response_path	vector	Server message transmission path, extracted from headers
user_agent	string	Contents of User-Agent: header from client
status_code	string	Status code returned by server
status_msg	string	Status message returned by server
warning	string	Contents of Warning: header
request_body_len	count	Contents of Content-Length: header from client
response_body_len	count	Contents of Content-Length: header from server
content_type	string	Contents of Content-Type: header from server

## dpd.log | Dynamic protocol detection failures

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp when protocol analysis failed
uid & id	string	Underlying connection info - See conn.log

## conn.log | IP, TCP, UDP, ICMP connection details

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp of first packet
uid	string	Unique identifier of connection
id	record	Connection's 4-tuple of endpoint addresses
> id.orig_h	addr	IP address of system initiating connection
> id.orig_p	port	Port from which the connection is initiated
> id.resp_h	addr	IP address of system responding to connection request
> id.resp_p	port	Port on which connection response is sent
proto	enum	Transport layer protocol of connection
service	string	Application protocol ID sent over connection
duration	interval	How long connection lasted
orig_bytes	count	Number of payload bytes originator sent
resp_bytes	count	Number of payload bytes responder sent
conn_state	string	Connection state (see conn.log - conn_state)
local_orig	bool	Value=T if connection originated locally
local_resp	bool	Value=T if connection responded locally
missed_bytes	count	Number of bytes missed (packet loss)
history	string	Connection state history (see conn.log - history)
orig_pkts	count	Number of packets originator sent
orig_ip_bytes	count	Number of originator IP bytes (via IP total length header field)
resp_pkts	count	Number of packets responder sent
resp_ip_bytes	count	Number of responder IP bytes (via IP total length header field)
tunnel_parents	table	If tunneled, connection UID value of encapsulating parent(s)
orig_l2_addr	string	Link-layer address of originator
resp_l2_addr	string	Link-layer address of responder
vlan	int	Outer VLAN for connection
inner_vlan	int	Inner VLAN for connection

## kerberos.log | Kerberos authentication

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp for when event happened
uid & id	string	Underlying connection info - See conn.log
request_type	string	Authentication Service (AS) or Ticket Granting Service (TGS)
client	string	Client
service	string	Service
success	bool	Request result

## dhcp.log | DHCP lease activity

FIELD	TYPE	DESCRIPTION
ts	time	Earliest Time DHCP message observed
client_addr	addr	Unique identifiers of DHCP connections
server_addr	addr	IP address of client
mac	addr	IP address of server hardware

## ntlm.log | NT LAN Manager (NTLM)

### http.log | HTTP request/reply details

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp for when request happened
uid & id	string	Underlying connection info - See conn.log
trans_depth	count	Pipelined depth into connection
host	string	Verb used in HTTP request (GET, POST, etc.)
uri	string	Value of Host header
referrer	string	URI used on request
version	string	Value of version header
user_agent	string	Value of User-Agent portion of request
origin	string	Value of Origin header from client
request_body_len	count	Uncompressed data size from client
response_body_len	count	Uncompressed data size from server
status_code	count	Status code returned by server
status_msg	string	Status message returned by server
info_code	count	Last seen 1xx info reply code from server
info_msg	string	Last seen 1xx info reply message from server
tags	table	Indicators of basic-auth performed for request
username	string	Username of various authentications discovered
password	string	Password of basic-auth performed for request
provided	table	All headers indicative of proxied request
orig_fuids	vector	Ordered vector of the unique IDs
orig_filenames	vector	Ordered vector of filenames from client
orig_mime_types	vector	Ordered vector of mime types
resp_fuids	vector	Ordered vector of file unique IDs
resp_filenames	vector	Ordered vector of filenames from server
resp_mime_types	vector	Ordered vector of mime types
headers	vector	Vector of HTTP header names sent by client
server_header_names	vector	Vector of HTTP header names sent by server
cookie_vars	vector	Variable names extracted from all cookies
url_vars	vector	Variable names from URI

## smtp.log | SMTP transactions

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp when message was first seen
uid & id	string	Underlying connection info - See conn.log
trans_depth	count	Transaction depth if there are multiple msgs
hello	string	Contents of Hello header
mailfrom	string	Email addresses found in From header
rcptto	table	Contents of Rcpt header
date	string	Contents of Date header
from	string	Contents of From header
cc	table	Contents of Cc header
reply_to	table	Contents of Reply-To header
msg_id	string	Contents of Message-ID header
in_reply_to	string	Contents of In-Reply-To header
subject	string	Contents of Subject header
x_originating_ip	addr	Contents of X-Originating-IP header
first_received	string	Contents of first received header
second_received	string	Contents of second received header
last_reply	string	Last message server sent to client
path	vector	Message transmission path, from headers
user_agent	string	Value of User-Agent header from headers
tls	bool	Indicates connection switched to using TLS
fuids	vector	File unique IDs attached to message (if message sent via webmail)
is_webmail	bool	

## irc.log | IRC communication details

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp when command seen
uid & id	string	Underlying connection info - See conn.log
nick	string	Nickname given for connection
user	string	Username given for connection
command	string	Command given by client
val	string	Value for command given by client
addl	string	Any additional data for command
cc_file_size	string	DCC filename requested
cc_mime_type	string	DCC transfer size as indicated by sender
cc_mime_type	string	Sniffed mime type of file
id	string	File unique ID

## snmp.log | SNMP messages

FIELD	TYPE	DESCRIPTION
ts	time	Timestamp of first packet of SNMP session
uid & id	string	Underlying connection info - See conn.log
duration	interval	Adjour of time between first packet belonging to SNMP session and latest seen
version	string	Version of SNMP being used
community	string	Community string of first SNMP packet associated with session
get_requests	count	Number of variable bindings in GetRequest/GetNextRequest PDUs seen for session
set_requests	count	Number of variable bindings in SetRequest/Response PDUs seen for session
get_responses	count	Number of variable bindings in GetResponse/Response PDUs seen for session
set_responses	count	Number of variable bindings in SetResponse/Response PDUs seen for session
display_string	string	System description of SNMP responder endpoint
up_time	time	Time at which SNMP responder endpoint came it's been up since

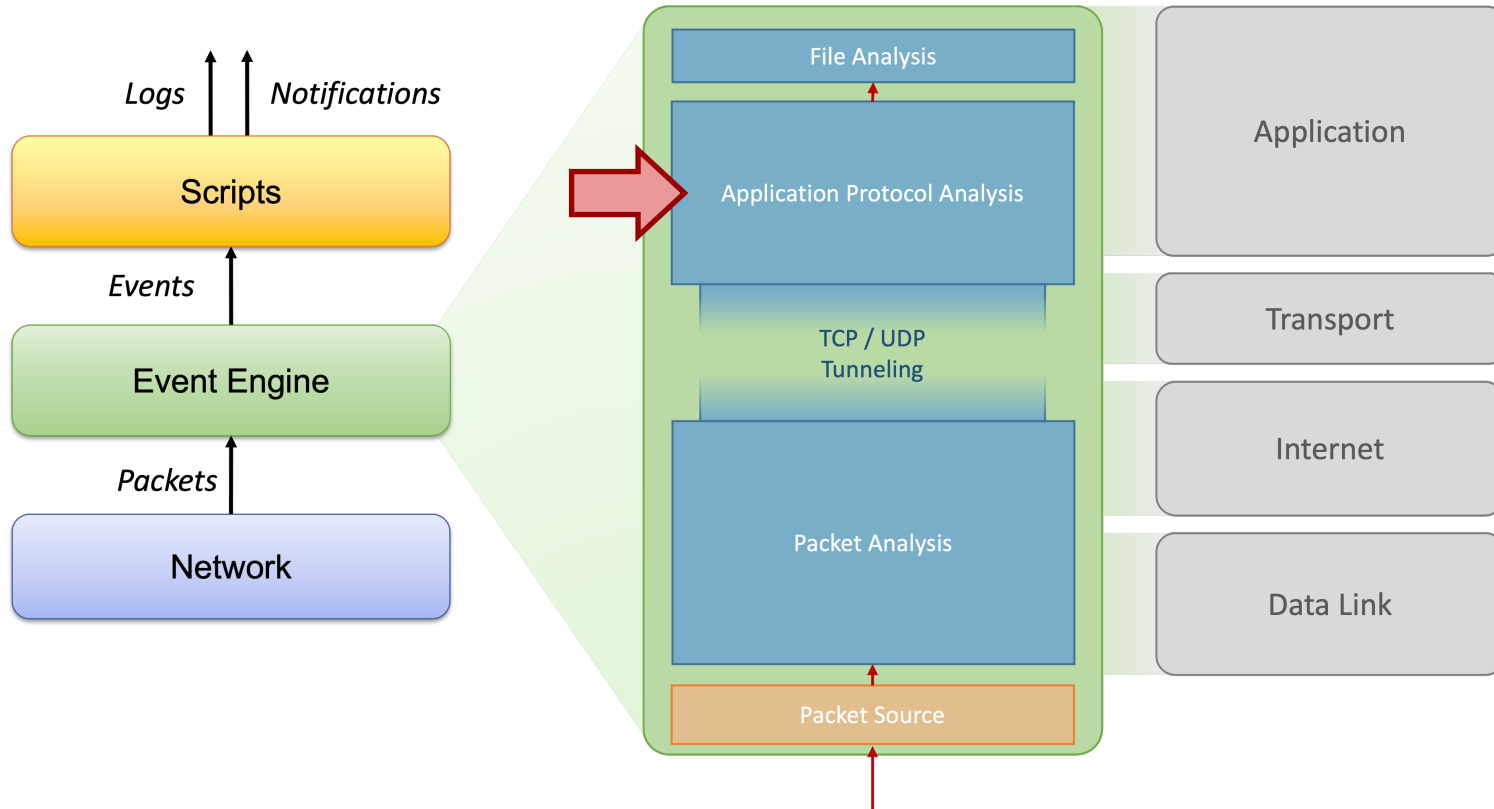
FIELD	TYPE	DESCRIPTION
data_channel	record	Expected FTP data channel
ts	time	Timestamp of event
uid & id	string	Underlying connection info - See conn.log
expected_data	string	Expected Data Channel
actual_data	string	Actual Data Channel
fuid	string	File unique ID

## conn\_state

- A summarized state for each connection
- S0 Connection attempt seen, no reply
  - S1 Connection established, not terminated (0 byte counts)
  - SF Normal establish & termination (>0 byte counts)
  - REJ Connection attempt rejected
  - S2 Established, Orig attempts close, no reply from Resp
  - S3 Established, Resp attempts close, no reply from Orig
  - RSTO Established, Orig aborted (RST)
  - RSTR Established, Resp aborted (RST)
  - RSTOS Orig sent SYN then RST; no Resp SYN-ACK
  - RSTRH Resp sent SYN-ACK then RST; no Orig SYN
  - SHR Orig sent SYN then FIN; no Resp SYN-ACK ("half-open")
  - SHR No SYN, not closed, Midstream traffic.
  - OTH Partial connection.

## history

- Orig uppercase, Resp lowercase, compressed
- S A SYN without the ACK bit set
  - H A SYN-ACK ("handshake")
  - A A pure ACK
  - D A packet with payload ("data")
  - F Packet with FIN bit set
  - R Packet with RST bit set
  - C Packet with a bad checksum
  - I Inconsistent packet (Both SYN & RST)
  - Q Multi-flag packet (SYN & FIN or SYN + RST)
  - T Retransmitted packet
  - W Packet with zero window advertisement
  - A Flipped connection



```
{  
  "ts": 629503200,  
  "uid": "CJKFoj4bpHEhTearRoj",  
  "id.orig_h": "10.0.0.1",  
  "id.orig_p": 51889,  
  "id.resp_h": "192.168.0.1",  
  "id.resp_p": 80,  
  "proto": "tcp",  
  "service": "http",  
  "duration": 0.0002599954605102539,  
  "orig_bytes": 18,  
  "resp_bytes": 12649,  
  "conn_state": "SF",  
  "local_orig": true,  
  "local_resp": true,  
  "missed_bytes": 0,  
  "history": "ShADadFf",  
  "orig_pkts": 15,  
  "orig_ip_bytes": 618,  
  "resp_pkts": 13,  
  "resp_ip_bytes": 13169  
}
```

The screenshot shows a network traffic analysis tool interface. The top part displays a list of packets with columns for No., Time, Source, SrcP, Destination, DstP, Protocol, Length, and Info. The bottom part shows a detailed view of a selected frame, including Ethernet II, Internet Protocol Version 4, and Transmission Control Protocol details.

No.	Time	Source	SrcP	Destination	DstP	Protocol	Length	Info
1	629503200.000000	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [SYN] Seq=0 Win=8192 Len=0
2	629503200.000010	192.168.0.1	80	10.0.0.1	51889	TCP	54	80 → 51889 [SYN, ACK] Seq=0 Ack=1 Win=8192 Len=0
3	629503200.000020	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=1 Ack=1 Win=8192 Len=0
4	629503200.000030	10.0.0.1	51889	192.168.0.1	80	HTTP	72	GET / HTTP/1.1
5	629503200.000040	192.168.0.1	80	10.0.0.1	51889	TCP	54	80 → 51889 [ACK] Seq=1 Ack=19 Win=8192 Len=0
6	629503200.000050	192.168.0.1	80	10.0.0.1	51889	TCP	85	80 → 51889 [ACK] Seq=1 Ack=19 Win=8192 Len=31 [TCP
7	629503200.000060	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=32 Win=8192 Len=0
8	629503200.000070	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=32 Ack=19 Win=8192 Len=1402
9	629503200.000080	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=1434 Win=8192 Len=0
10	629503200.000090	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=1434 Ack=19 Win=8192 Len=1402
11	629503200.000100	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=2836 Win=8192 Len=0
12	629503200.000110	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=2836 Ack=19 Win=8192 Len=1402
13	629503200.000120	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=4238 Win=8192 Len=0
14	629503200.000130	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=4238 Ack=19 Win=8192 Len=1402
15	629503200.000140	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=5640 Win=8192 Len=0
16	629503200.000150	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=5640 Ack=19 Win=8192 Len=1402
17	629503200.000160	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=7042 Win=8192 Len=0
18	629503200.000170	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=7042 Ack=19 Win=8192 Len=1402
19	629503200.000180	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=8444 Win=8192 Len=0
20	629503200.000190	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=8444 Ack=19 Win=8192 Len=1402
21	629503200.000200	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=9846 Win=8192 Len=0
22	629503200.000210	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=9846 Ack=19 Win=8192 Len=1402
23	629503200.000220	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=11248 Win=8192 Len=0
24	629503200.000230	192.168.0.1	80	10.0.0.1	51889	TCP	1456	80 → 51889 [ACK] Seq=11248 Ack=19 Win=8192 Len=1402
25	629503200.000240	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=19 Ack=12650 Win=8192 Len=0
26	629503200.000250	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [FIN, ACK] Seq=19 Ack=12650 Win=8192 Len=0
27	629503200.000260	192.168.0.1	80	10.0.0.1	51889	HTTP	54	HTTP/1.1 200 OK
28	629503200.000270	10.0.0.1	51889	192.168.0.1	80	TCP	54	51889 → 80 [ACK] Seq=20 Ack=12651 Win=8192 Len=0

Frame 1: 54 bytes on wire (432 bits), 54 bytes captured (432 bits)  
> Ethernet II, Src: 02:aa:00:00:00:01 (02:aa:00:00:00:01), Dst: 02:bb:00:00:00:00 (02:bb:00:00:00:00)  
> Internet Protocol Version 4, Src: 10.0.0.1, Dst: 192.168.0.1  
> Transmission Control Protocol, Src Port: 51889, Dst Port: 80, Seq: 0, Len: 0

0000 02 bb 00 00 00 00 02 aa 00 00 00 01 08 00 45 00 .....E:  
0010 00 28 00 01 00 00 40 06 b0 25 0a 00 00 01 c0 a8 .....@:~&.....  
0020 00 01 ca b1 0c 50 00 00 03 e8 00 00 00 00 50 02 .....P.....P:  
0030 20 00 f6 4c 00 00 .....N.....

# Protocol in conn.log – What changed?

```
{
  "ts": 629503200,
  "uid": "CJKFoj4bpHEhTearRoj",
  "id.orig_h": "10.0.0.1",
  "id.orig_p": 51889,
  "id.resp_h": "192.168.0.1",
  "id.resp_p": 2323,
  "proto": "tcp",
  "service": "http",
  "duration": 0.0002599954605102539,
  "orig_bytes": 18,
  "resp_bytes": 12649,
  "conn_state": "SF",
  "local_orig": true,
  "local_resp": true,
  "missed_bytes": 0,
  "history": "ShADadFf",
  "orig_pkts": 15,
  "orig_ip_bytes": 618,
  "resp_pkts": 13,
  "resp_ip_bytes": 13169
}
```

No.	Time	Source	SrcP	Destination	DstP	Protocol	Length	Info
1	629503200.000000	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [SYN] Seq=0 Win=8192 Len=0
2	629503200.000010	192.168.0.1	2323	10.0.0.1	51889	TCP	54	2323 → 51889 [SYN, ACK] Seq=0 Ack=1 Win=8192 Len=0
3	629503200.000020	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=1 Ack=1 Win=8192 Len=0
4	629503200.000030	10.0.0.1	51889	192.168.0.1	2323	HTTP	72	GET / HTTP/1.1
5	629503200.000040	192.168.0.1	2323	10.0.0.1	51889	TCP	54	2323 → 51889 [ACK] Seq=1 Ack=19 Win=8192 Len=0
6	629503200.000050	192.168.0.1	2323	10.0.0.1	51889	TCP	85	2323 → 51889 [ACK] Seq=1 Ack=19 Win=8192 Len=31 [RST] Win=0
7	629503200.000060	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=32 Win=8192 Len=0
8	629503200.000070	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=32 Ack=19 Win=8192 Len=1402
9	629503200.000080	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=1434 Win=8192 Len=0
10	629503200.000090	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=1434 Ack=19 Win=8192 Len=14
11	629503200.000100	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=2836 Win=8192 Len=0
12	629503200.000110	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=2836 Ack=19 Win=8192 Len=14
13	629503200.000120	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=4238 Win=8192 Len=0
14	629503200.000130	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=4238 Ack=19 Win=8192 Len=14
15	629503200.000140	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=5640 Win=8192 Len=0
16	629503200.000150	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=5640 Ack=19 Win=8192 Len=14
17	629503200.000160	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=7042 Win=8192 Len=0
18	629503200.000170	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=7042 Ack=19 Win=8192 Len=14
19	629503200.000180	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=8444 Win=8192 Len=0
20	629503200.000190	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=8444 Ack=19 Win=8192 Len=14
21	629503200.000200	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=9846 Win=8192 Len=0
22	629503200.000210	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=9846 Ack=19 Win=8192 Len=14
23	629503200.000220	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=11248 Win=8192 Len=0
24	629503200.000230	192.168.0.1	2323	10.0.0.1	51889	TCP	1456	2323 → 51889 [ACK] Seq=11248 Ack=19 Win=8192 Len=14
25	629503200.000240	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=19 Ack=12650 Win=8192 Len=0
26	629503200.000250	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [FIN, ACK] Seq=19 Ack=12650 Win=8192 Len=0
27	629503200.000260	192.168.0.1	2323	10.0.0.1	51889	HTTP	54	HTTP/1.1 200 OK
28	629503200.000270	10.0.0.1	51889	192.168.0.1	2323	TCP	54	51889 → 2323 [ACK] Seq=20 Ack=12651 Win=8192 Len=0



```
{  
  "ts": 629503200,  
  "uid": "CJKFoj4bpHEhTearoj",  
  "id.orig_h": "10.0.0.1",  
  "id.orig_p": 51889,  
  "id.resp_h": "192.168.0.1",  
  "id.resp_p": 2323,  
  "proto": "tcp",  
  "service": "http",  
  "duration": 0.0002599954605102539,  
  "orig_bytes": 18,  
  "resp_bytes": 12649,  
  "conn_state": "SF",  
  "local_orig": true,  
  "local_resp": true,  
  "missed_bytes": 0,  
  "history": "ShADadFf",  
  "orig_pkts": 15,  
  "orig_ip_bytes": 618,  
  "resp_pkts": 13,  
  "resp_ip_bytes": 13169,  
  "ip_proto": 6  
}
```

- Dynamic Protocol Detection (DPD)  
→ **Trial and Error** approach
  - Just run every analyzer?  
😓 too resource intensive!
- Attach analyzers based on heuristics
  - **Port**-based  
Well-known/registered ports (e.g., 80 for HTTP)
  - **Content**-based  
Look for expected patterns (e.g., HTTP Methods)

 Dynamic Application-Layer Protocol Analysis for Network Intrusion Detection

- Challenge: Signature sensitivity
  - Pattern accuracy → specific vs. lax
  - Require matches in originator and responder direction? (requires-reverse-signature)

```
signature dpd_http_client {
  ip-proto == tcp
  payload /^[[:space:]]*(OPTIONS|GET|HEAD|POST|PUT
  tcp-state originator
  enable "http"
}

signature dpd_http_server {
  ip-proto == tcp
  payload /^HTTP\[0-9]/
  tcp-state responder
  enable "http"
}
```

<scripts/base/protocols/http/dpd.sig>

```
GET /index.html
```

```
HTTP/1.0 200 OK
Hello World!
```

```
CUSTOMMETHOD /index.html HTTP/1.1
```

```
HTTP/1.1 200 OK
Hello customized World!
```

- Challenge: Signature sensitivity
  - Pattern accuracy → lax vs. specific
  - Require matches in originator and responder direction? (requires-reverse-signature)
- Content buffered to replay on signature match
  - `dpd_buffer_size = 1024`
  - `dpd_max_packets = 100`
  - DPD stops if buffer is exhausted

```
GET /index.html
```

```
HTTP/1.0 200 OK  
Hello World!
```

```
CUSTOMMETHOD /index.html HTTP/1.1
```

```
HTTP/1.1 200 OK  
Hello customized World!
```

```
POST /index.html HTTP/1.1  
Content-Type: application/octet-stream  
Content-Length: 64000
```

```
Data...
```

```
HTTP/1.1 200 OK  
Finally, hello World!
```


```
{
  "ts": 629503200,
  "uid": "CJKFoj4bpHEhTeaRoj",
  "id.orig_h": "10.0.0.1",
  "id.orig_p": 51889,
  "id.resp_h": "192.168.0.1",
  "id.resp_p": 2323,
  "proto": "tcp",
  "service": "http",
  "duration": 0.0002599954605102539,
  "orig_bytes": 18,
  "resp_bytes": 12649,
  "conn_state": "SF",
  "local_orig": true,
  "local_resp": true,
  "missed_bytes": 0,
  "history": "ShADadFf",
  "orig_pkts": 15,
  "orig_ip_bytes": 618,
  "resp_pkts": 13,
  "resp_ip_bytes": 13169,
  "ip_proto": 6
}
```

## • Dynamic Protocol Detection (DPD)

### → Trial and Error approach


- Run all potential analyzers in parallel
- Analyzers run in trial mode until
  - **Confirmation** → `service` field set in `conn.log`
  - **Violation** → Analyzer is detached if `threshold crossed`
- Confirmed analyzers are detached upon violation
  - ~~Protocol removed from `service` in `conn.log`~~
  - Protocol is kept in `service` in `conn.log`

Changed in Zeek 7.2



```
{
  "ts": 629503200,
  "uid": "CJKFoj4bpHEhTeaRoj",
  "id.orig_h": "10.0.0.1",
  "id.orig_p": 51889,
  "id.resp_h": "192.168.0.1",
  "id.resp_p": 2323,
  "proto": "tcp",
  "service": "http",
  "duration": 0.0002599954605102539,
  "orig_bytes": 18,
  "resp_bytes": 12649,
  "conn_state": "SF",
  "local_orig": true,
  "local_resp": true,
  "missed_bytes": 0,
  "history": "ShADadFf",
  "orig_pkts": 15,
  "orig_ip_bytes": 618,
  "resp_pkts": 13,
  "resp_ip_bytes": 13169,
  "ip_proto": 6
}
```

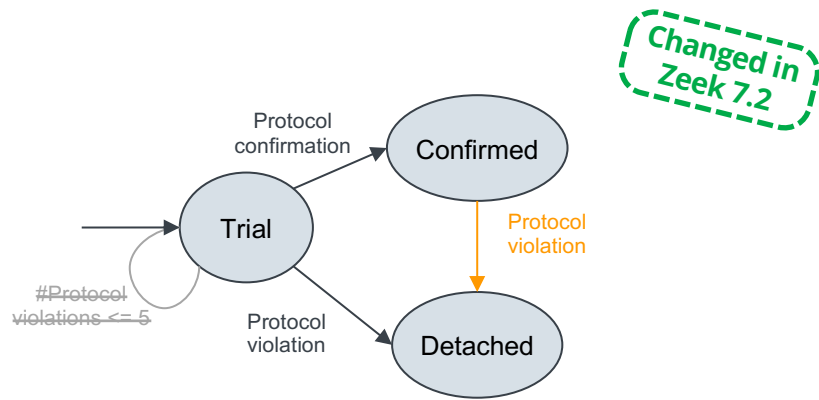
- ⚠ Protocol-specific log generation independent of confirmation
- Configurable cut-offs
  - ~~DPD::max\_violations: table[Analyzer::Tag] of count = table() &default = 5~~
  - DPD::ignore\_violations: set[Analyzer::Tag] = set()
  - DPD::ignore\_violations\_after = 10 \* 1024



Changed in Zeek 7.2


```
{
  "ts": 629503200,
  "uid": "CJKFoj4bpHEhTeaRoj",
  "id.orig_h": "10.0.0.1",
  "id.orig_p": 51889,
  "id.resp_h": "192.168.0.1",
  "id.resp_p": 2323,
  "proto": "tcp",
  "service": "http",
  "duration": 0.0002599954605102539,
  "orig_bytes": 18,
  "resp_bytes": 12649,
  "conn_state": "SF",
  "local_orig": true,
  "local_resp": true,
  "missed_bytes": 0,
  "history": "ShADadFf",
  "orig_pkts": 15,
  "orig_ip_bytes": 618,
  "resp_pkts": 13,
  "resp_ip_bytes": 13169,
  "ip_proto": 6
}
```

- How to tell an analyzer *confirmed* and then got *detached*?
  - **Zeek 5.1 - 7.1:** Protocol won't show up in **service** field
  - **Zeek 7.2:** Protocol remains in **service** field



- Only generated for violations of previously confirmed analyzers
  - ⚠ Analyzer got detached → no further analysis

```
{  
  "ts": 1669822677.033301,  
  "uid": "CJKFoj4bpHEhTeaRoj",  
  "id.orig_h": "192.168.12.5",  
  "id.orig_p": 51792,  
  "id.resp_h": "192.0.78.212",  
  "id.resp_p": 80,  
  "proto": "tcp",  
  "analyzer": "HTTP",  
  "failure_reason": "not a http request line"  
}
```



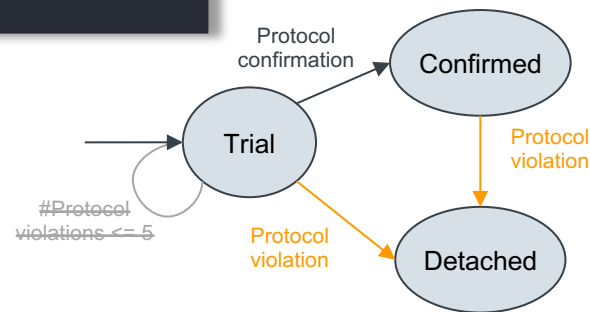
- Generated for all protocol violations for all analyzers (incl. file analyzers)
  - Primarily intended for debugging of analyzers in development phase

```

{
  "ts": 629503200.00015,
  "cause": "violation",
  "analyzer_kind": "protocol",
  "analyzer_name": "SSH",
  "uid": "CJKFoj4bpHEhTeaRoj",
  "fuid":
  "id.orig_h": "10.0.0.1",
  "id.orig_p": 51889,
  "id.resp_h": "192.168.0.1",
  "id.resp_p": 22,
  "failure_reason": "Binpac exception: binpac exception: invalid index for case: SSH_Key_Exchange: 3",
  "failure_data":
}
    
```

Customizable

- Analyzer::Logging::include\_confirmations = F
- Analyzer::Logging::include\_disabling = F
- Analyzer::Logging::failure\_data\_max\_size = 40
- Analyzer::Logging::ignore\_analyzers = set()




```

{
  "ts": 1669822676.914779,
  "uid": "CJKFoj4bpHEhTearoj",
  "id.orig_h": "192.168.12.5",
  "id.orig_p": 51792,
  "id.resp_h": "192.0.78.212",
  "id.resp_p": 80,
  "proto": "tcp",
  1. "service": "http,-http",
  "duration": 0.41950106620788574,
  "orig_bytes": 71,
  "resp_bytes": 699,
  "conn_state": "SF",
  "local_orig": true,
  "local_resp": false,
  "missed_bytes": 0,
  "history": "ShADadfF",
  "orig_pkts": 11,
  "orig_ip_bytes": 531,
  "resp_pkts": 9,
  "resp_ip_bytes": 1071,
  "ip_proto": 6,
  2. "failed_service": [
    "http"
  ]
}

```

- **New in Zeek 7.2:** Keeping track of detached, previously confirmed analyzers in conn.log

1. Add to **service** field using "-" prefix  
 → DPD::track\_removed\_services\_in\_connection = T
2. Add **failed\_service** field  
 → @load protocols/conn/failed-service-logging



## HTTP traffic on port 22 (SSH)

Expectations?

```
{
  "ts": 1705192016.281275,
  "uid": "CJKFoj4bpHEhTeaRoj",
  "id.orig_h": "127.0.0.1",
  "id.orig_p": 42906,
  "id.resp_h": "127.0.0.1",
  "id.resp_p": 8888,
  "proto": "tcp",
  "service": "http,websocket,ssh",
  "duration": 8.016665935516357,
  "orig_bytes": 4117,
  "resp_bytes": 2017,
  "conn_state": "RST0",
  "local_orig": true,
  "local_resp": true,
  "missed_bytes": 0,
  "history": "ShADadR",
  "orig_pkts": 20,
  "orig_ip_bytes": 5165,
  "resp_pkts": 22,
  "resp_ip_bytes": 3169,
  "ip_proto": 6
}
```

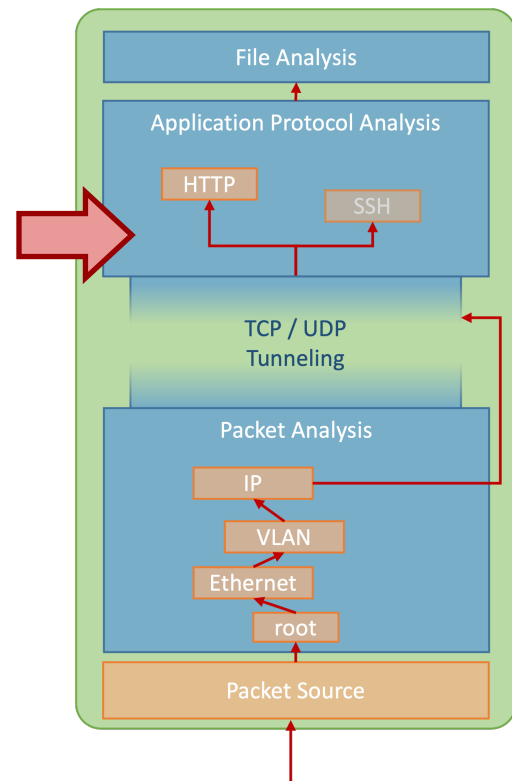
- **New in Zeek 7.2:** The **service** field is ordered  
→ reflects the order in which protocols got **confirmed**
  - Note: Confirmation order does not necessarily reflect nesting of protocol layers

- Continuous matching to determine "speculative service"
  - Buffering is stopped once DPD-buffer is filled, and data gets deleted
  - @load protocols/conn/speculative-service
    - adds column to conn.log
      - dpd\_match\_only\_beginning = F
        - ⚠ applies to all signatures
      - dpd\_late\_match\_stop = T
    - No replay & no analyzer attached (resynchronization required)
    - Independent of previous successful DPD

```
{
  "ts": 629503200,
  "uid": "CJKFoj4bpHEhTeaRoj",
  "id.orig_h": "10.0.0.1",
  "id.orig_p": 51889,
  "id.resp_h": "192.168.0.1",
  "id.resp_p": 2323,
  "proto": "tcp",
  "service": ,
  "duration": 0.0010809898376464844,
  "orig_bytes": 2137,
  "resp_bytes": 68729,
  "conn_state": "SF",
  "local_orig": true,
  "local_resp": true,
  "missed_bytes": 0,
  "history": "ShADadFf",
  "orig_pkts": 56,
  "orig_ip_bytes": 4377,
  "resp_pkts": 54,
  "resp_ip_bytes": 70889,
  "speculative_service": "http"
}
```

## Application Protocol Analysis → Dynamic Protocol Detection

- Trial and error approach using Ports and Signatures
  - Trade-offs can be tuned (buffer capacity, violation limit)
  - `dpd.log` and `analyzer.log` inform about failed attempts
- Zeek implements clever approaches to increase visibility
- Zeek tells you what it “did not understand”



- Logs protocol IDs packet analyzers cannot map  
→ @load policy/misc/unknown-protocols

```
{  
  "ts": 1342606813.729856,  
  "analyzer": "ETHERNET",  
  "protocol_id": "0x8809",  
  "f{  
    }  
  }  
  {  
    "ts": 1725534240.604223,  
    "analyzer": "IP",  
    "protocol_id": "0x58",  
    "first_bytes": "0205e0a9000000000000",  
    "analyzer_history": [  
      "ETHERNET",  
      "VLAN",  
      "IP",  
      "GRE",  
      "IPTUNNEL",  
      "IP"  
    ]  
  }  
}
```

- UnknownProtocol::sampling\_threshold = 3  
→ enter sampling mode if (analyzer, ID) pair was seen 3 times
- UnknownProtocol::sampling\_rate = 100,000  
→ in sampling mode log every 100,000 occurrence of an (analyzer, ID) pair
- UnknownProtocol::sampling\_duration = 1 hr  
→ expire sampling after 1 hr without observation of an (analyzer, ID) pair

← New in Zeek 7.1: Analyzer History