

# Parsing with Spicy

Suricon Athens 2022

Benjamin Banner

Corelight

2022-11-10

# Zeek



<https://zeek.org>

<https://github.com/zeek/zeek>

Zeek taps into network traffic and logs *interesting information*.

- used and deployed in large networks for 20+ years
- core developed in C++
- extensible via its own scripting language, or C++ code
- released under a permissive license

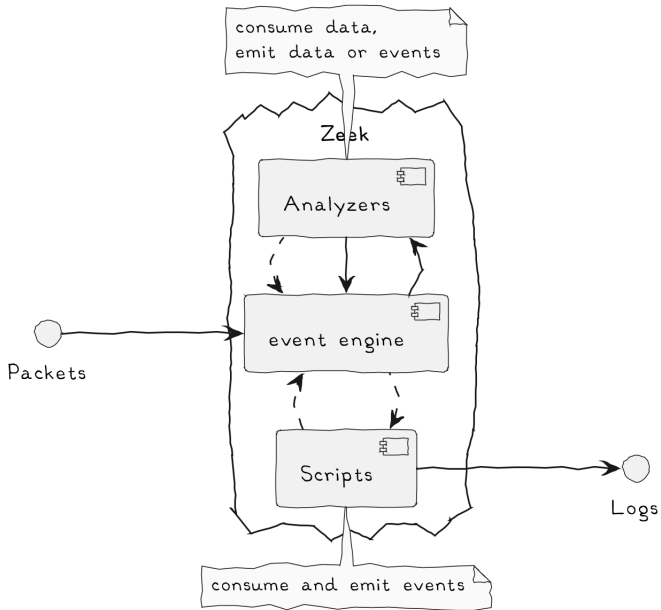


Figure 1: Simplified Zeek protocol analysis workflow

## Adding new analyzers to Zeek

- historically required C++ implementation
- C++ has a significant barrier of entry
- code can be pretty low level hiding business logic
- C++ is not a safe language

# High-level scripting language for parsing: Spicy



## Goals

- robust language with support for both syntax and logic (no need for C++!)
- efficient enough to parse many concurrent connections with low latency and memory overhead
- transparent integration in Zeek

## Project

- started as ICSI research prototype
- actively used in the Zeek ecosystem to write packet, protocol and file analyzers
- <https://github.com/zeek/spicy>, #spicy

# Trivial File Transfer Protocol (TFTP)

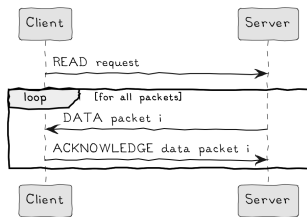


Figure 2: Read




Figure 3: Write

*# Figure 5-1: RRQ/WRQ packet*

*#*

*# 2 bytes      string      1 byte      string      1 byte*

*# -----*

*# | Opcode | Filename | 0 | Mode | 0 |*

*# -----*

```
public type Packet = unit {
```

```
  op: uint16
```

```
  rrq: ReadRequest;
```

```
  # Hook triggered once this unit has finished parsing.
```

```
  on %done { print self; }
```

```
};
```

```
type ReadRequest = unit {
```

```
  filename: bytes &until=b"\x00";
```

```
  mode:      bytes &until=b"\x00";
```

```
};
```

```

# Figure 5-1: RRQ/WRQ packet
#
# 2 bytes      string      1 byte      string      1 byte
# -----
# | Opcode | Filename | 0 | Mode | 0 |
# -----

```

```

public type Packet = unit {
  op: uint16 &convert=Opcode($$);
  switch ( self.op ) {
    Opcode::RRQ   -> rrq:   Request(True);
    Opcode::WRQ   -> wrq:   Request(False);
    # TODO: Add handling for other opcodes.
  };

  # Hook triggered once this unit has finished parsing.
  on %done { print self; }
};

type Request = unit(is_read: bool) { # Parameter is implicitly `in`.
  filename: bytes &until=b"\x00";
  mode:     bytes &until=b"\x00";
};

type Opcode = enum { RRQ = 0x01, WRQ = 0x02, DATA = 0x03,
  ACK = 0x04, ERROR = 0x05 };

```



## Recovering from input gaps/parse errors

```
type X = unit {  
  a: b"begin-of-X";  
  b: bytes &size=10;  
};
```

```
type Foo = unit {  
  xs: (X &synchronize)[];  
};
```

.....

```
type Bar = unit {  
  a: A;  
  b: B;  
  c: C &synchronize;  
  d: D;  
};
```

More low-level manual backtracking is also supported.

## Reshaping data

```
# csv_naive.spicy
module csv_naive;

public type CSV = unit {
    rows: Row[];
};

type Row = unit {
    cols: bytes &until=b"\n" &convert=$$.split(b",");
} &convert=self.cols;

on CSV::%done {
    print self;
}

.....

$ printf '1,2,3\nA,B,C\n' | spicy-driver csv_naive.spicy
[$rows=[[b"1", b"2", b"3"], [b"A", b"B", b"C"]]]
```

# Extensibility

```
// mylibrary.cc  
namespace MyLibrary {  
std::string rotl3(const std::string& in) {  
    // ...  
} // MyLibrary
```

```
# mylibrary.spicy  
module MyLibrary;
```

```
public function rotl3(s: string) : string &cxxname="MyLibrary::rotl3";
```

```
# example.spicy  
module Example;  
import MyLibrary;  
const data = "Hello, world!";  
print "%s' -> '%s' -> '%s'" % (data,  
                               MyLibrary::rotl3(data),  
                               MyLibrary::rotl3(MyLibrary::rotl3(data)));
```

```
.....  
$ spicyc -j rotl3.spicy mylibrary.spicy mylibrary.cc  
'Hello, world!' -> 'Uryyb, jbeyq!' -> 'Hello, world!'
```

# Spicy processing pipeline

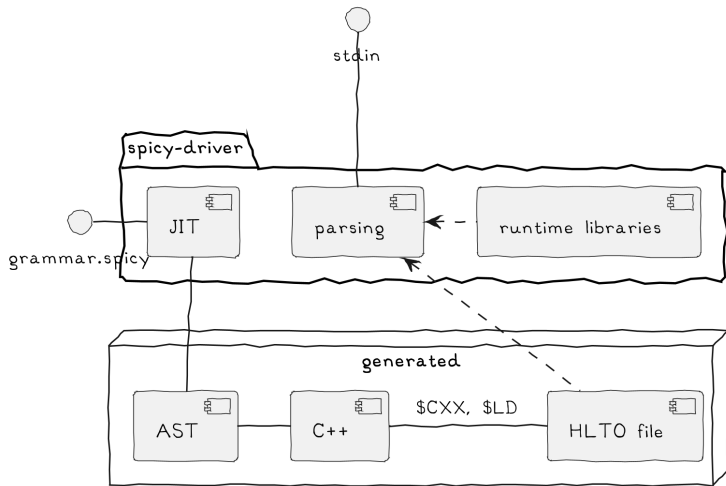


Figure 4: spicy-driver workflow

```
# foo.spicy
# ...
public type X = unit {
  a: uint8;
  b: bytes &eod;
};
```

```
.....

$ spicyc -c foo.spicy
namespace <internal> {
struct X: // ... {
  // ...
  static auto parse3(::hilti::rt::ValueReference<::spicy::rt::ParsedUnit>& gunit,
                    ::hilti::rt::ValueReference<::hilti::rt::Stream>& data,
                    const std::optional<::hilti::rt::stream::View>& cur,
                    const std::optional<::spicy::rt::UnitContext>& context)
    -> ::hilti::rt::stream::View;

  std::optional<::hilti::rt::integer::safe<uint8_t>> a{};
  std::optional<::hilti::rt::Bytes> b{};
};
}
```

*ParsedUnit* can be reflected on.

ttf\_rrq.pcap

Apply a display filter ... <⌘/>

No.	Time	Source	Destination	Protoc	Length	Info
1	0.000000	192.168.0.253	192.168.0.10	TFTP	62	Read Request, File: rfc1350.txt, Transfer t
2	0.104391	192.168.0.10	192.168.0.253	TFTP	558	Data Packet, Block: 1
3	0.108938	192.168.0.253	192.168.0.10	TFTP	60	Acknowledgement, Block: 1
4	0.113448	192.168.0.10	192.168.0.253	TFTP	558	Data Packet, Block: 2
5	0.116109	192.168.0.253	192.168.0.10	TFTP	60	Acknowledgement, Block: 2
6	0.116143	192.168.0.10	192.168.0.253	TFTP	558	Data Packet, Block: 3
7	0.118794	192.168.0.253	192.168.0.10	TFTP	60	Acknowledgement, Block: 3
8	0.118823	192.168.0.10	192.168.0.253	TFTP	558	Data Packet, Block: 4
9	0.121531	192.168.0.253	192.168.0.10	TFTP	60	Acknowledgement, Block: 4
10	0.121564	192.168.0.10	192.168.0.253	TFTP	558	Data Packet, Block: 5
11	0.124141	192.168.0.253	192.168.0.10	TFTP	60	Acknowledgement, Block: 5

▶ Frame 7: 60 bytes on wire (480 bits), 60 bytes captured (480 bits)

▶ Ethernet II, Src: Cisco\_18:9a:40 (00:0b:be:18:9a:40), Dst: AbitComp\_d7:8b:43 (00:50:8d:d7:8b:43)

▶ Internet Protocol Version 4, Src: 192.168.0.253, Dst: 192.168.0.10

▶ User Datagram Protocol, Src Port: 50618, Dst Port: 3445

▼ Trivial File Transfer Protocol

  Opcode: Acknowledgement (4)

  [Source File: rfc1350.txt]

  Block: 3

```

0000  00 50 8d d7 8b 43 00 0b  be 18 9a 40 08 00 45 00  .P...C...@..E.
0010  00 20 00 03 00 00 ff 11  39 72 c0 a8 00 fd c0 a8  ..9r.....
0020  00 0a c5 ba 0d 75 00 0c  aa 47 00 04 00 03 00 00  ...u..G...
0030  00 00 00 00 00 00 00 00  00 00 00 00
  
```

Block number (tftp.block), 2 bytes

Packets: 99 · Displayed: 99 (100.0%)

Profile: Default

## Features not touched on

- standard library of frequently used types
- *preprocessing* data through filter abstraction
- support for *reassembly* with sink abstraction
- unit context to attach additional information (e.g., to *share state between both sides of connection*)
- *look-ahead parsing* for variant parsing, lists of unknown size, ...
- *control data* a field is parsed from with attributes
- *random access* support to completely change input position ...

## Use in Zeek

- Spicy knows nothing about Zeek
- Zeek-specific *spicy-plugin* glue layer adapts types and matches up Spicy hooks with Zeek events<sup>1</sup>:

```
# tftp.evt
protocol analyzer spicy::TFTP over UDP:
    parse with TFTP::Packet,
    port 69/udp;

import TFTP;

on TFTP::Request if ( is_read )    -> event tftp::read_request(
                                    $conn, $is_orig, self.filename, self.mode);
on TFTP::Request if ( ! is_read ) -> event tftp::write_request(
                                    $conn, $is_orig, self.filename, self.mode);
# ...
```

- parser developers typically write a Spicy grammar and an EVT file starting from a template, and implement additional analyses in Zeek scripts

---

<sup>1</sup>Work on automatically exposing Spicy types to Zeek underway.



# Sample STUN analyzer

<https://github.com/corelight/zeek-spicy-stun>

```
$ for f in *(spicylevtlzeeklsig); do echo $f $(rg -v '^(#I$)' $f | wc -l); done | column -t
analyzer.spicy      160
analyzer.evt       70
zeek_analyzer.spicy 9
dpd.sig            10
consts.zeek        79
main.zeek          179
__load__.zeek      3
```

# Sample STUN analyzer

<https://github.com/corelight/zeek-spicy-stun>

```
const MAGIC_COOKIE = /\x21\x12\xA4\x42/;
public type STUNPacket = unit {
    var M: uint16;
    var C: uint16;

    msgtype: bitfield(16) {
        m0: 0..3;
        c0: 4;
        m1: 5..7;
        c1: 8;
        m2: 9..13;
        zeros: 14..15 &requires=($$ == 0);
    } {
        self.M = self.msgtype.m0 + 16*self.msgtype.m1 + 128*self.msgtype.m2;
        self.C = self.msgtype.c0 + 2*self.msgtype.c1;
    }

    msg_length: uint16;
    : MAGIC_COOKIE;
    trans_id: bytes &size=12;
    attributes: Attribute(self.M, self.C, self.trans_id)[];
};
```

Spicy successfully enabled more users to contribute parsers to Zeek.  
We will gradually move bundled Zeek parsers to Spicy.  
Consider starting your Zeek parser from Spicy *zkg* templates.

Let's discuss whether we as a community could use Spicy to share parsers across different consumers.



<https://github.com/zeek/spicy>  
#spicy