



HILTI: An Abstract Execution Environment for Deep, Stateful Network Traffic Analysis

Robin Sommer

ICSI / LBNL

Matthias Vallentin

UC Berkeley

Lorenzo De Carli

University of Wisconsin-Madison

Vern Paxson

ICSI / UC Berkeley

A Tale of Three Open-Source IDS



Shared
functionality?



A Tale of Three Open-Source IDS



Shared
functionality?
Essentially none.



A Tale of Three Open-Source IDS



Shared
functionality?
Essentially none.



Same for packet filters, firewalls, proxies, routers, switches, OS stack ...

DPI Architecture

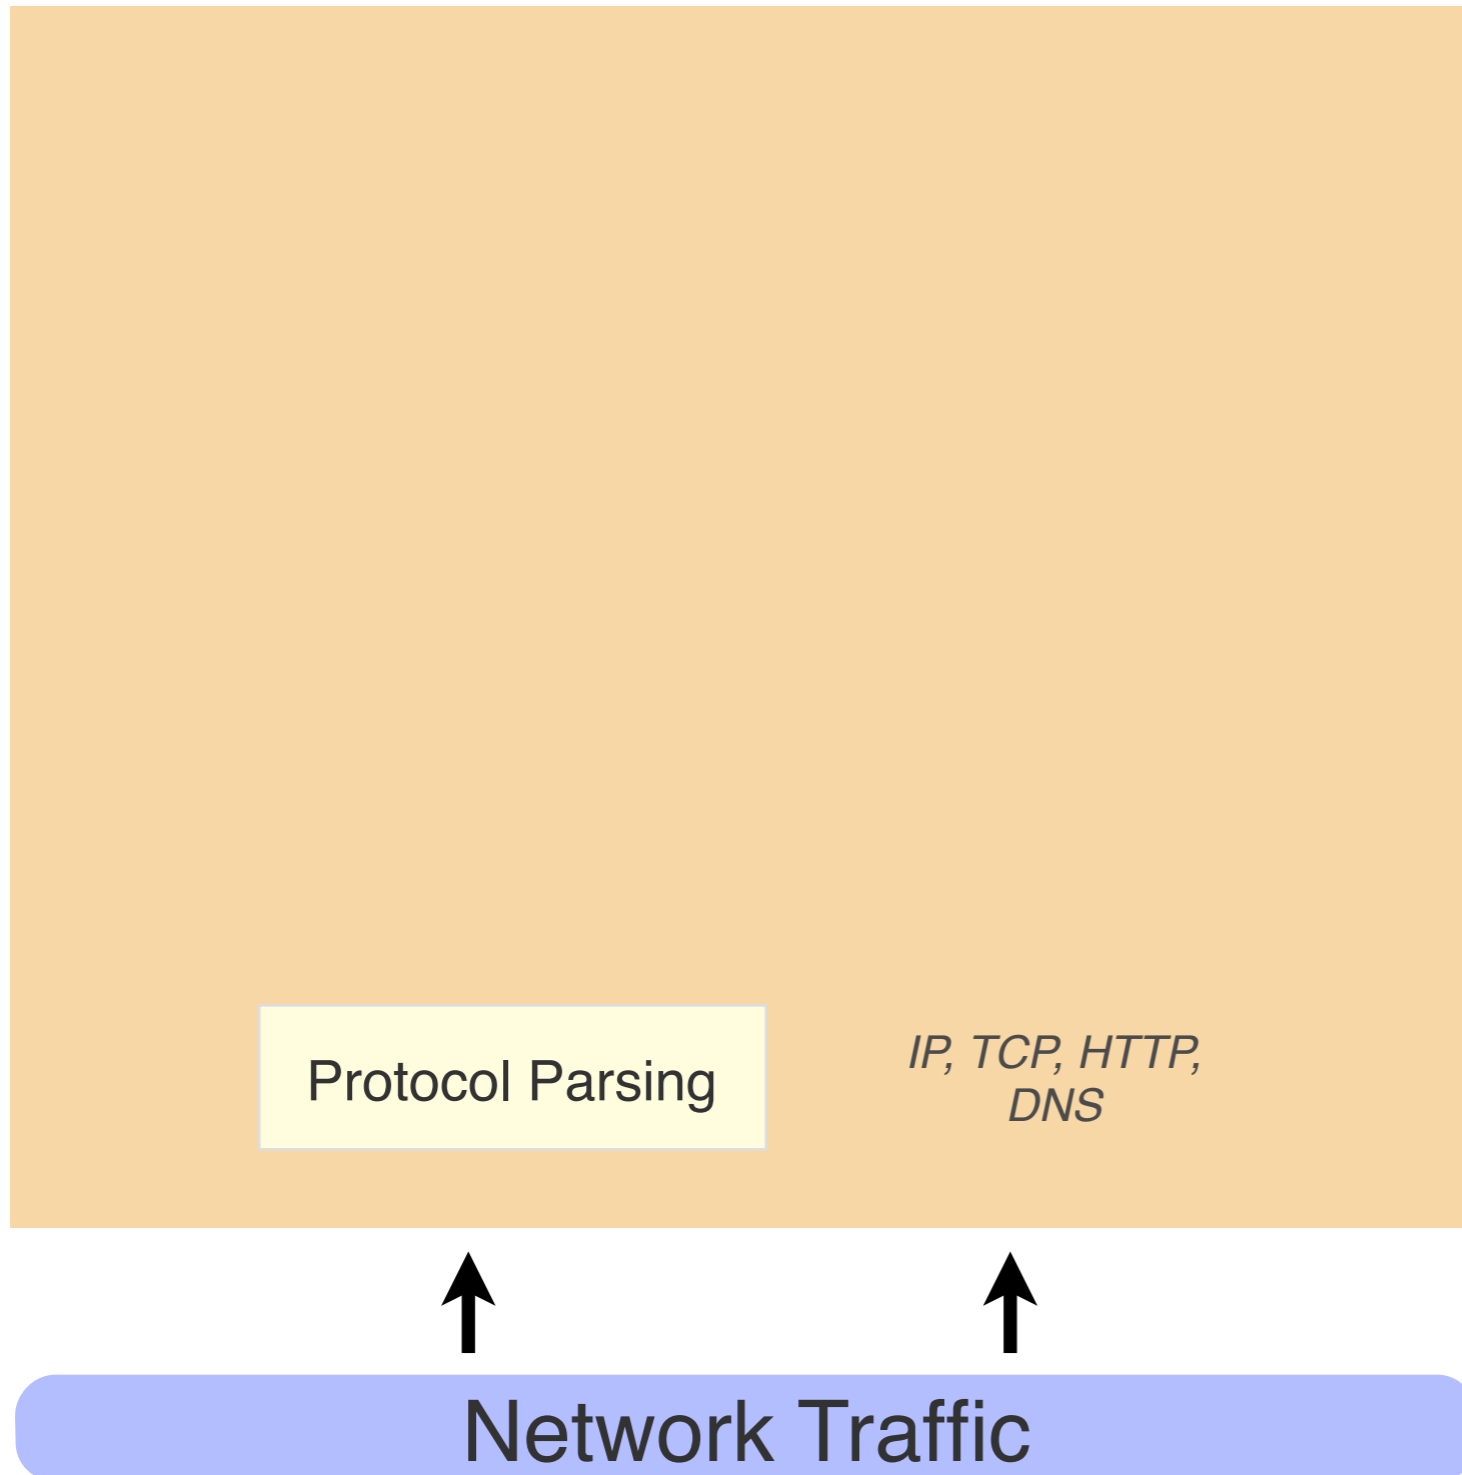
Application



Network Traffic

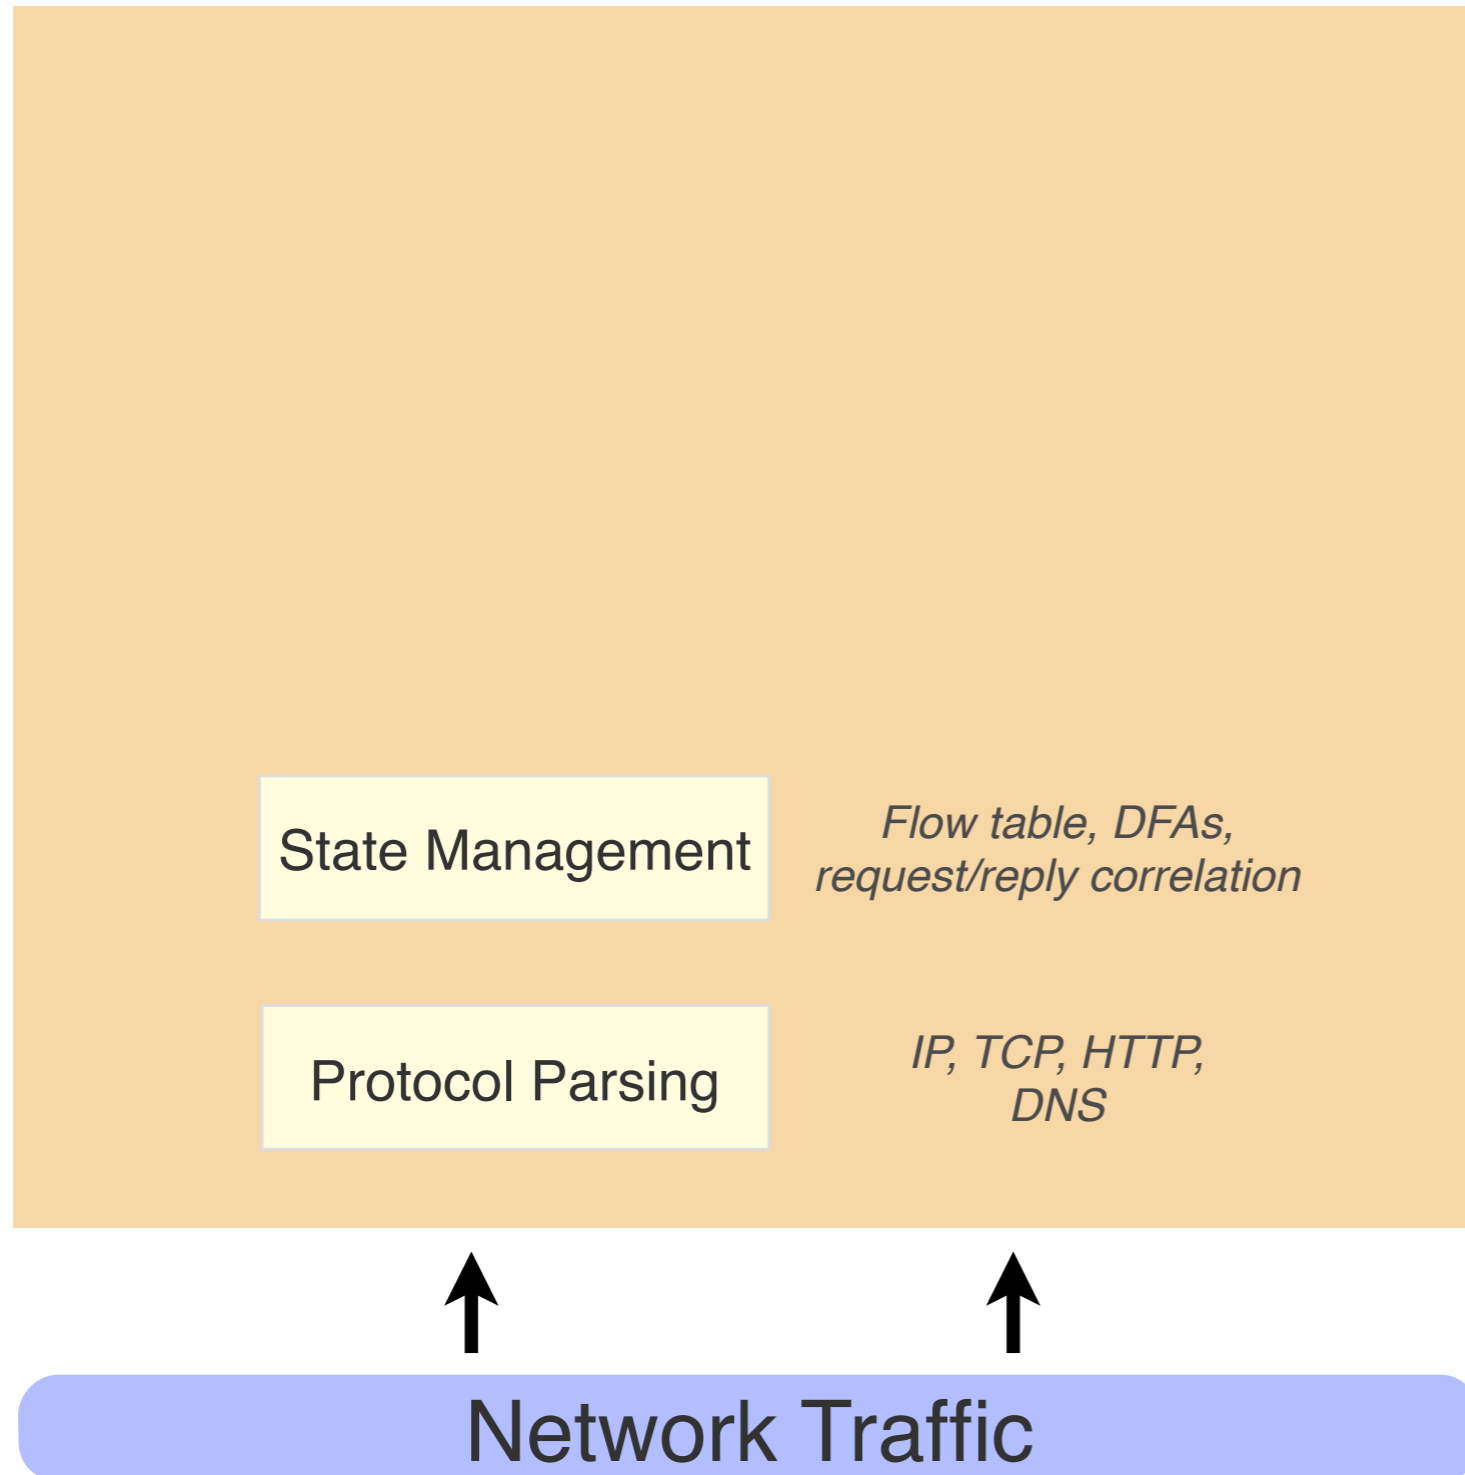
DPI Architecture

Application



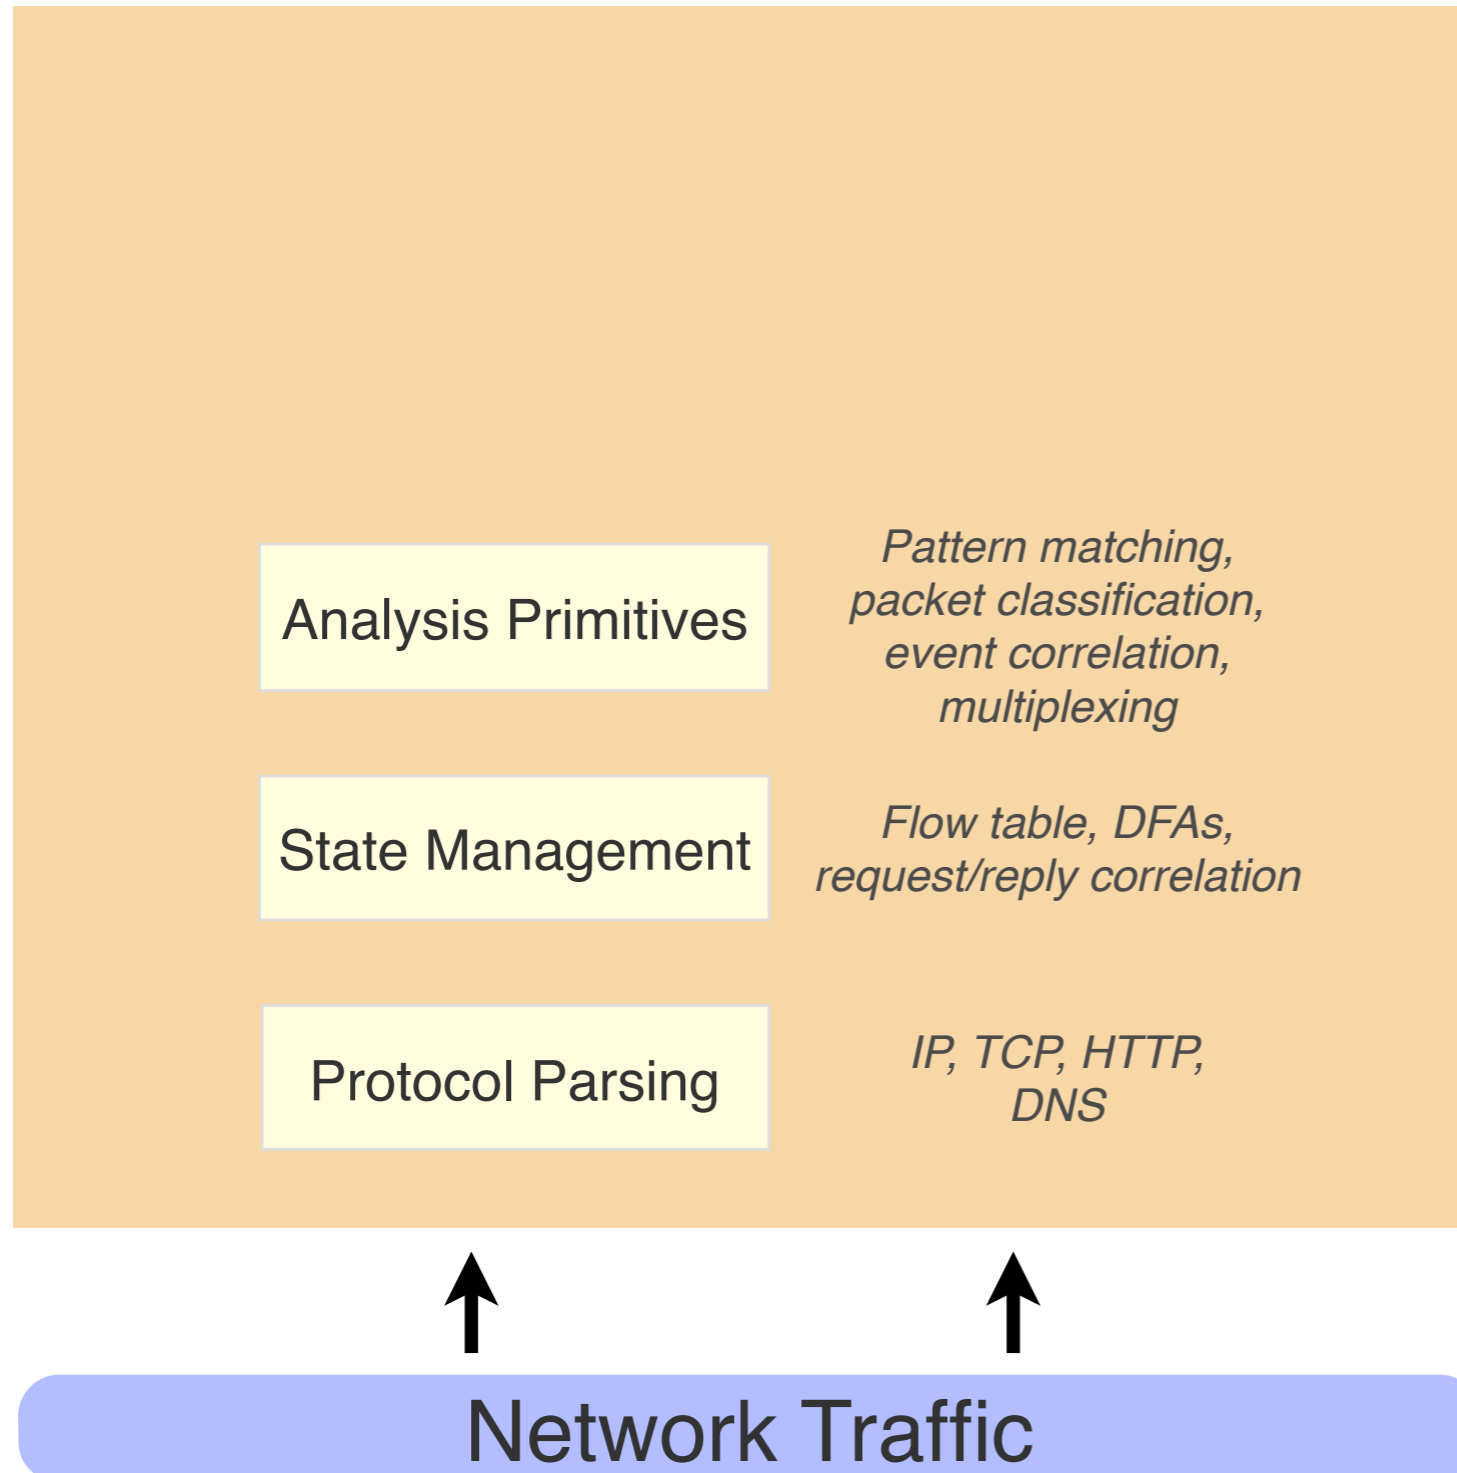
DPI Architecture

Application



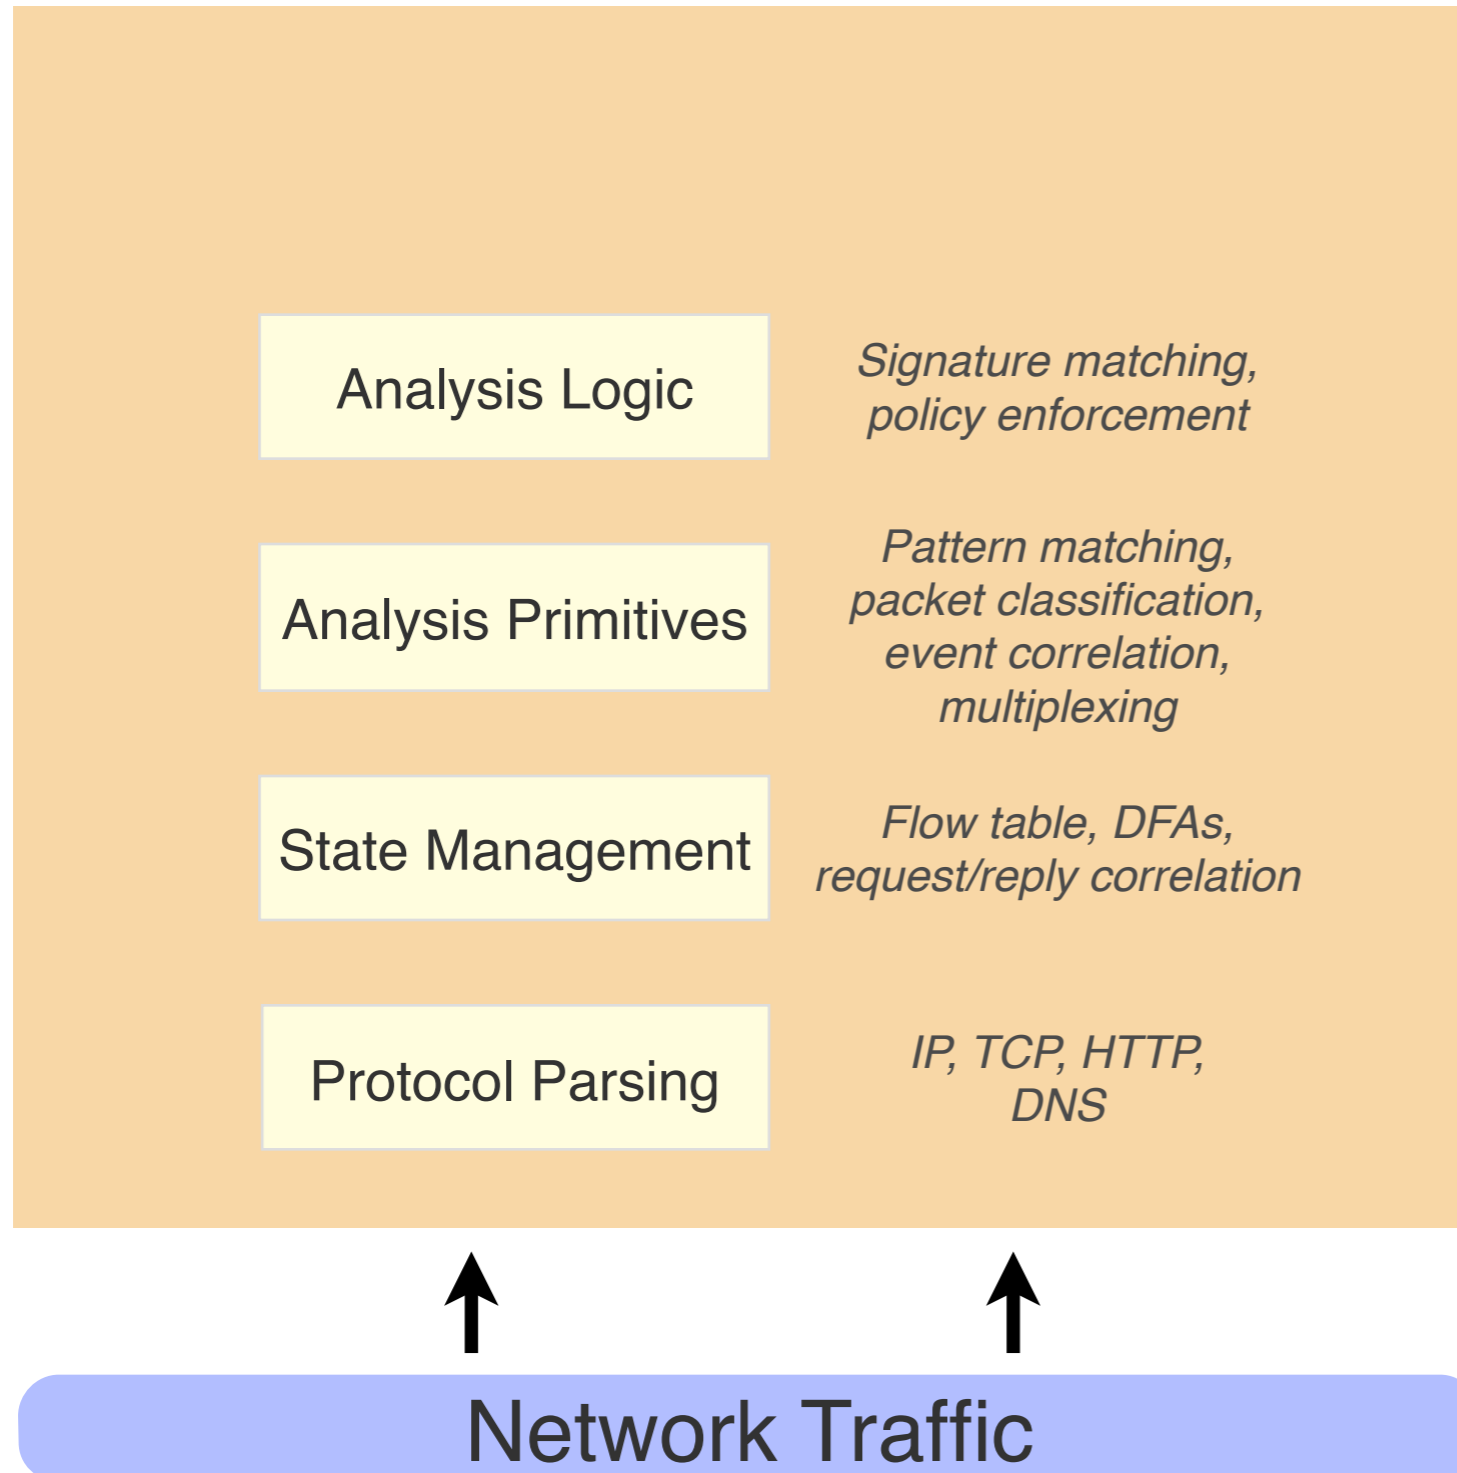
DPI Architecture

Application



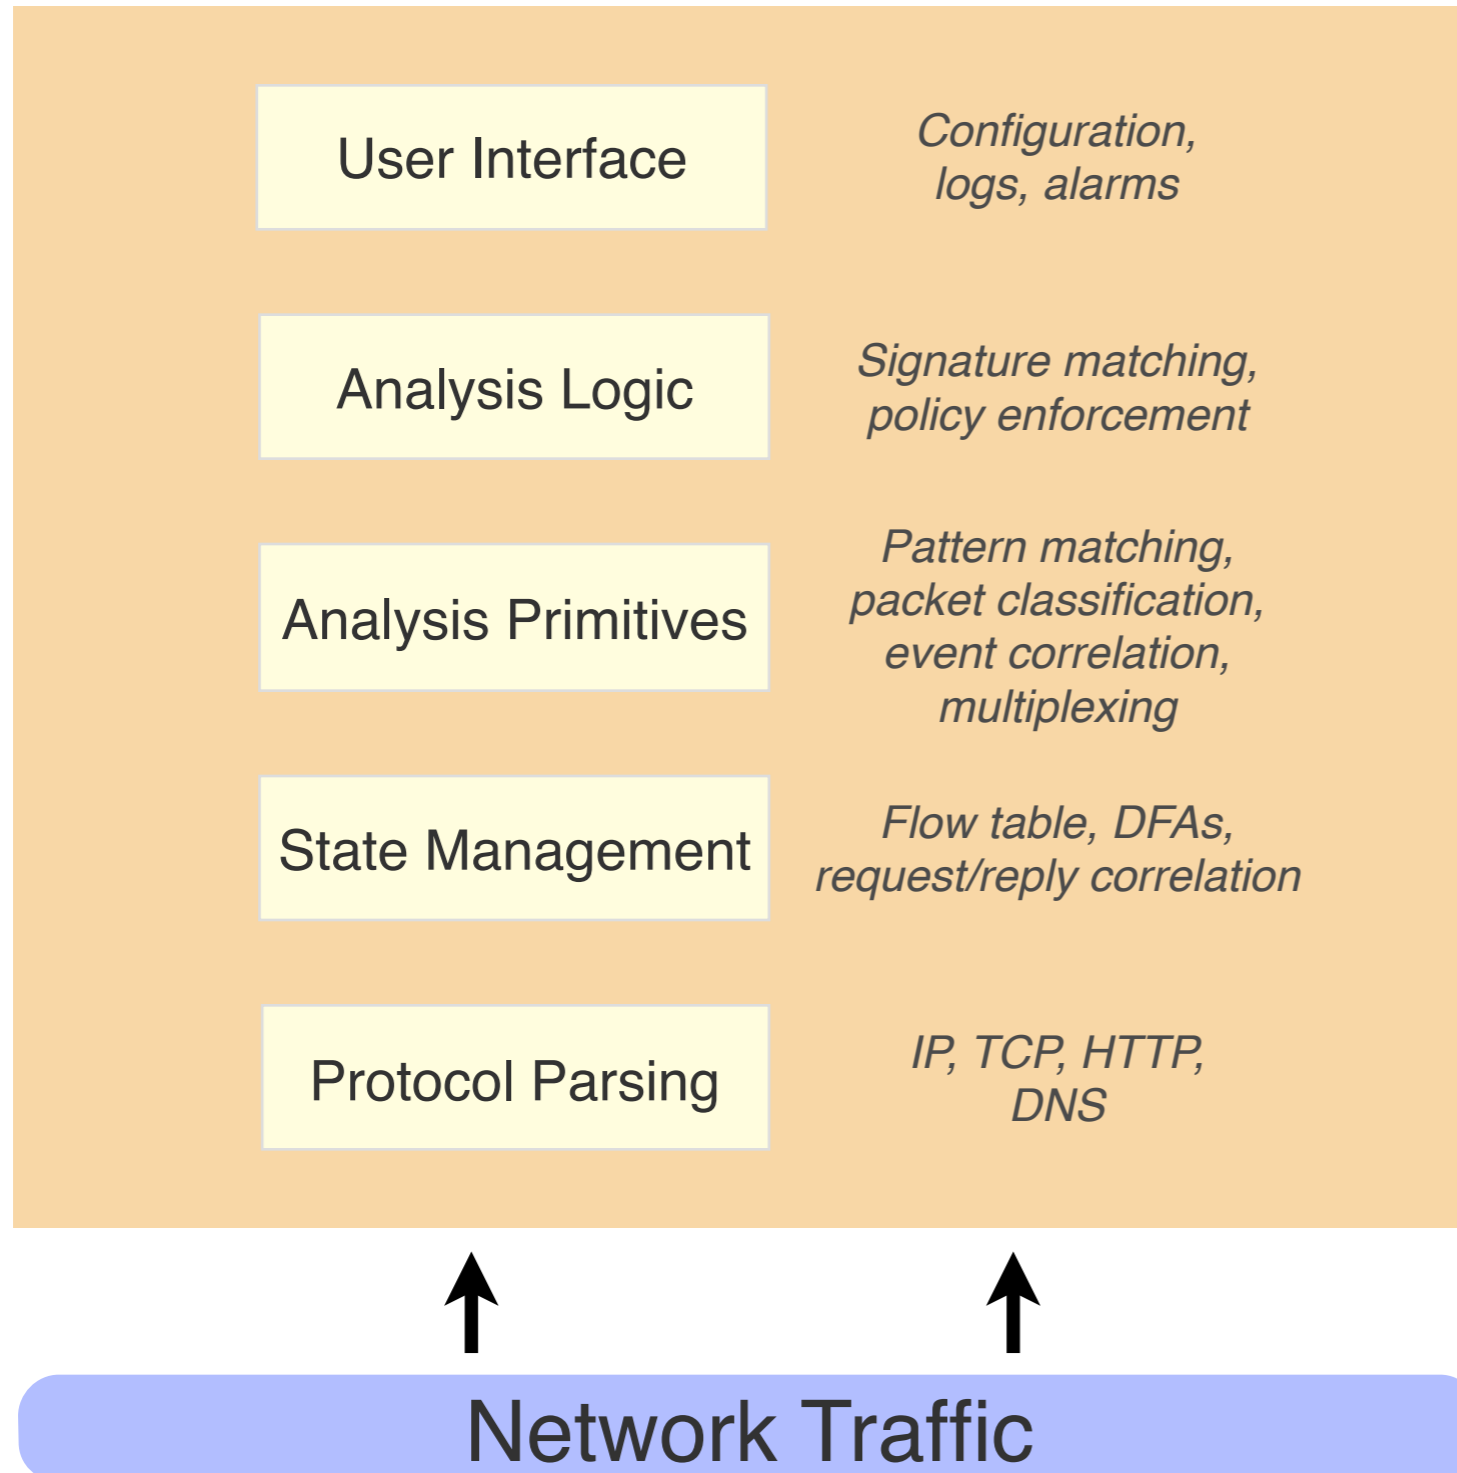
DPI Architecture

Application



DPI Architecture

Application

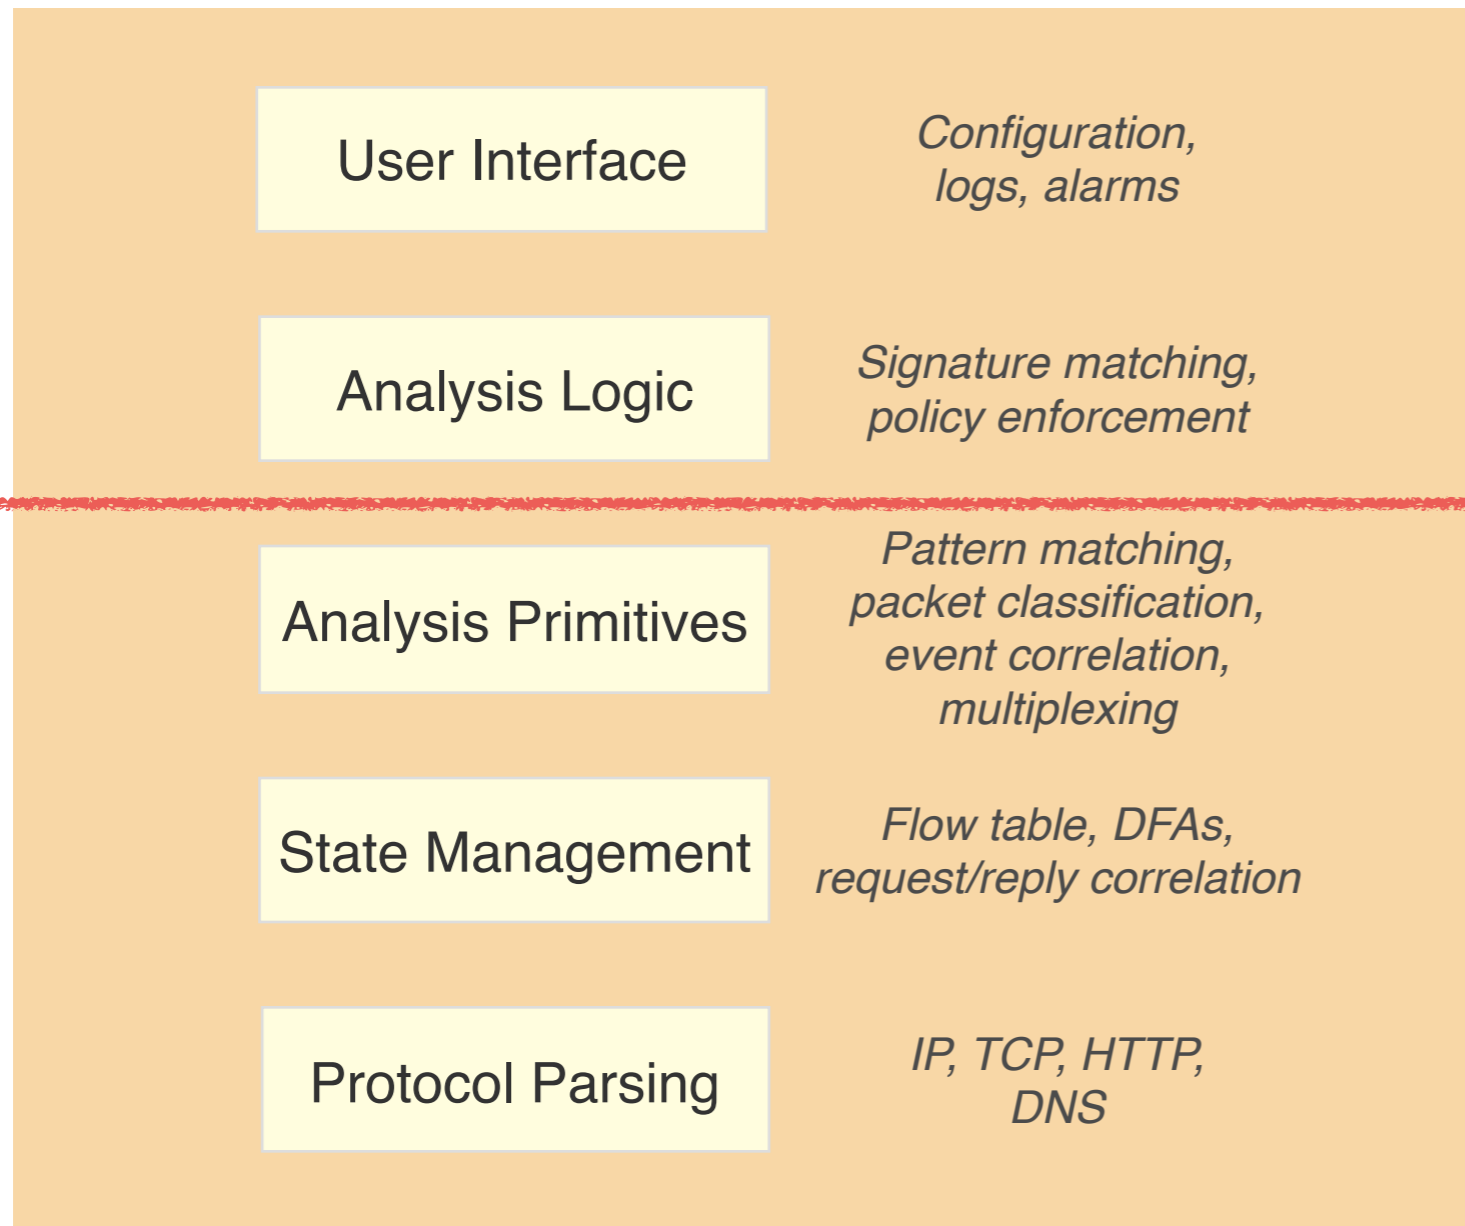


DPI Architecture

Application

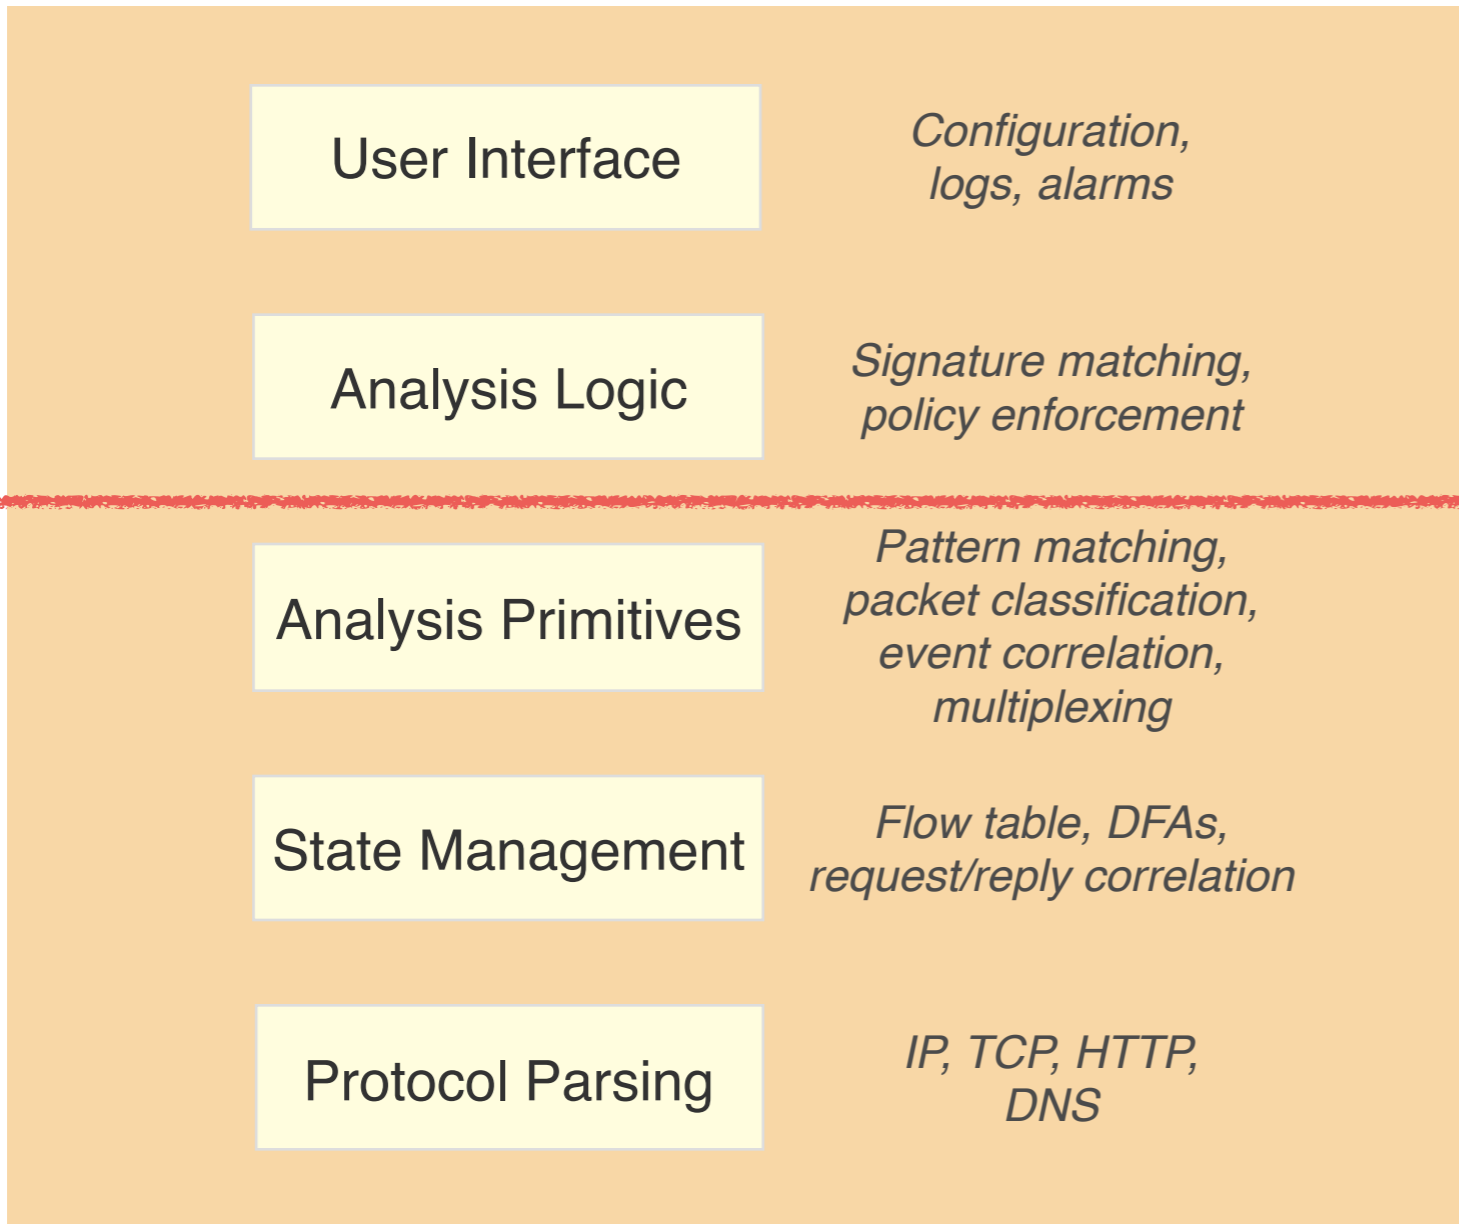
Common primitives & idioms — but hardly any reuse ...

... and that even though this stuff is hard.



DPI Architecture

Application



Common primitives & idioms — but hardly any reuse ...

... and that even though this stuff is hard.

Why?


Different low-level structure & data flows.



Network Traffic

DPI Architecture

Application



Common primitives & idioms — but hardly any reuse ...

... and that even though this stuff is hard.

Why?

Different low-level structure & data flows.

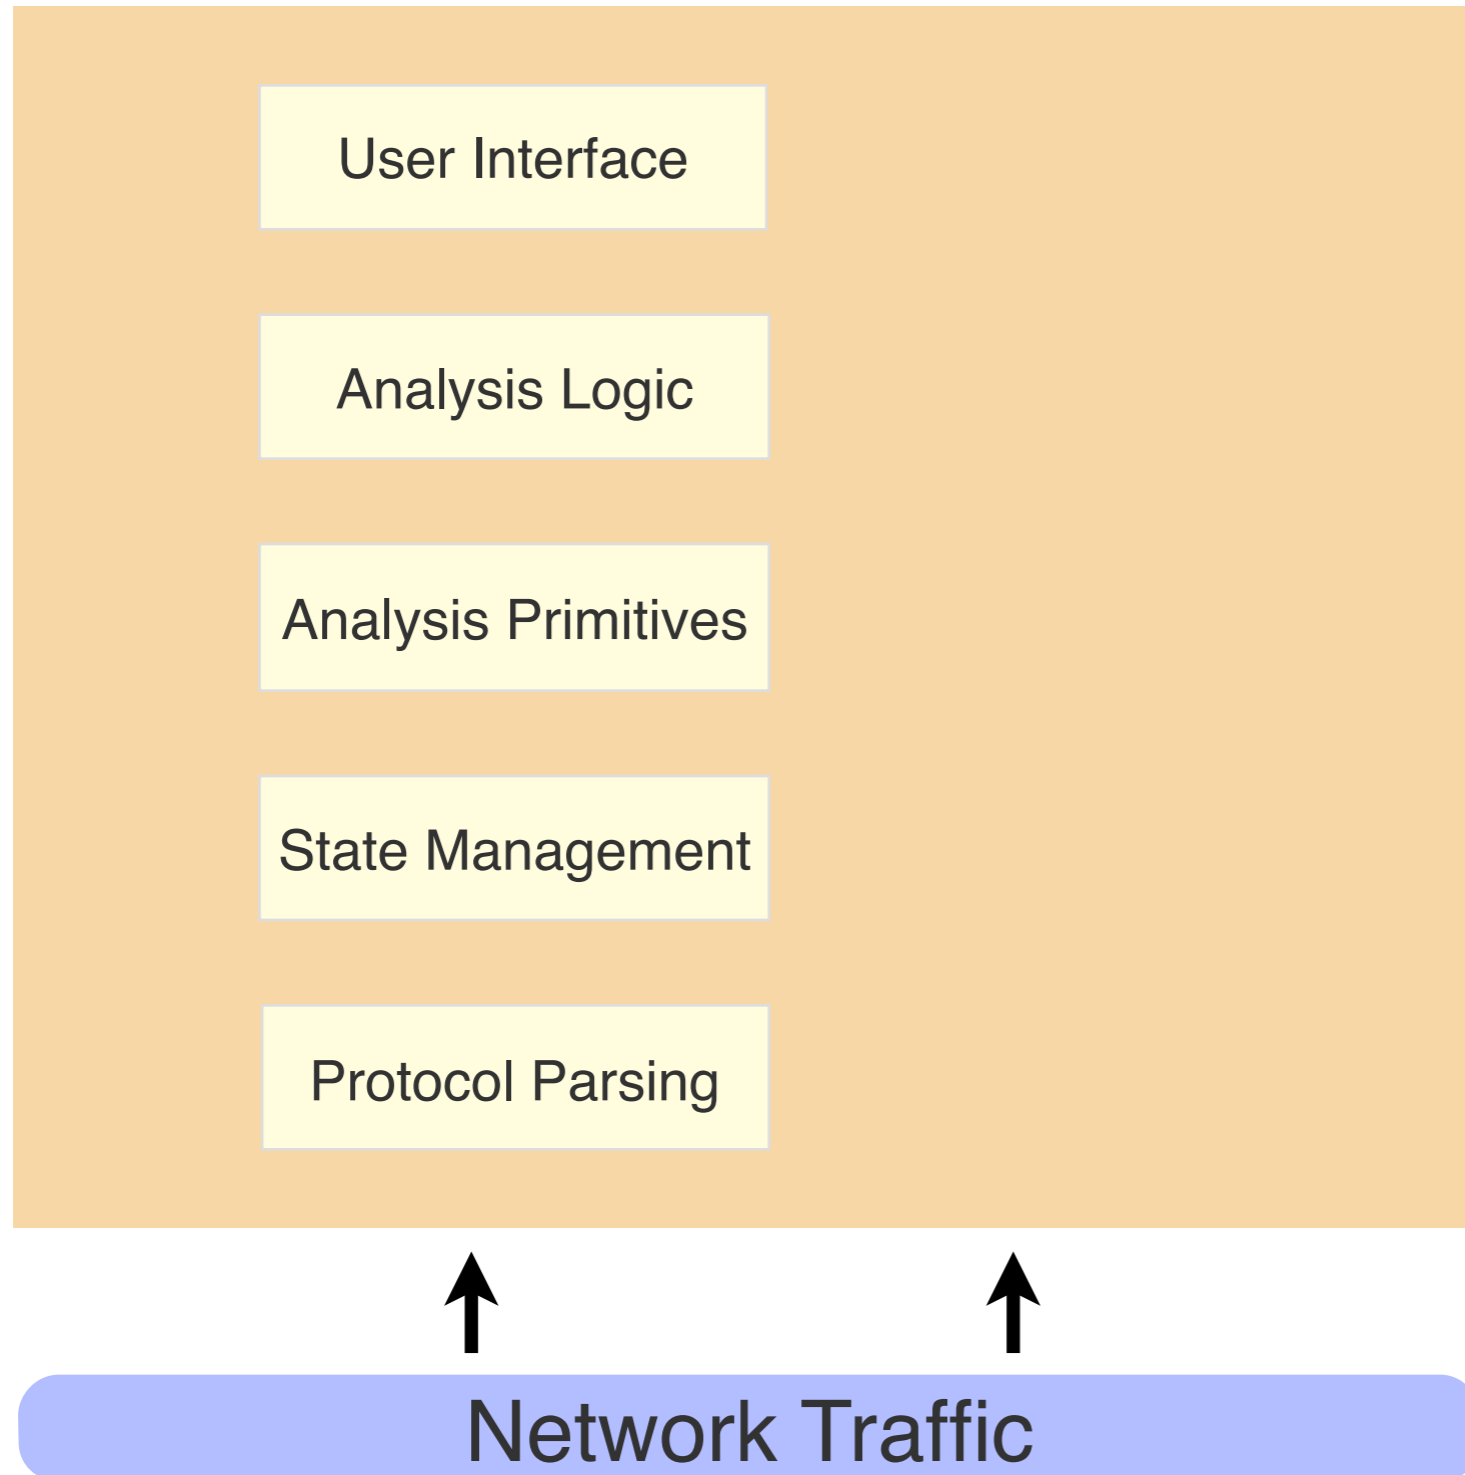
No “common language”.



Network Traffic

A High-Level Intermediary Language for Traffic Inspection

Application



A High-Level Intermediary Language for Traffic Inspection

Application

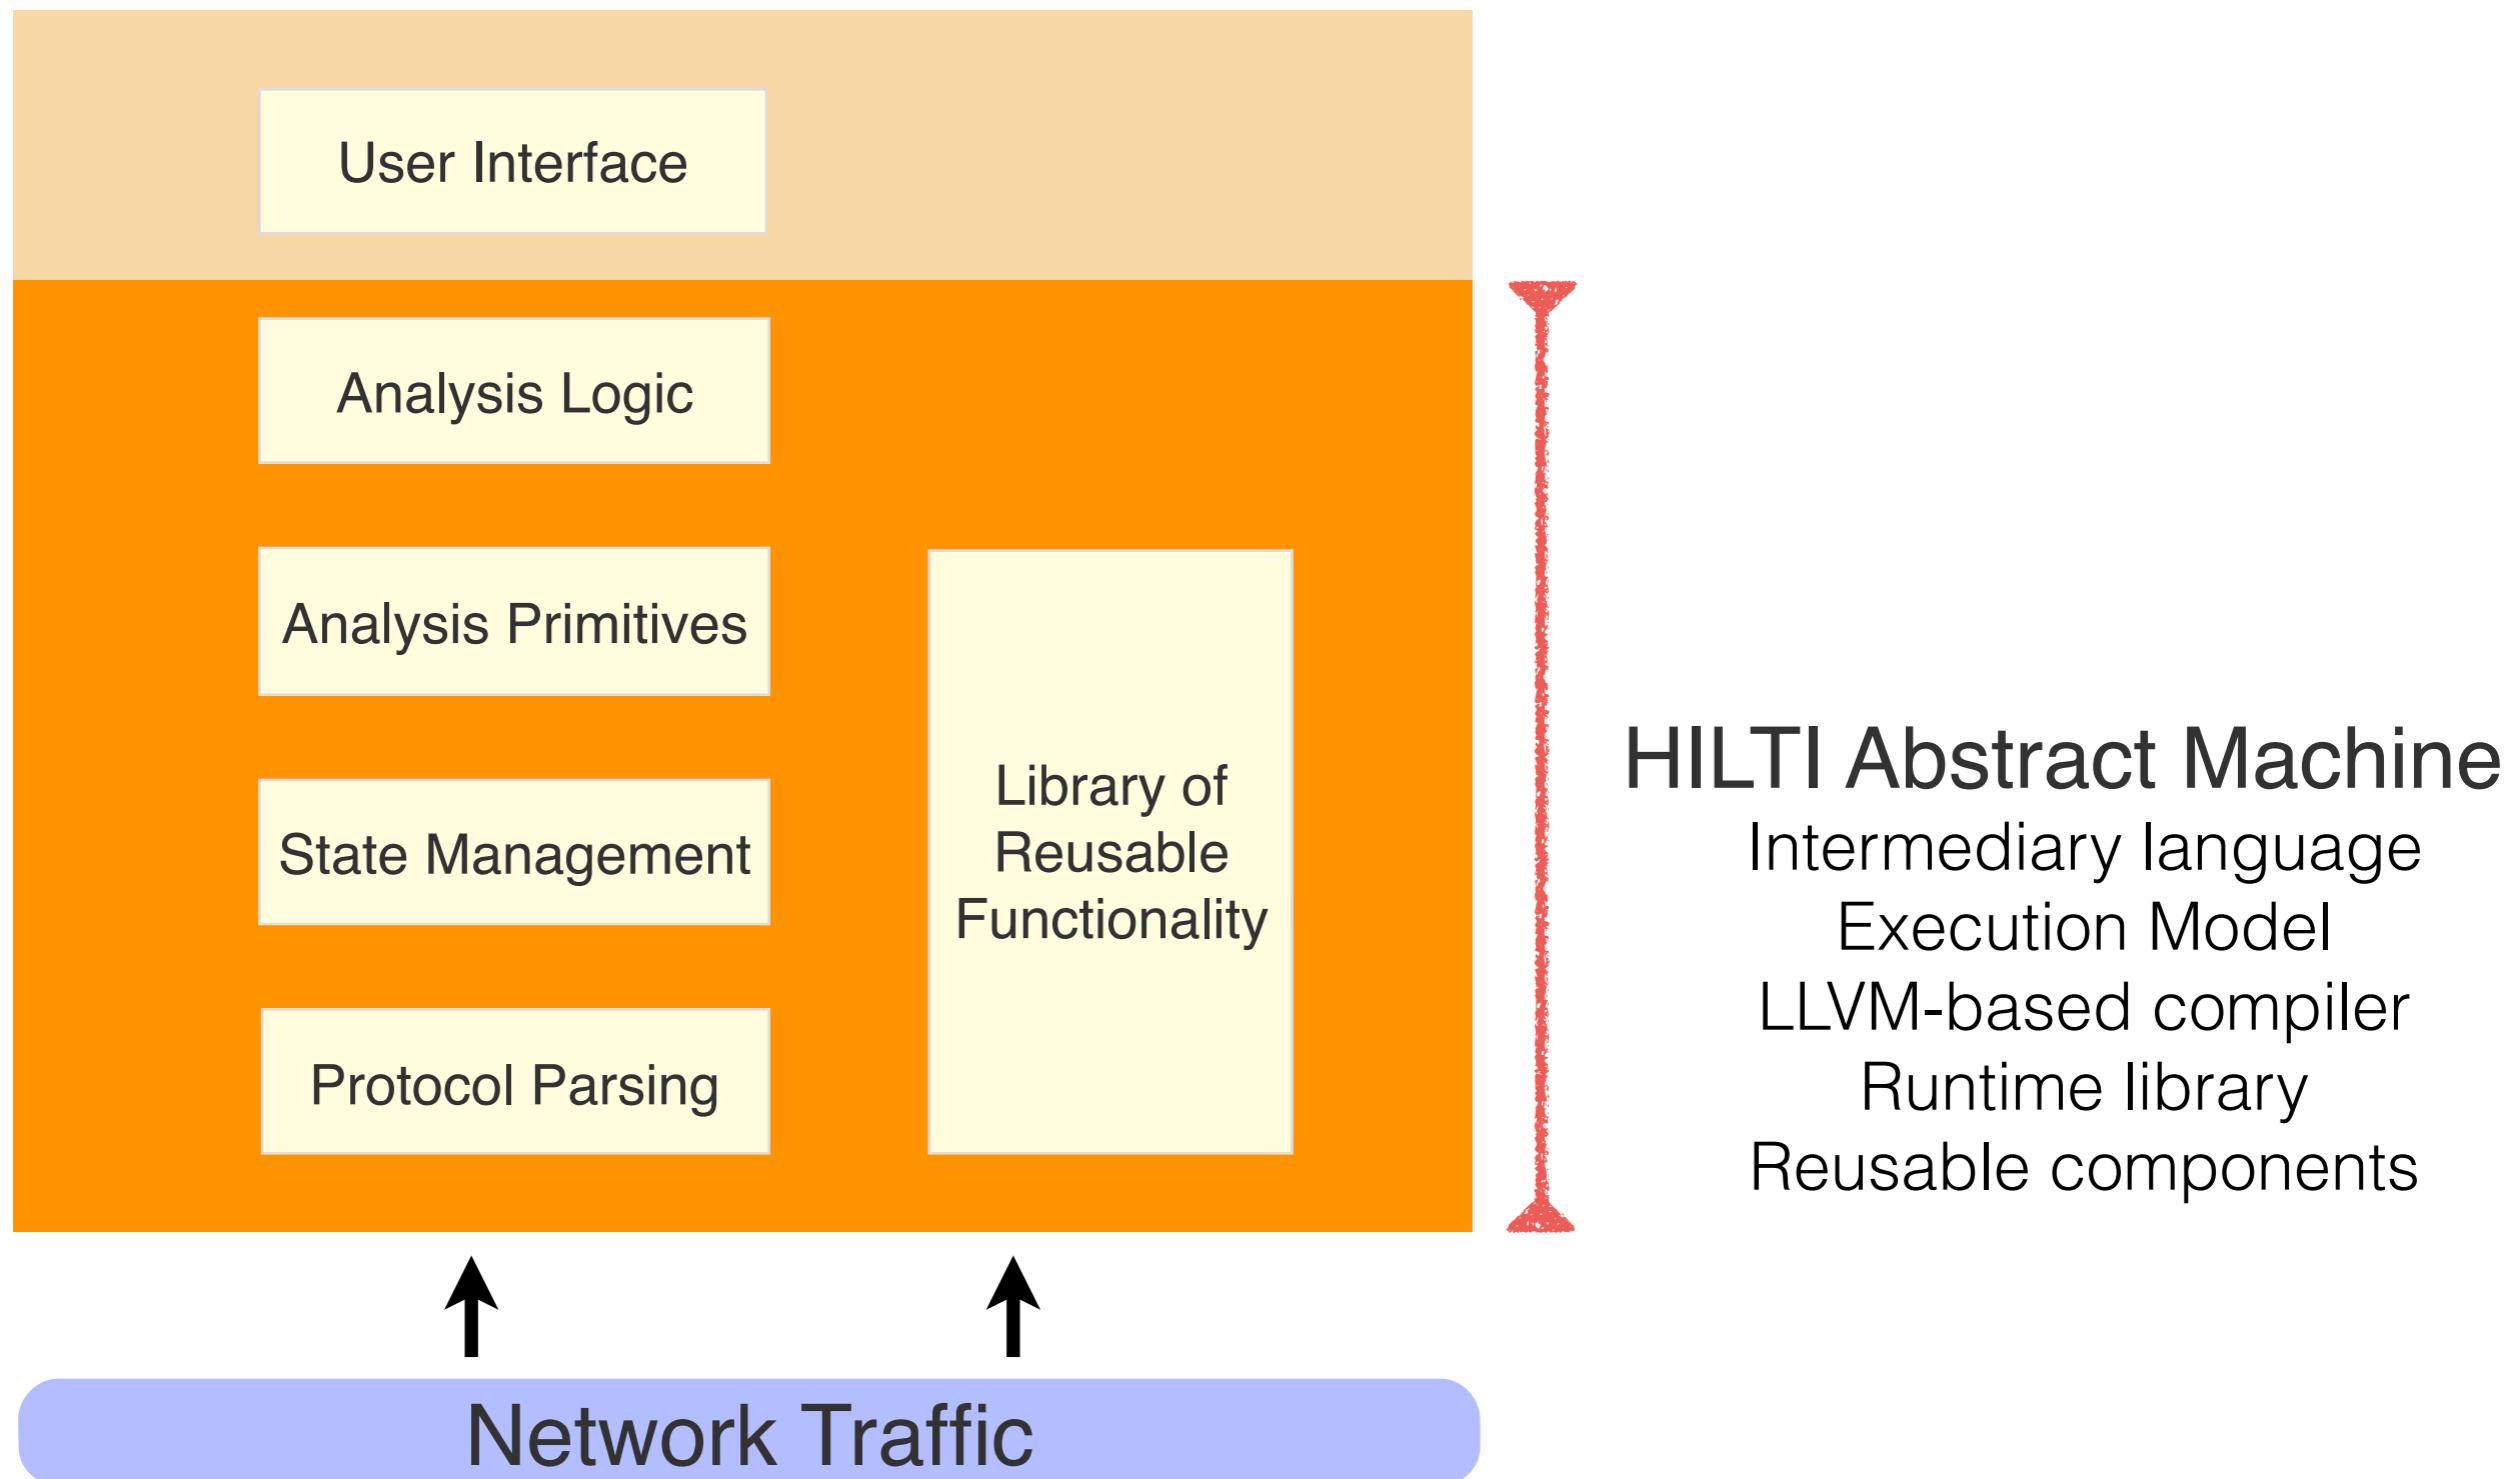


HILTI Abstract Machine

- Intermediary language
- Execution Model
- LLVM-based compiler
- Runtime library
- Reusable components

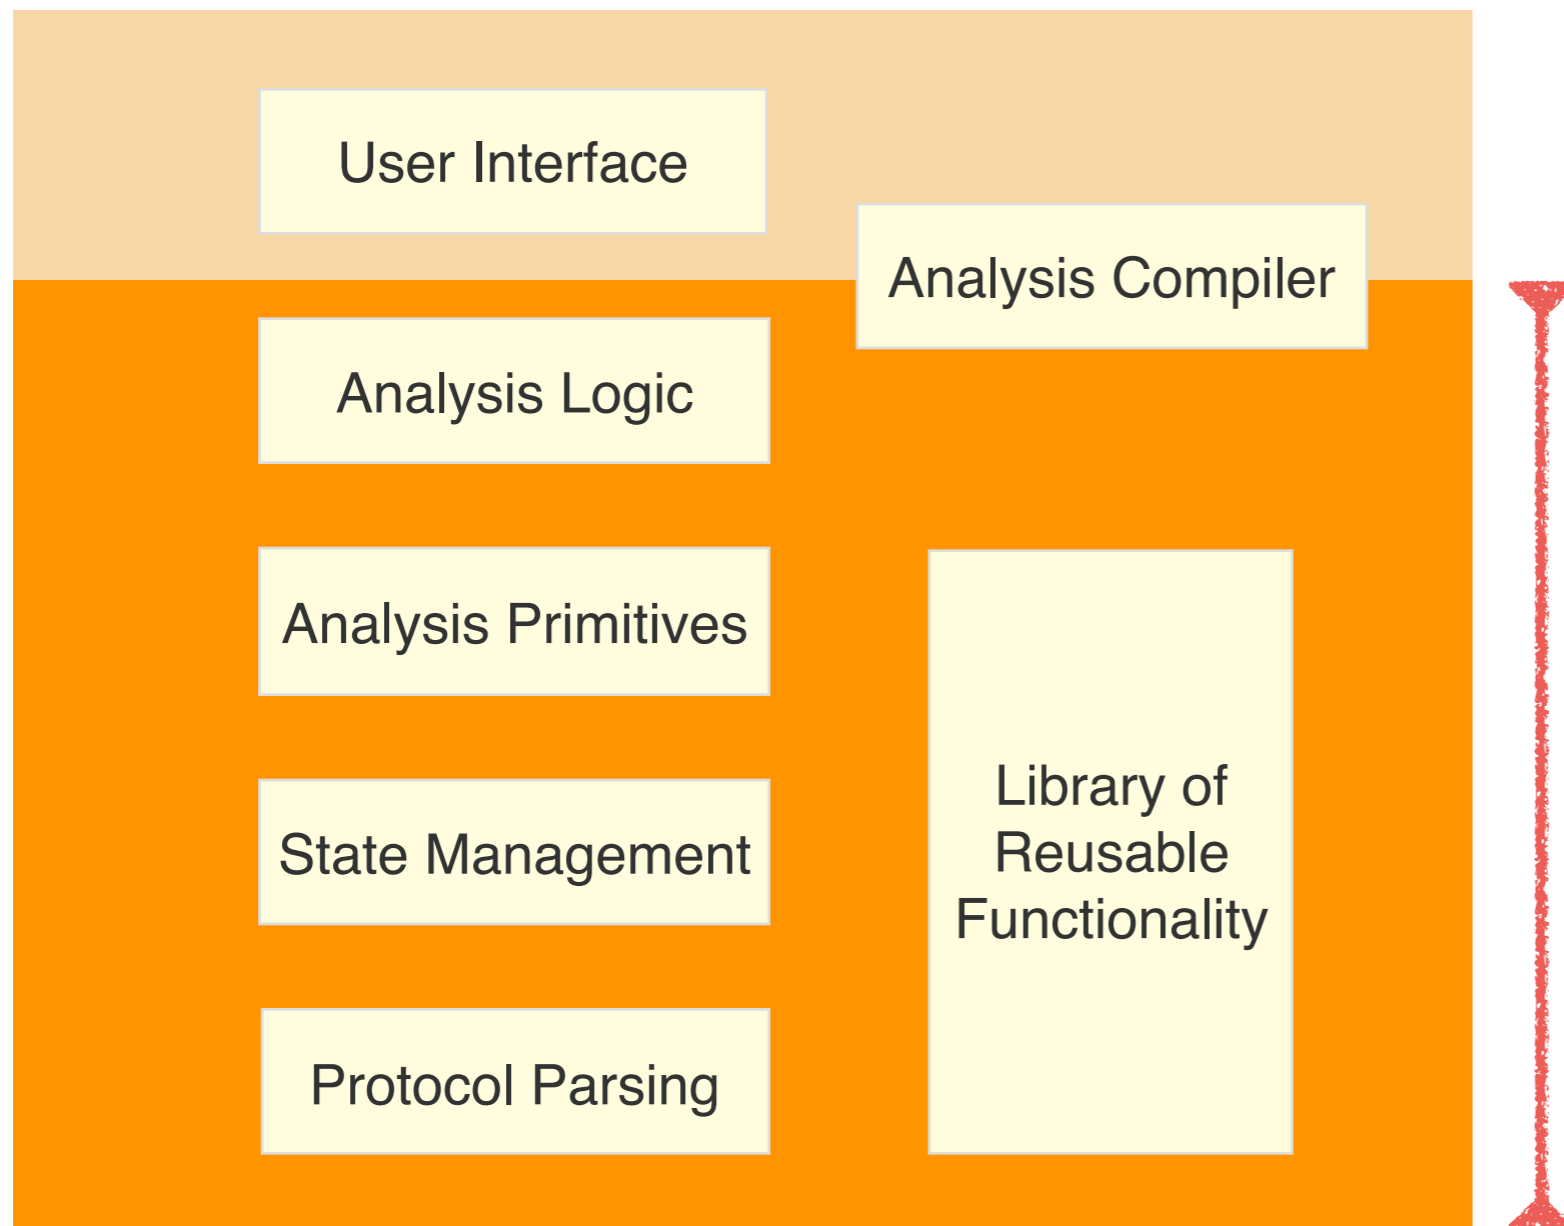
A High-Level Intermediary Language for Traffic Inspection

Application



A High-Level Intermediary Language for Traffic Inspection

Application



Host Application

Firewall rules,
IDS signatures,
forwarding rules, ...

HILTI Abstract Machine

Intermediary language
Execution Model
LLVM-based compiler
Runtime library
Reusable components



Network Traffic

Example: BPF Filters

```
host 192.168.1.1 or src net 10.0.5.0/24
```

Example: BPF Filters

```
host 192.168.1.1 or src net 10.0.5.0/24
```

```
type IP::Header = overlay {  
  version: int<8> at 0 unpack UInt8InBigEndian (4, 7),  
  hdr_len: int<8> at 0 unpack UInt8InBigEndian (0, 3),  
  [...]   
  src:      addr      at 12 unpack IPv4InNetworkOrder,  
  dst:      addr      at 16 unpack IPv4InNetworkOrder  
}
```

Example: BPF Filters

```
host 192.168.1.1 or src net 10.0.5.0/24
```

```
type IP::Header = overlay {  
  version: int<8> at 0 unpack UInt8InBigEndian (4, 7),  
  hdr_len: int<8> at 0 unpack UInt8InBigEndian (0, 3),  
  [...]  
  src:      addr      at 12 unpack IPv4InNetworkOrder,  
  dst:      addr      at 16 unpack IPv4InNetworkOrder  
}
```

```
bool filter(ref<bytes> packet) {  
  local addr a1, a2  
  local bool b1, b2, b3  
  
  a1 = overlay.get IP::Header src packet  
  b1 = equal a1 192.168.1.1  
  a1 = overlay.get IP::Header dst packet  
  b2 = equal a2 192.168.1.1  
  b1 = or b1 b2  
  b2 = equal 10.0.5.0/24 a1  
  b3 = or b1 b2  
  return b3  
}
```

HILTI Machine Model

Focus Areas

Rich Domain-specific Data Types

Flexible Control Flow

Concurrent Analysis

Robust & Secure Execution

Comprehensive Host Interface

Real-time Performance

Debugging & Profiling Support

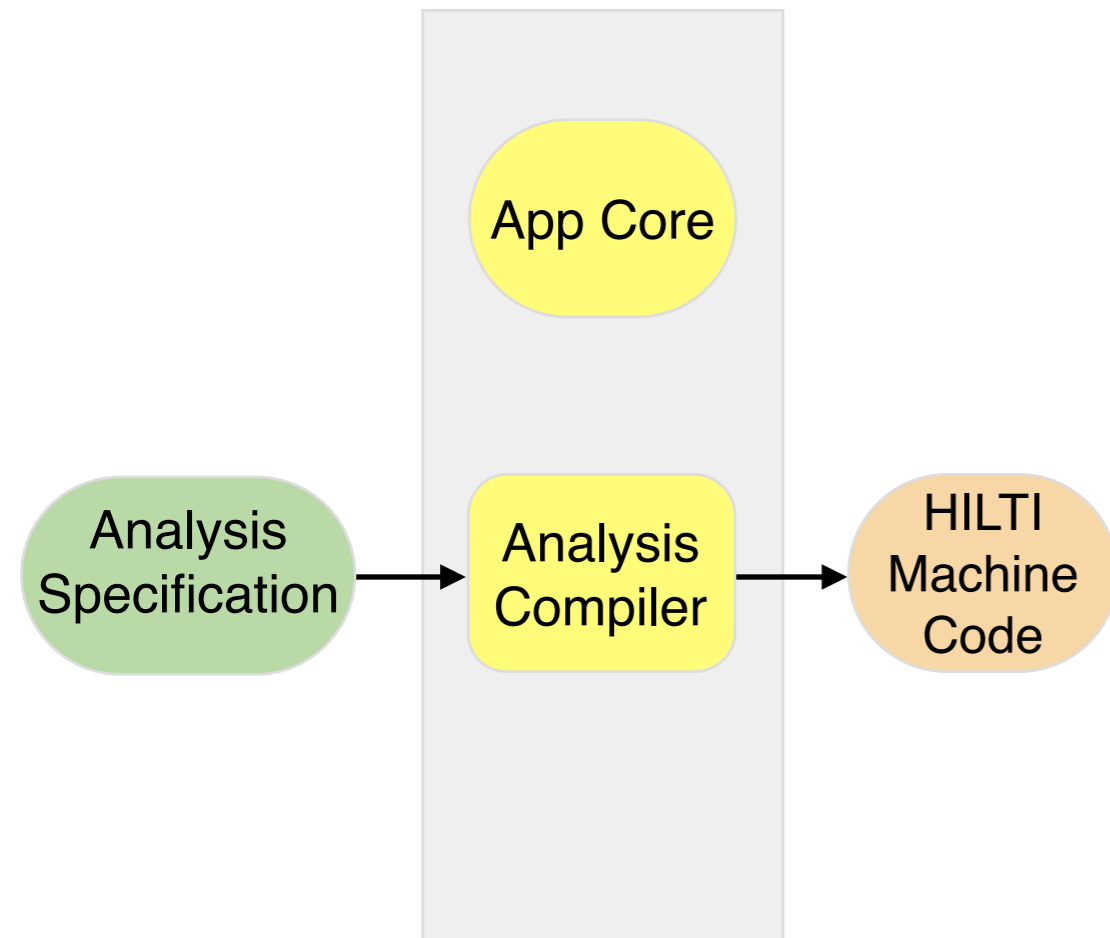
High-level Optimization

Instruction Set

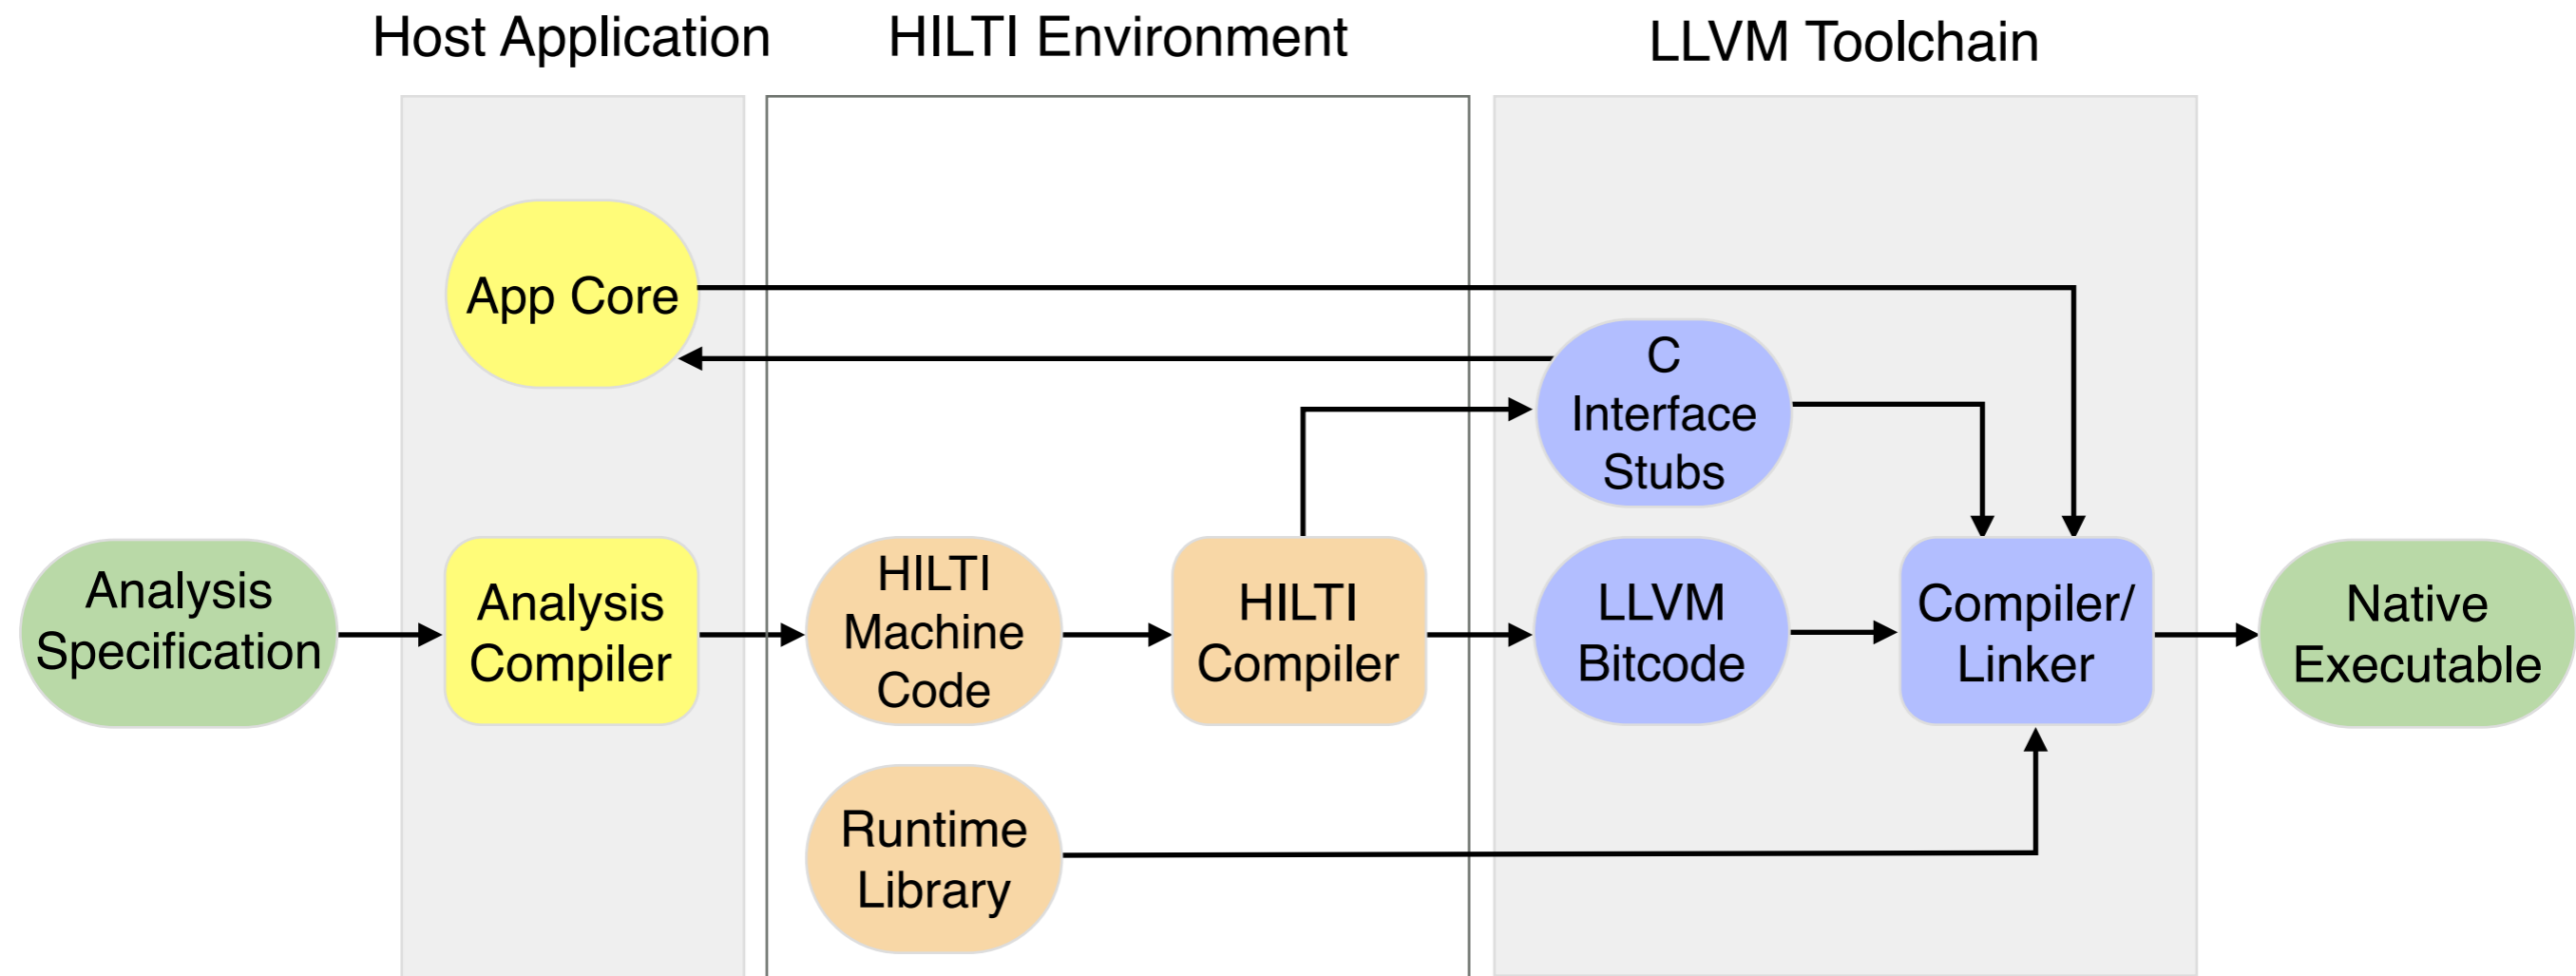
Bitsets	Packet input
Booleans	Packet classification
CIDR masks	Packet dissection
Callbacks	Ports
Closures	Profiling
Channels	Raw data
Debug support	References
Doubles	Regular expressions
Enumerations	Strings
Exceptions	Structs
File i/o	Unions
Flow control	Time intervals
Hashmaps	Times
Hashsets	Timers
IP addresses	Threads
Integers	Tuples
Lists	Vectors/arrays

Implementation: The HILTI Toolchain

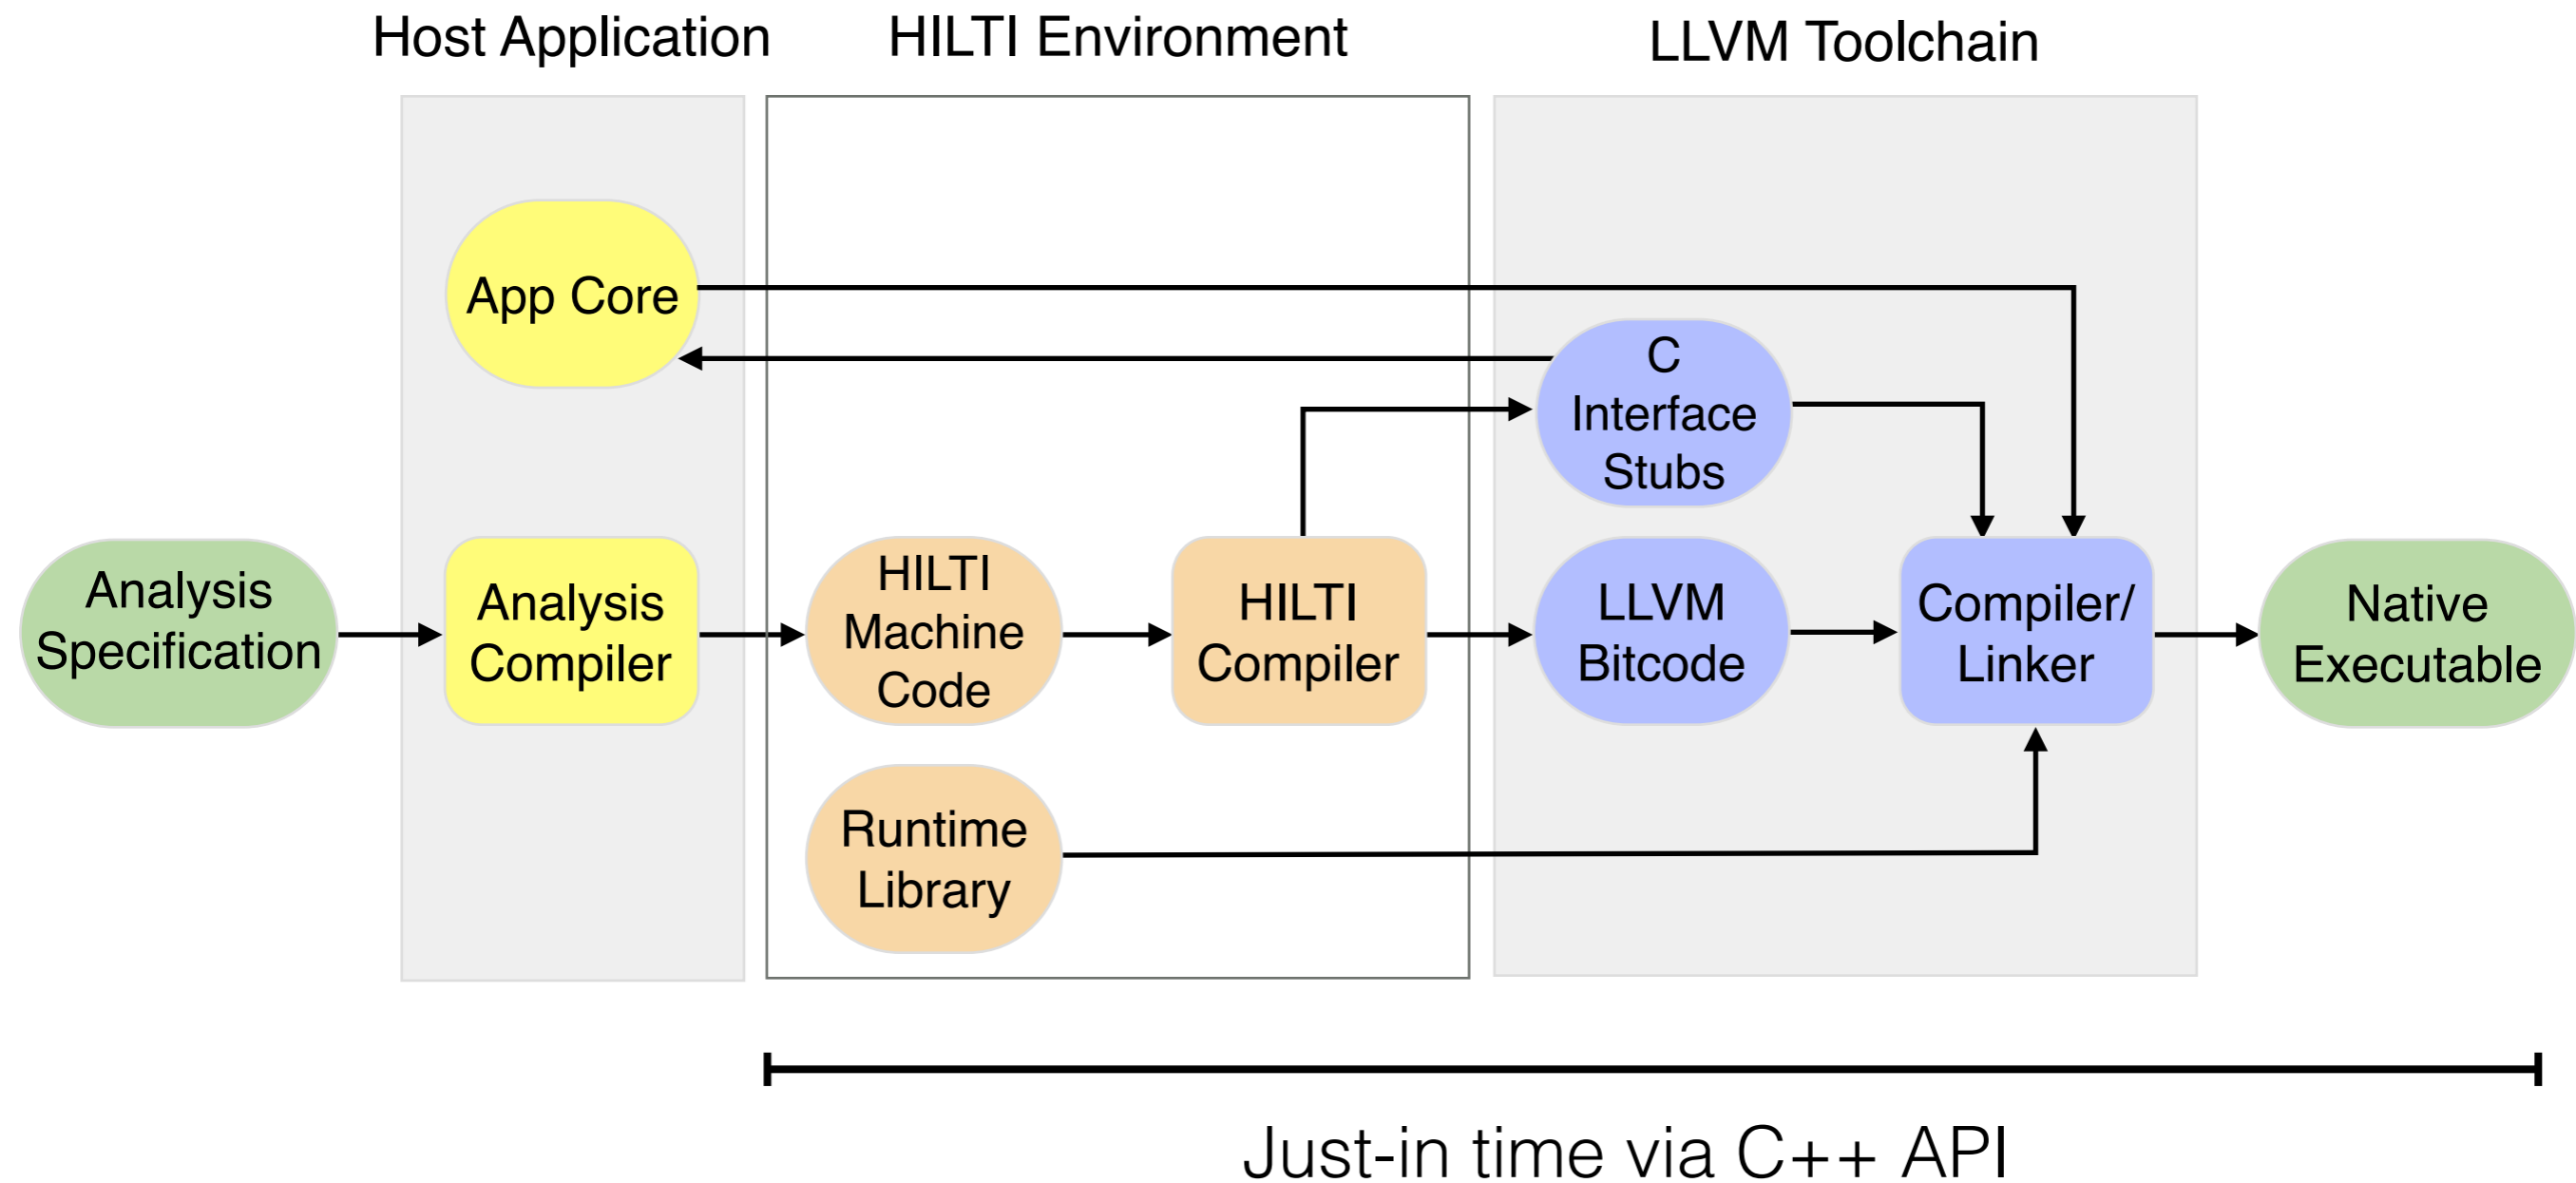
Host Application



Implementation: The HILTI Toolchain



Implementation: The HILTI Toolchain



Can HILTI support complex applications?

Application Case Studies

BPF Filter

Stateful Firewall

Protocol Parsing

Bro Script Execution

Application Case Studies

BPF Filter

Stateful Firewall

Protocol Parsing

Bro Script Execution

BinPAC - A Yacc for Network Protocols

Grammar example: Parsing SSH banners.

```
type SSH::Banner = unit {  
  magic      : /SSH-;/  
  version    : /^[^-]*/;  
  dash       : /-;/  
  software:  /^[^\r\n]*/;  
}
```

BinPAC - A Yacc for Network Protocols

Grammar example: Parsing SSH banners.

```
type SSH::Banner = unit {  
  magic      : /SSH-;/  
  version    : /^[^-]*/;  
  dash       : /-;/  
  software:  /^[^\r\n]*/;  
}
```

BinPAC compiles grammar into HILTI parser.

HILTI compiles parser into executable code just-in-time.

BinPAC - A Yacc for Network Protocols

Grammar example: Parsing SSH banners.

```
type SSH::Banner = unit {  
  magic      : /SSH-;/  
  version    : /^[^-]*/;  
  dash       : /-;/  
  software:  /^[^\r\n]*/;  
}
```

BinPAC compiles grammar into HILTI parser.

HILTI compiles parser into executable code just-in-time.

Bro plugin integrates parsers at startup.

Bro Scripts

Script example: A simple scan detector.

```
global attempts: table[addr] of count &default=0;

event connection_rejected(c: connection)
{
    local orig = c$id$orig_h;    # Get originator address.

    local n = ++attempts[orig]; # Increase counter.

    if ( n == SOME_THRESHOLD ) # Check for threshold.
        NOTICE(...);        # Alarm.
}
```

Bro Scripts

Script example: A simple scan detector.

```
global attempts: table[addr] of count &default=0;

event connection_rejected(c: connection)
{
    local orig = c$id$orig_h;    # Get originator address.

    local n = ++attempts[orig]; # Increase counter.

    if ( n == SOME_THRESHOLD )  # Check for threshold.
        NOTICE(...);         # Alarm.
}
```

Bro plugin compiles scripts into HILTI code.
HILTI compiles that into executable code just-in-time.

Evaluation

Use HILTI plugin for Bro to compare parsing & script execution with a native Bro.

Traces: HTTP: 1/25 of Berkeley port 80 traffic.
30GB trace, 52min, 340k messages.

DNS: Full Berkeley port 53 traffic.
1GB trace, 10min, 65M messages.

Evaluation

Use HILTI plugin for Bro to compare parsing & script execution with a native Bro.

Traces: HTTP: 1/25 of Berkeley port 80 traffic.
30GB trace, 52min, 340k messages.

DNS: Full Berkeley port 53 traffic.
1GB trace, 10min, 65M messages.

Correctness

HILTI captures semantics correctly.

Evaluation

Use HILTI plugin for Bro to compare parsing & script execution with a native Bro.

Traces: HTTP: 1/25 of Berkeley port 80 traffic.
30GB trace, 52min, 340k messages.

DNS: Full Berkeley port 53 traffic.
1GB trace, 10min, 65M messages.


Correctness

HILTI captures semantics correctly.


Performance

Let's see.

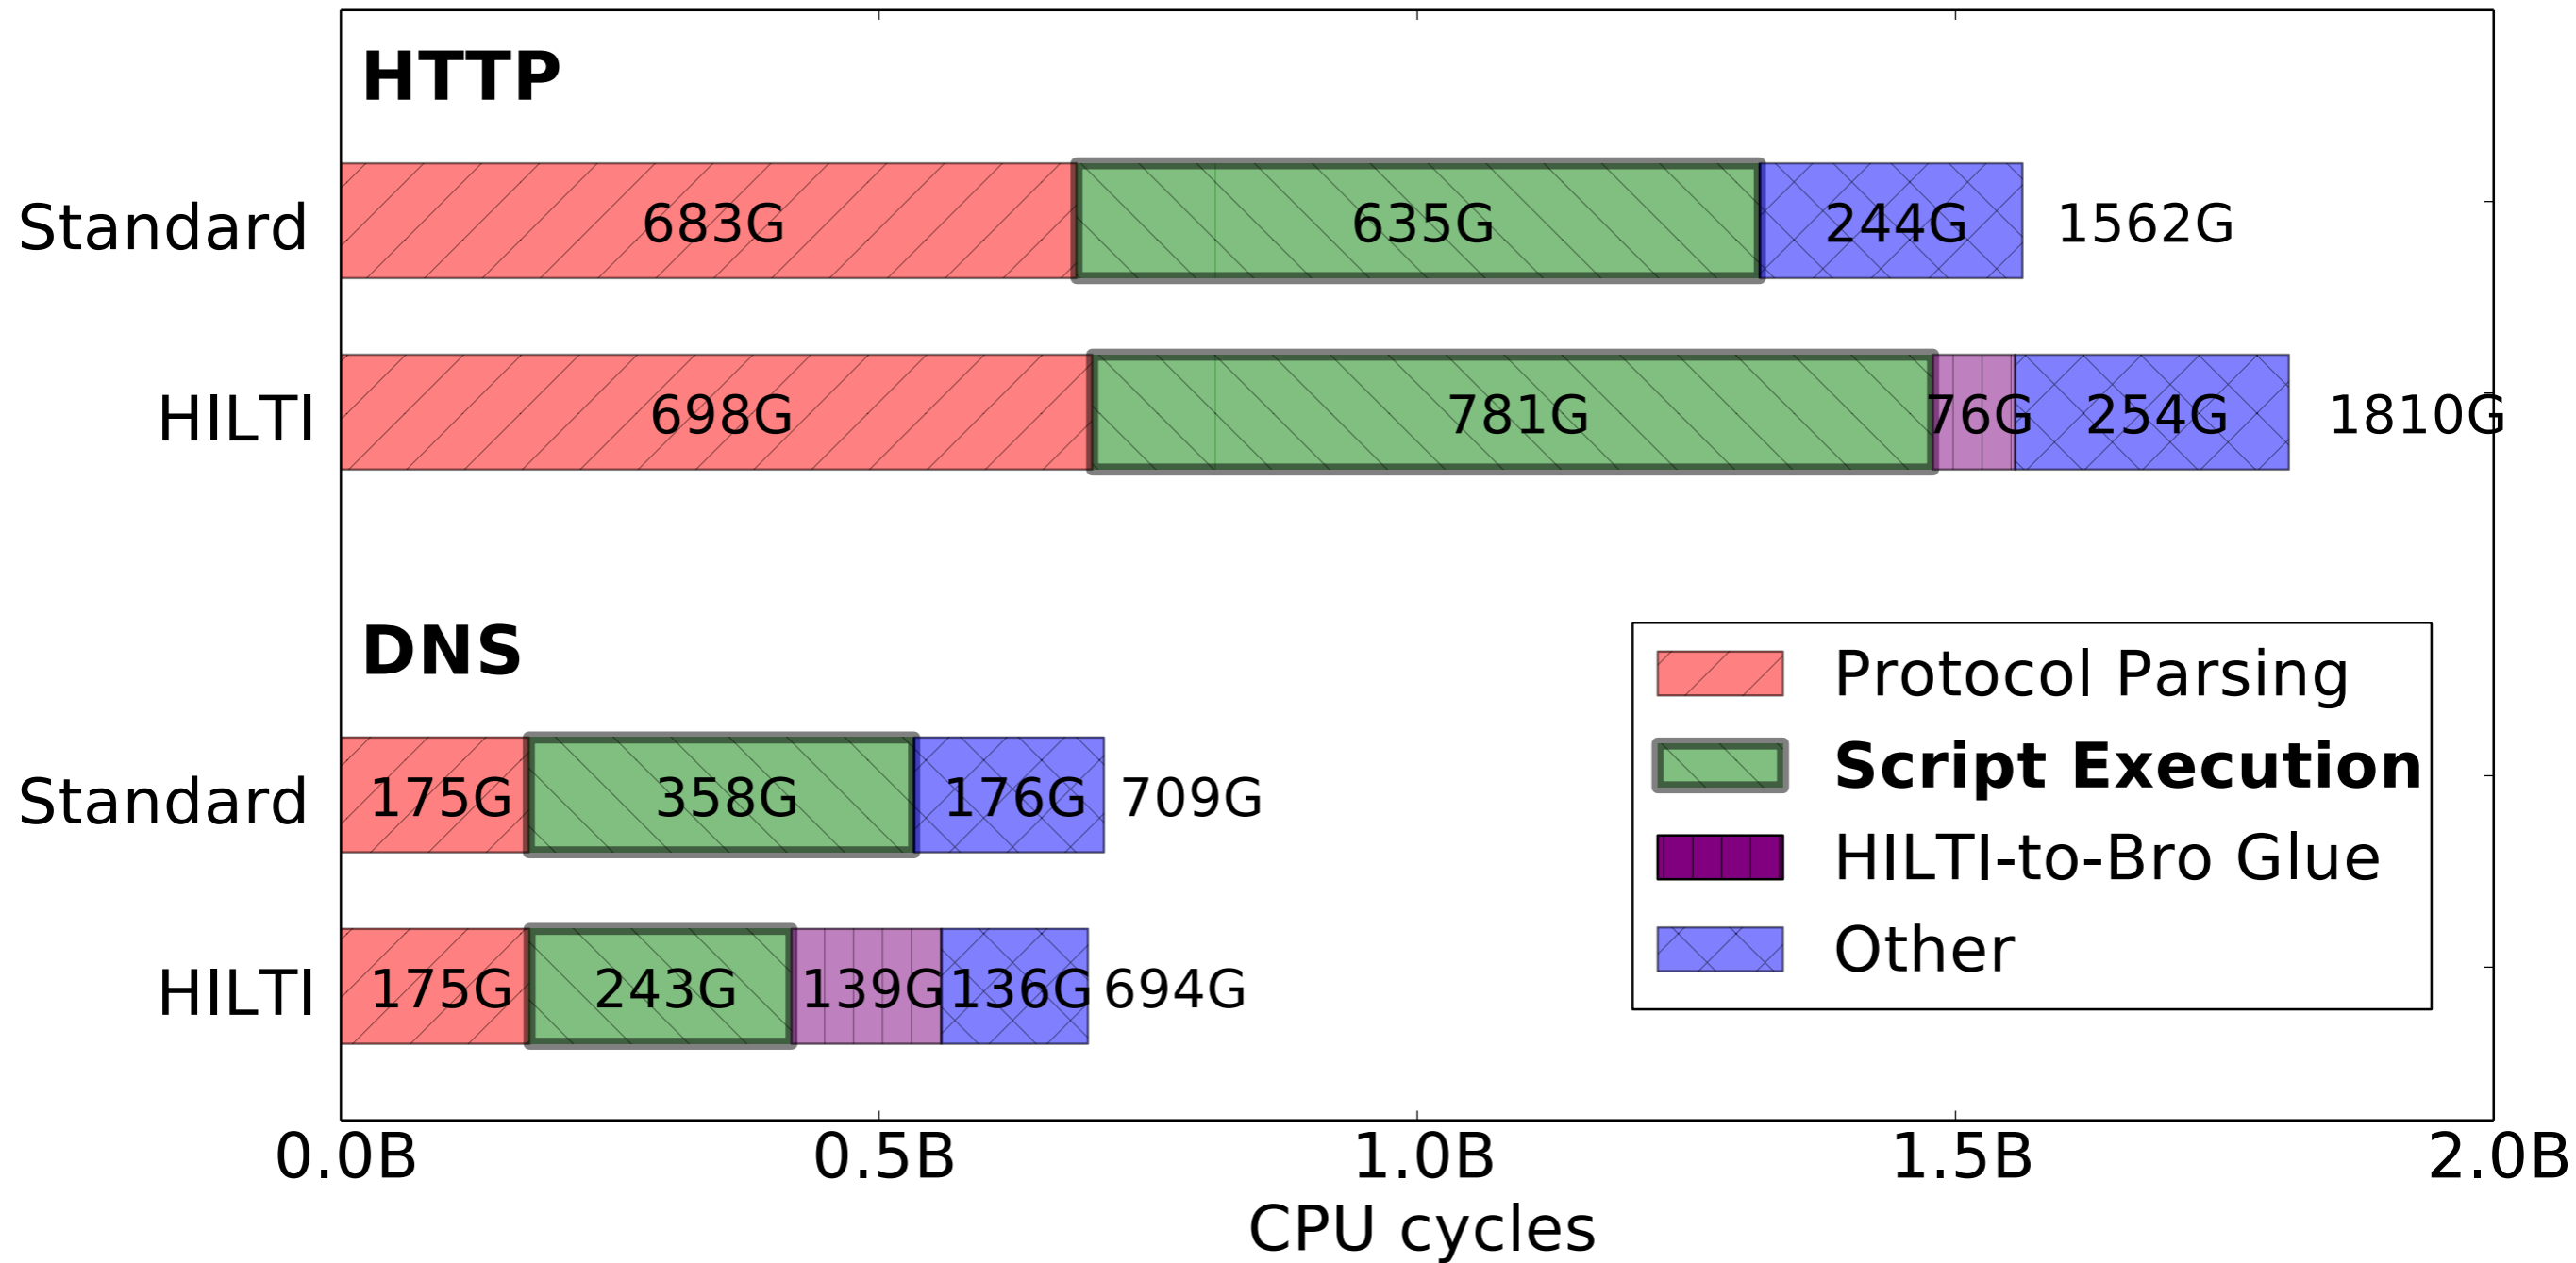
Protocol Parsing




Protocol Parsing



Bro Scripts



Bro Scripts



Summary

HILTI: A new platform for network traffic analysis.

A compiler-target for host applications to leverage.

Provides common data structures and control flow primitives.

Case studies demonstrate aptness of design.

Packet filter, stateful firewall, protocol parsing, Bro scripts.

Initial performance experiments encouraging.

Not too different from native applications.

It's still a prototype, with lots of potential.

The HILTI Vision

Performance
via Abstraction

Transparent improvement under the hood.
Integration of non-standard hardware.
High-level compiler optimizations.
Automatic parallelization.*

Facilitate Reuse

Means and glue to share functionality.
HILTI library of common high-level components.

* De Carli/Sommer/Jha: “Beyond Pattern Matching: A Concurrency Model for Stateful Deep Packet Inspection”. ACM CCS 2014.

The HILTI Vision

Performance
via Abstraction

Transparent improvement under the hood.
Integration of non-standard hardware.
High-level compiler optimizations.
Automatic parallelization.*

Facilitate Reuse

Means and glue to share functionality.
HILTI library of common high-level components.



HILTI is available under BSD license at

<http://www.icir.org/hilti>

* De Carli/Sommer/Jha: “Beyond Pattern Matching: A Concurrency Model for Stateful Deep Packet Inspection”. ACM CCS 2014.