

# Formal Verification of Cloud Services

Bryan Parno



<https://www.chase.com/>


Personal | Business | Commercial

CHASE 

CHASE 


### Online and Mobile Security

- We periodically review our operations and business practices to make sure they comply with the corporate policies and procedures we follow to protect confidential information




# Large-Scale Data Breaches Affect Millions of Users

Number of compromised records in recent large-scale data breaches



# The Ironclad Project



Chris Hawblitzel



Jon Howell



Manos Kapritsos



Rustan Leino




Jay Lorch



Arjun Narayan




**Bryan Parno**



Michael Lowell Roberts



Srinath Setty



Danfeng Zhang



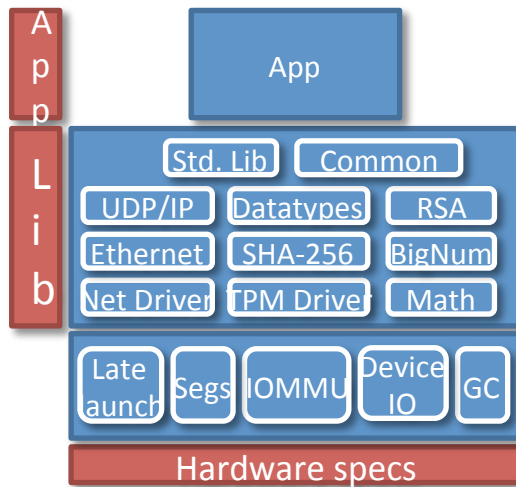
Brian Zill



# The Ironclad Project


## Ironclad Apps

[OSDI 2014]




## IronFleet

[SOSP 2015]



## Everest HTTPS







An *Ironclad app* guarantees to remote parties that every instruction it executes adheres to a high-level security spec.



# Ironclad combines:

- Late launch
- Secure hardware
- Software verification




```
push ebp  
mov ebp, esp  
sub esp, 4  
mov eax, 8
```



# Ironclad combines:

- Late launch
- Secure hardware
- Software verification

Secure Remote Equivalence





```
push ebp  
mov ebp, esp  
sub esp, 4  
mov eax, 8
```



Reasonable effort

Entire software stack





# Verification goals

- End-to-end security
  - Complete
  - Low-level
- Rapid development

• Non-goal: Verify existing code

• Long-term: Performance matches unsafe code




```
push ebp  
mov ebp, esp  
sub esp, 4  
mov eax, 8
```



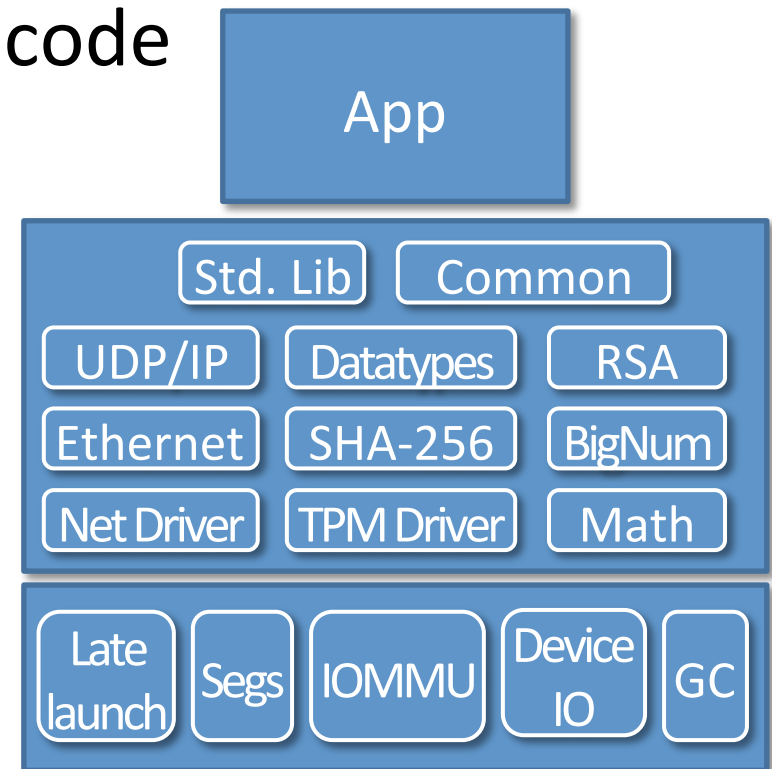
# Prior work

- Small components
  - E.g., RSA-OAEP
  - Localized guarantee
- Single layers
  - No end-to-end guarantee

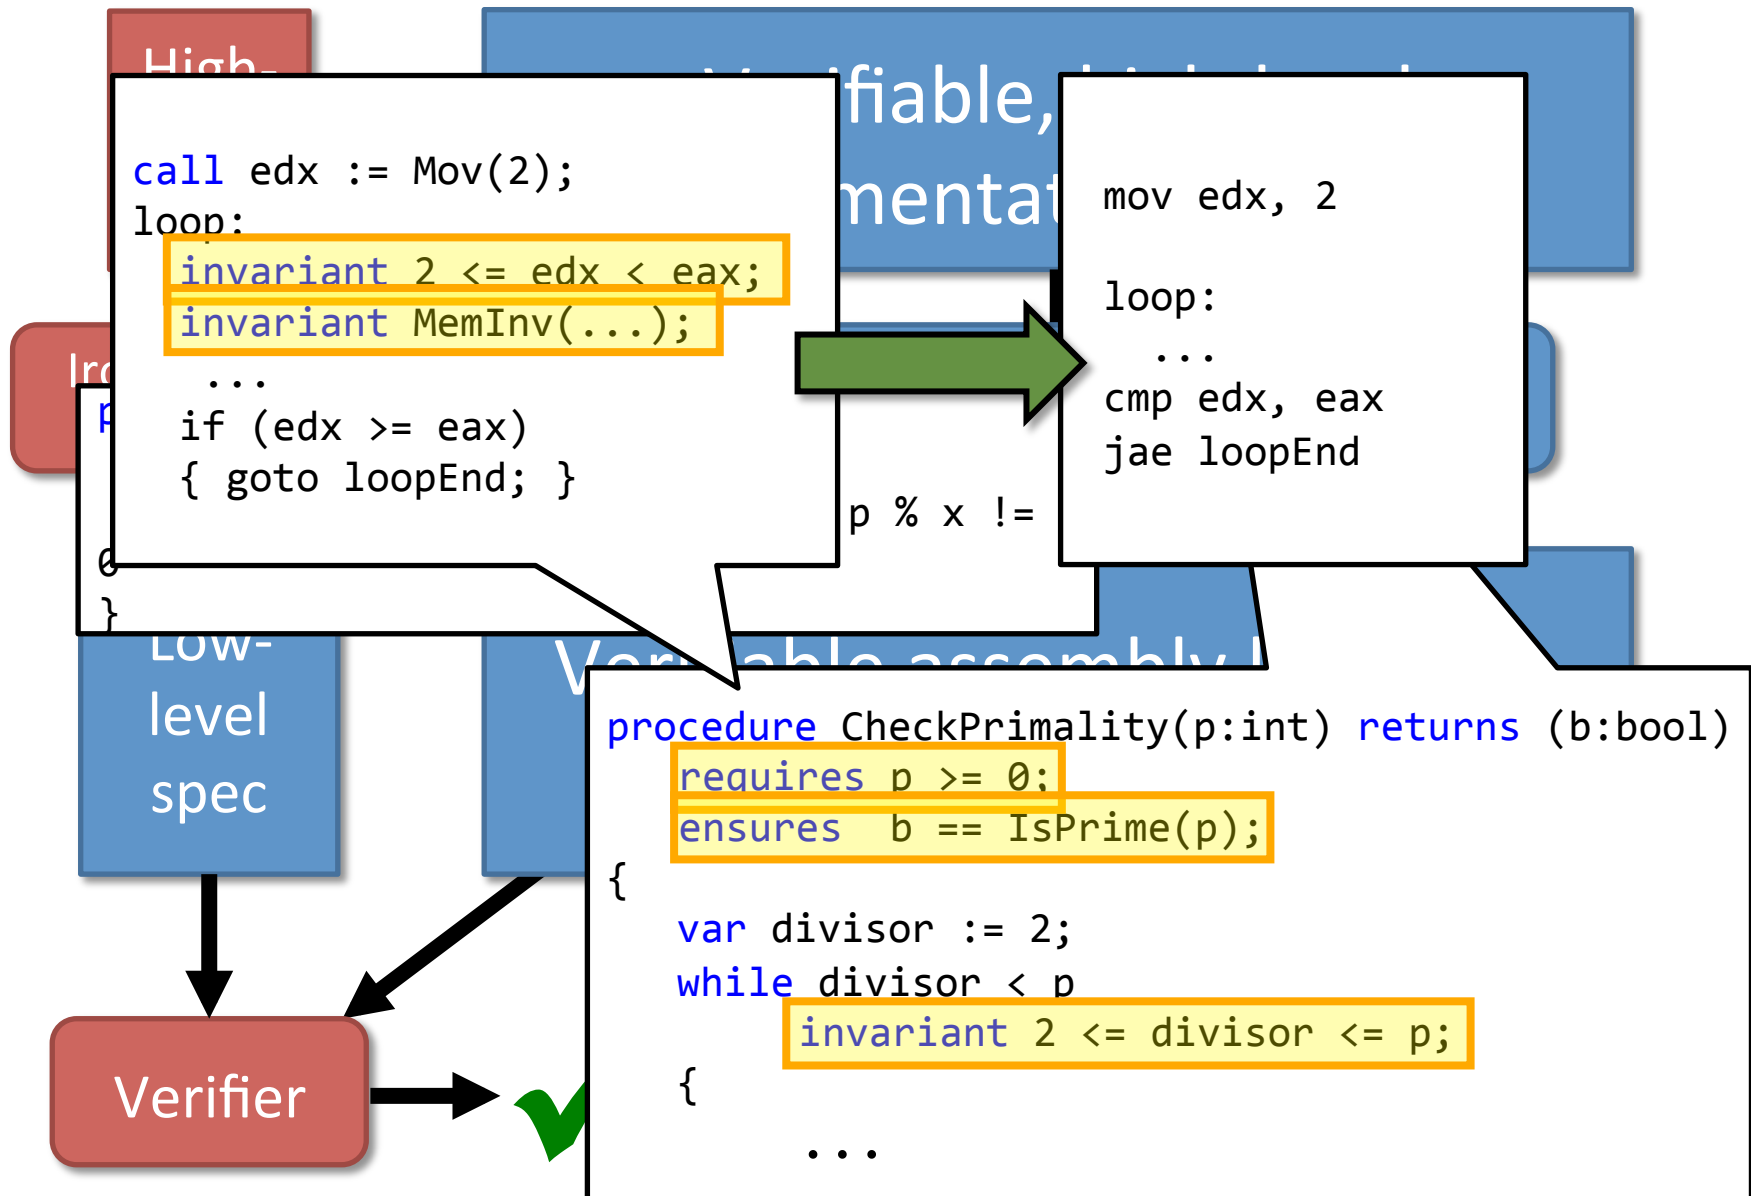



# Our formal, end-to-end guarantee

- End-to-end secure communication with provably secure assembly code
- Implies:
  - No buffer overflows
  - No code injection
  - No type-safety flaws
  - No information disclosures
  - No crypto impl flaws
  -

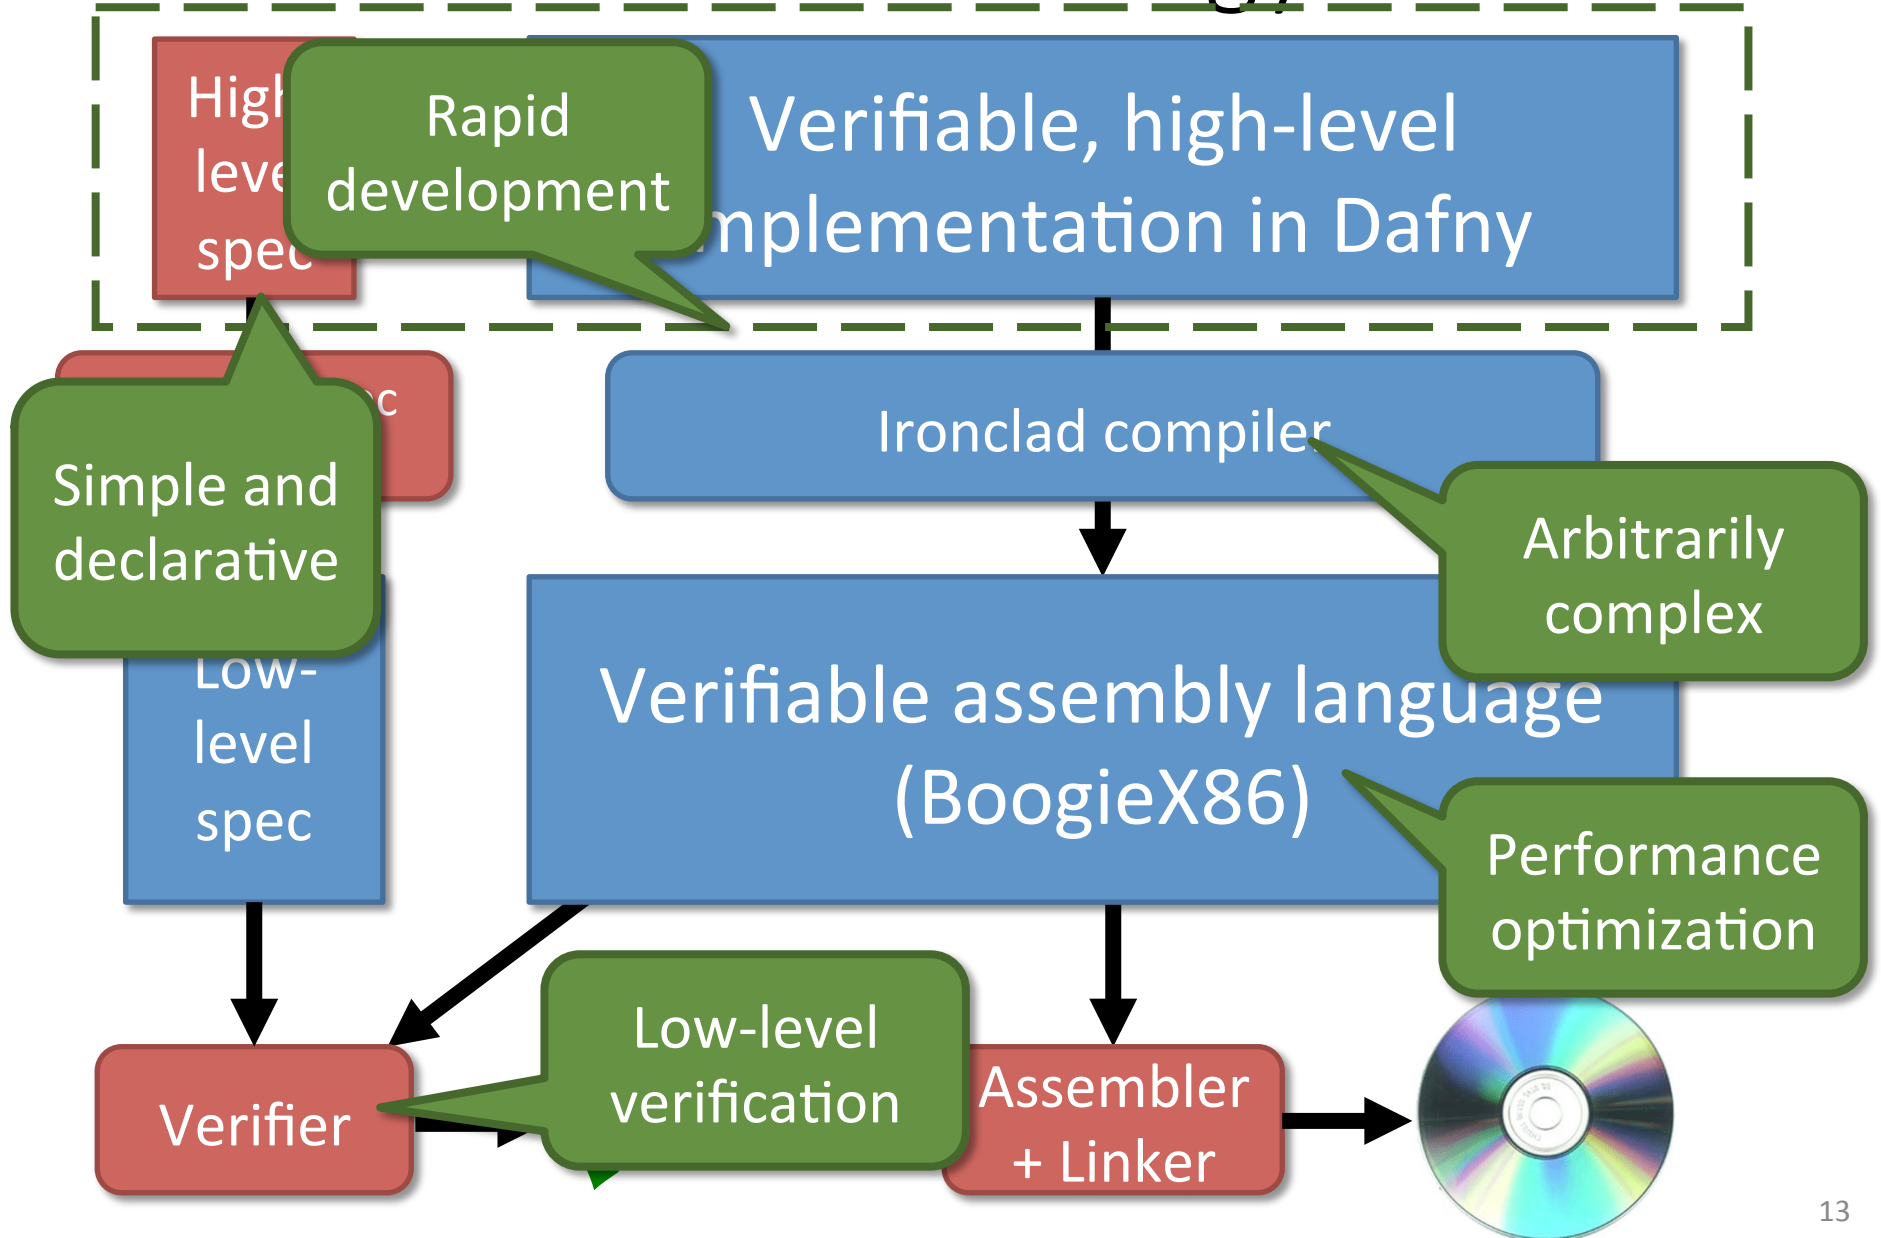


# Verification methodology





# Verification methodology: Benefits



# Architecture

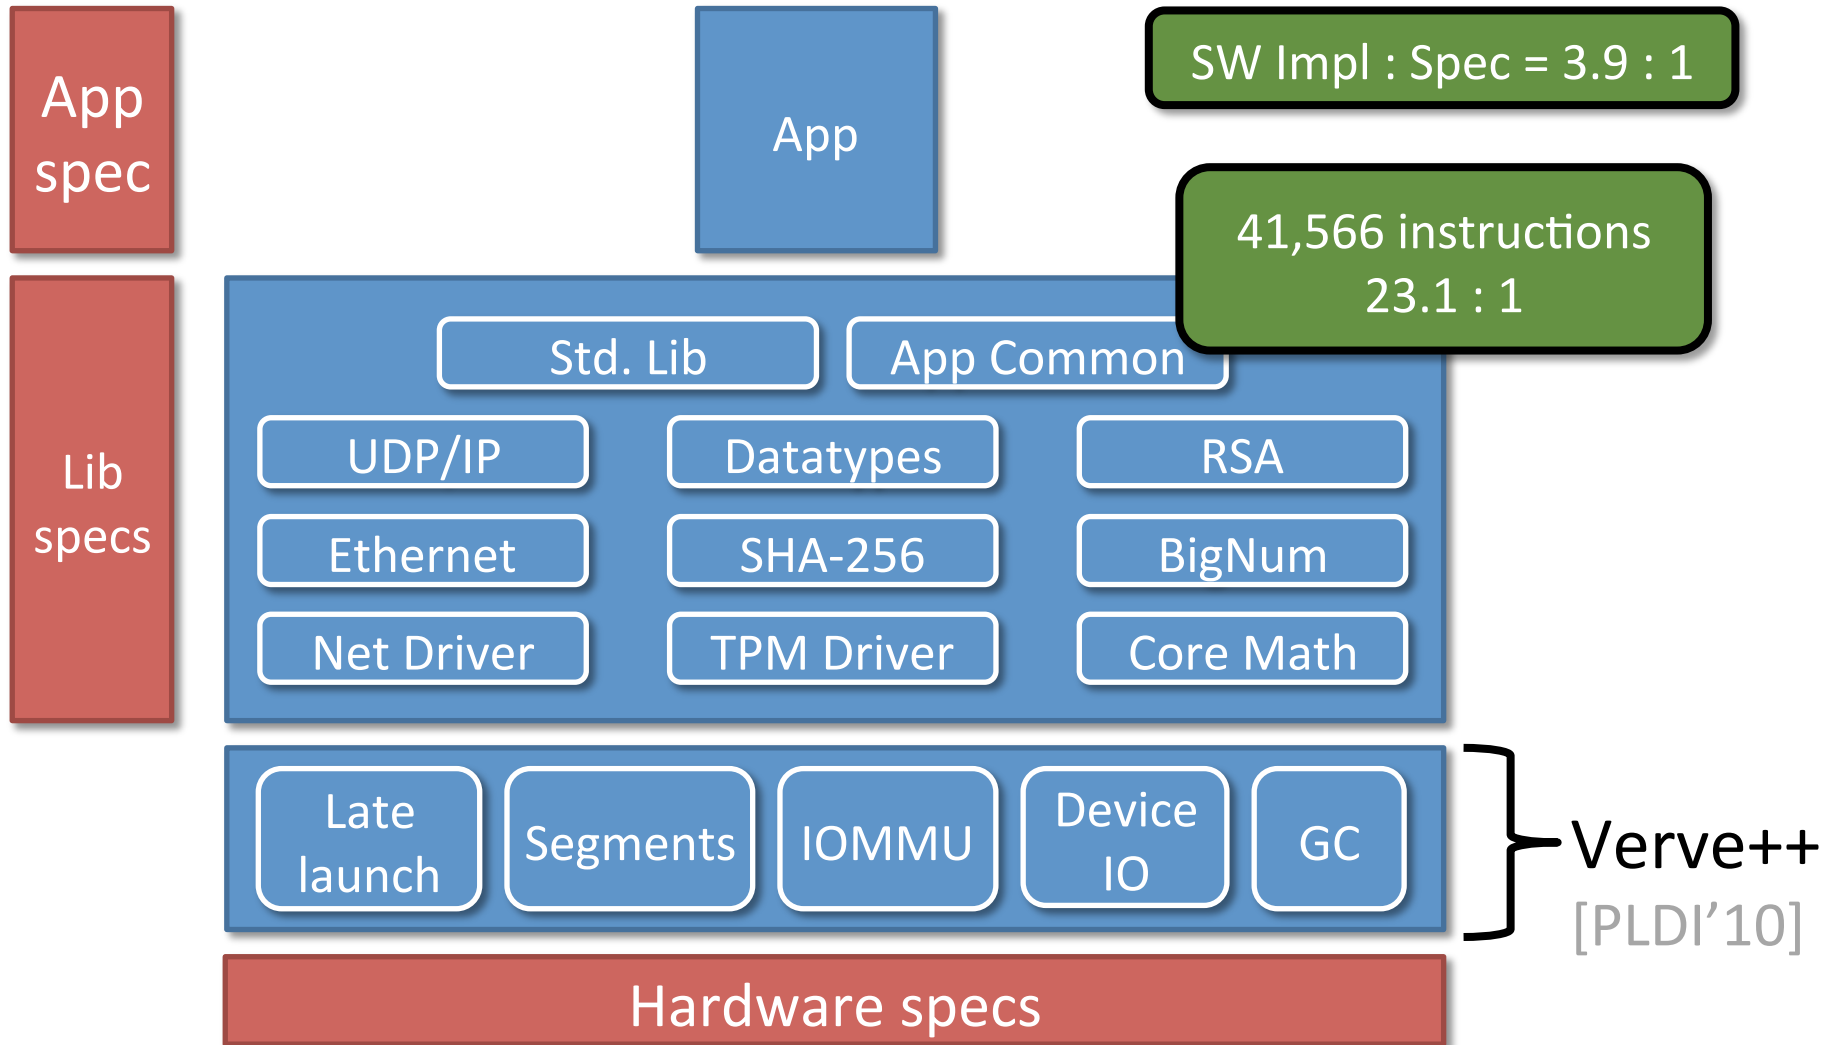
Trusted Spec  
Implementation

## Software

1796 LoC  
6971 LoC

## Hardware

1750 LoC



# Ironclad Apps

## Password Protector


password letmein  
1234567890 dragon  
12345678901111  
abc123456789 baseball  
monkeys loveyou  
qwerty trustno1



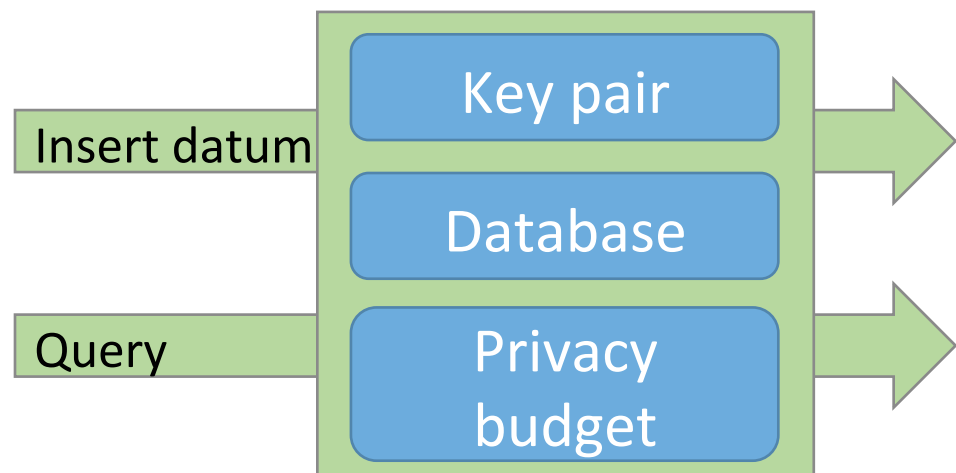
## Notary



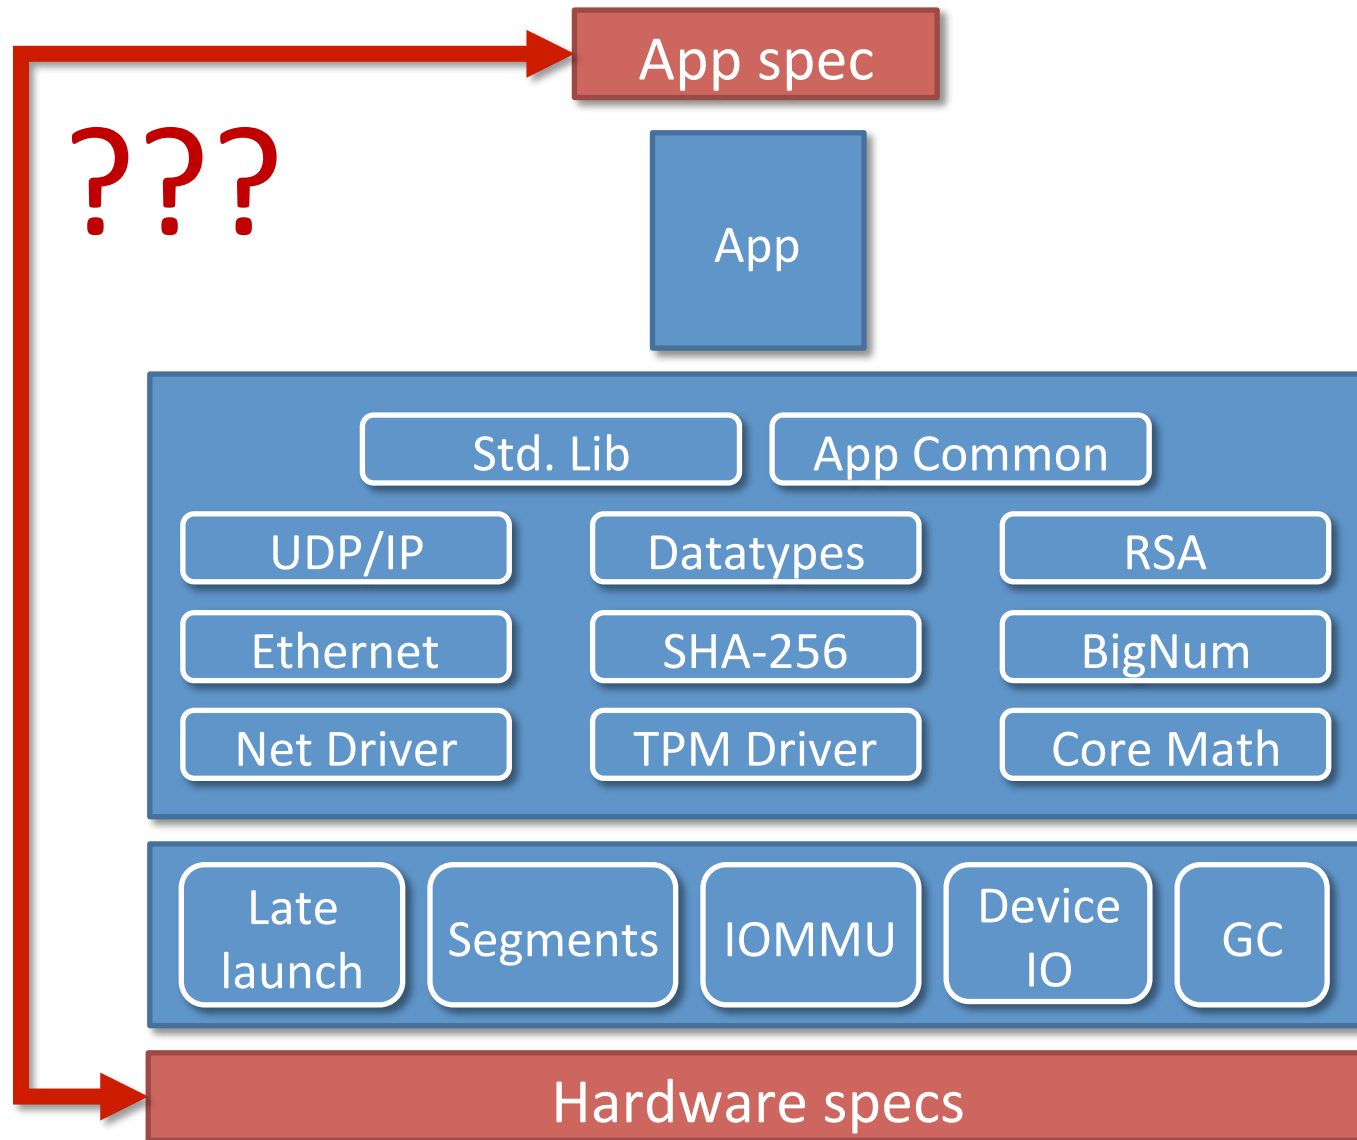
## Trusted Incrementer



## Differentially Private DB




# Key Challenge: Specifying security



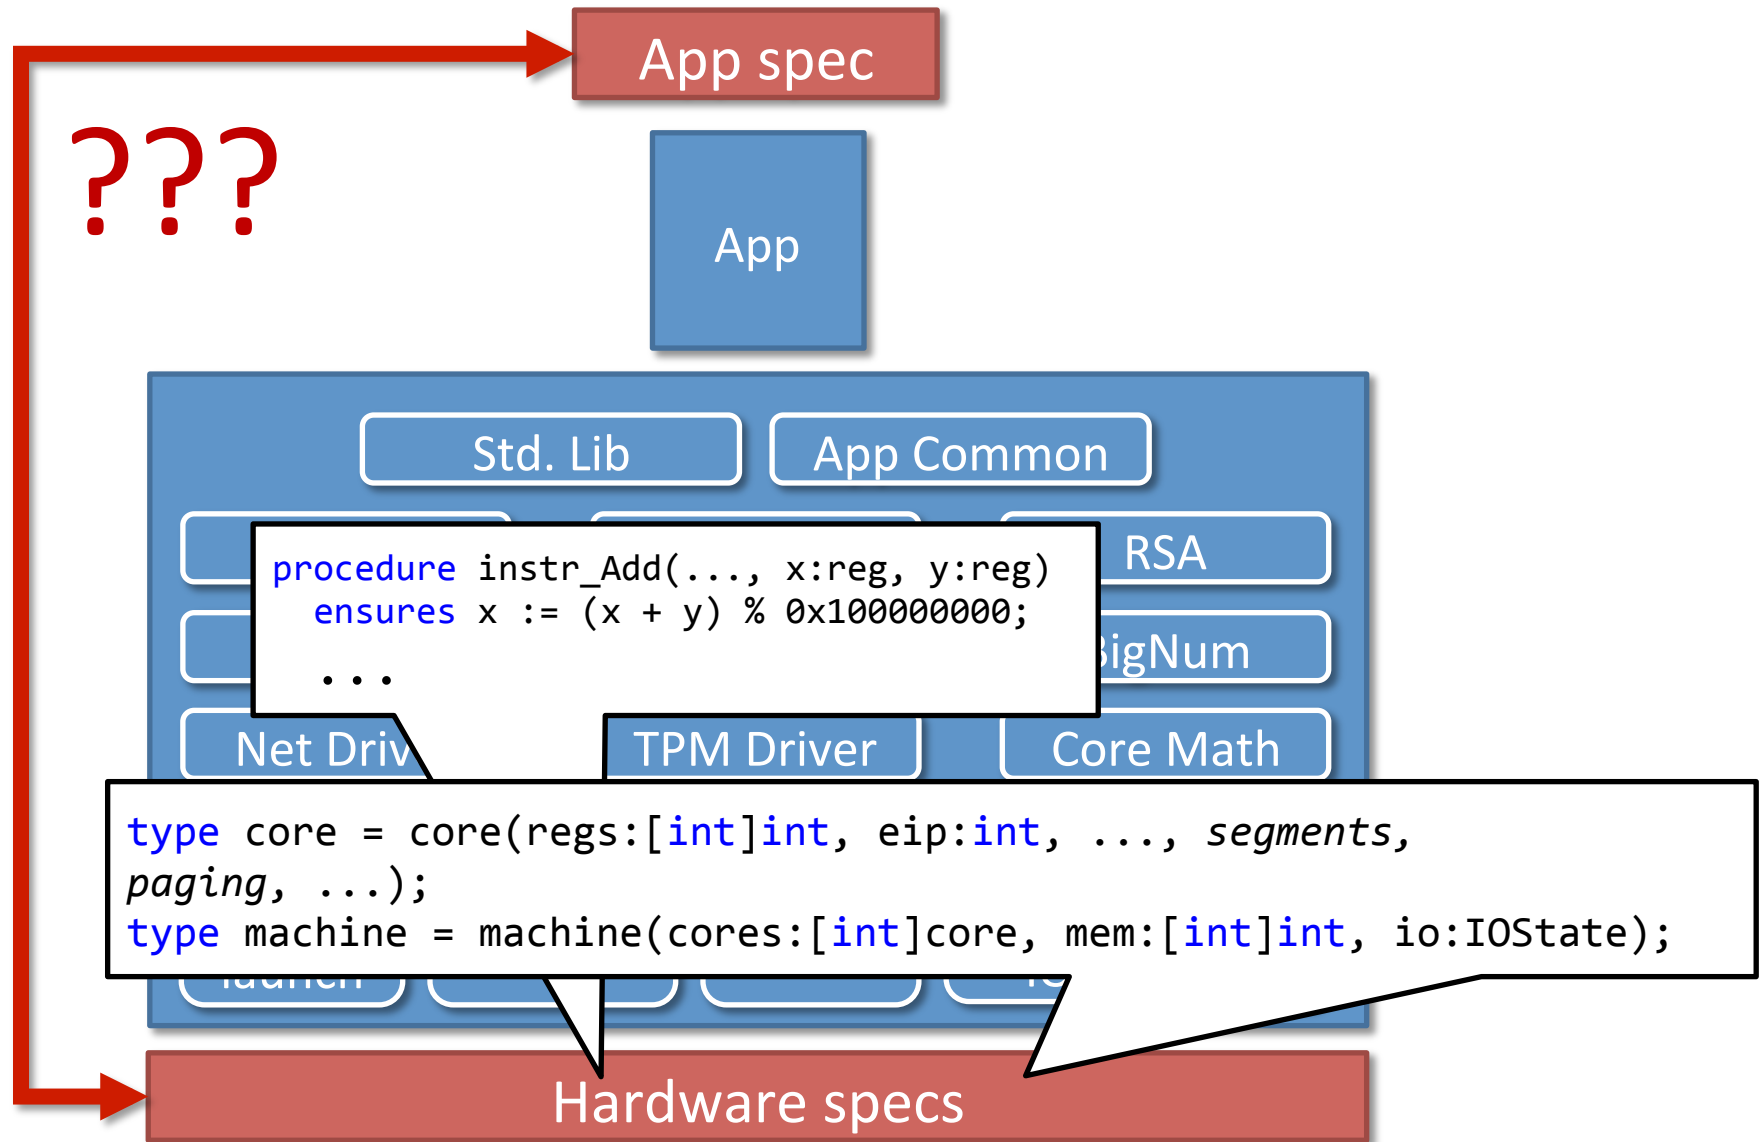
# Specifying the DiffPriv DB

```
datatype DiffPrivState(  
  keys:RSA,  
  budget:real,  
  db:seq<Row>  
)  
  
predicate Transition(old_state:DiffPrivState,  
  new_state:DiffPrivState,  
  request:Request,  
  reply:Reply)  
{  
  match request  
  ...  
  case Query =>  
    request.spend < old_state.budget  
  
    && new_state.budget == old_state.budget - request.spend  
  
    && reply.answer == RunQuery(request.query, old_state.db)  
      + ComputeNoise(request.spend)  
  
    && ...  
}
```

App spec




# Key Challenge: Specifying security




# Connecting the app to the hardware

```
procedure instr_inb(..., x:reg)
```



```
procedure instr_outb(..., x:reg)  
requires ????
```

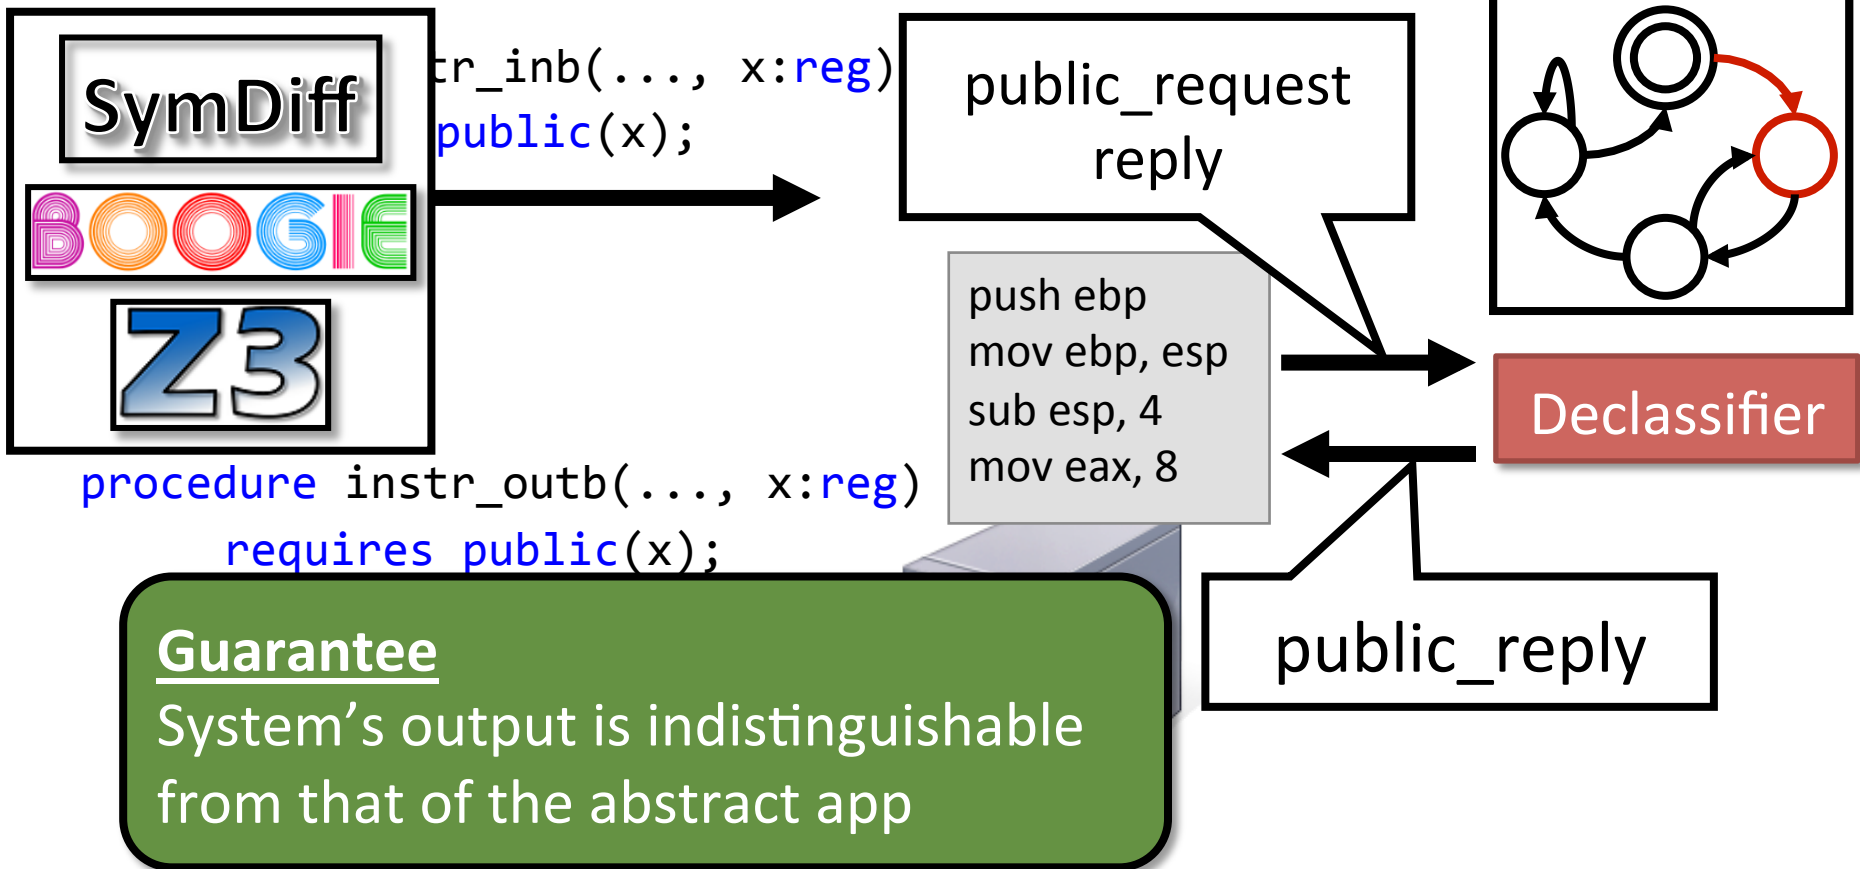
```
push ebp  
mov ebp, esp  
sub esp, 4  
mov eax, 8
```



# Connecting the app to the hardware

## Key idea

State-machine-based relational verification



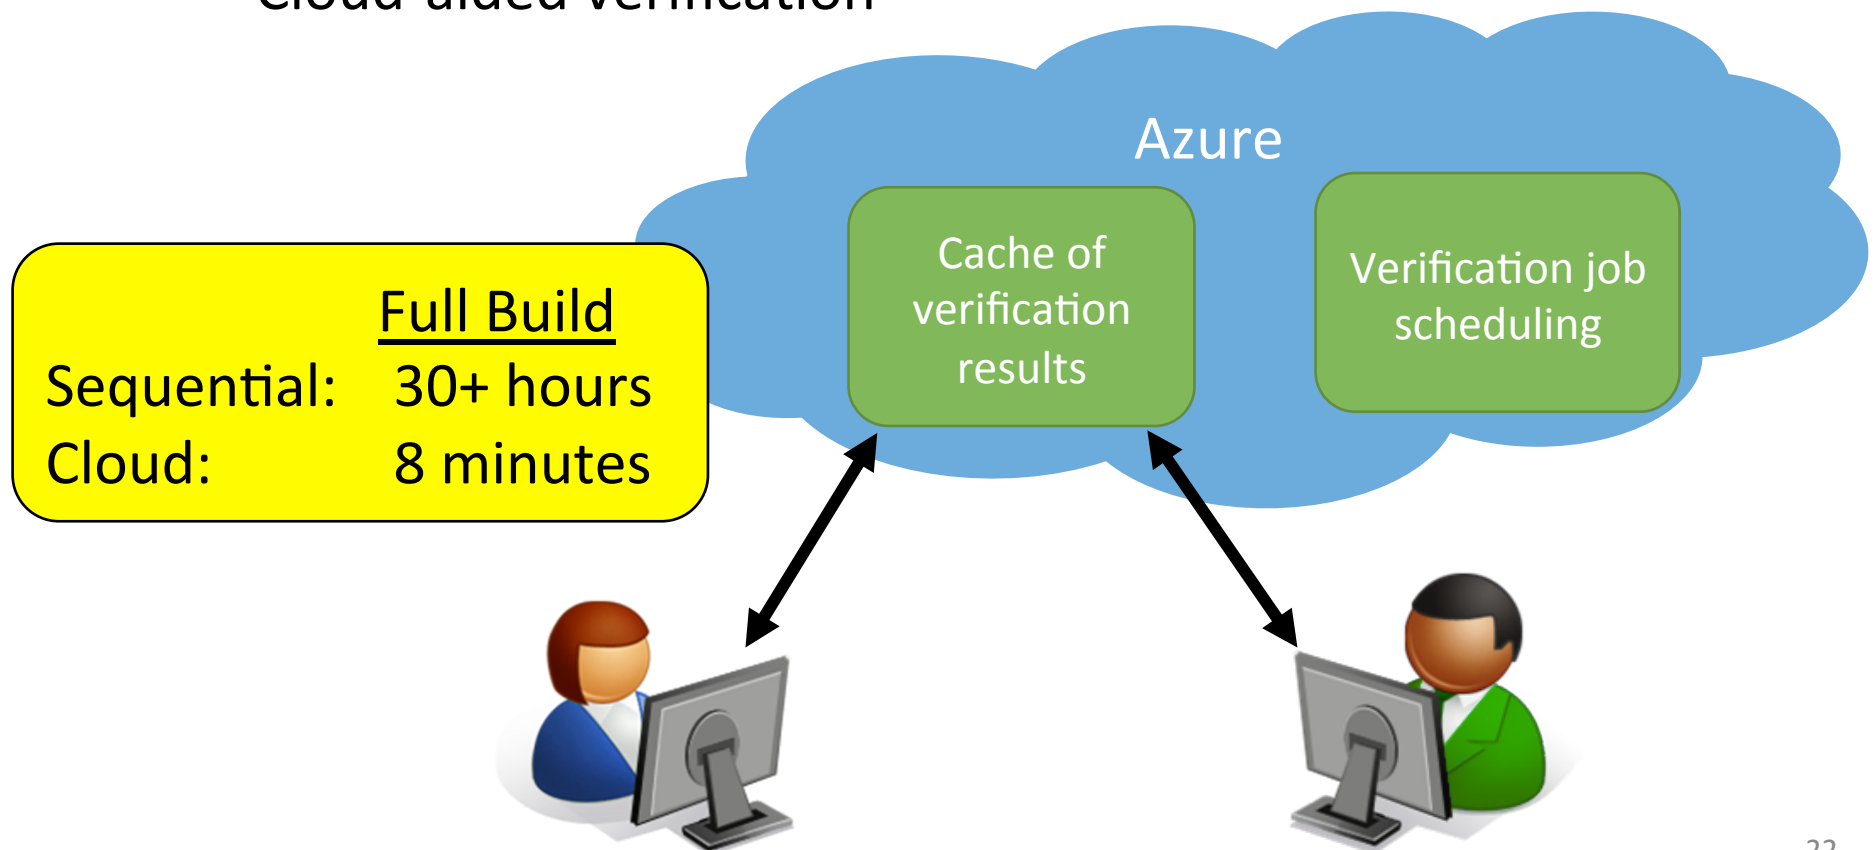


# Evaluation questions

- How does verification affect development?
- How does verification affect performance?

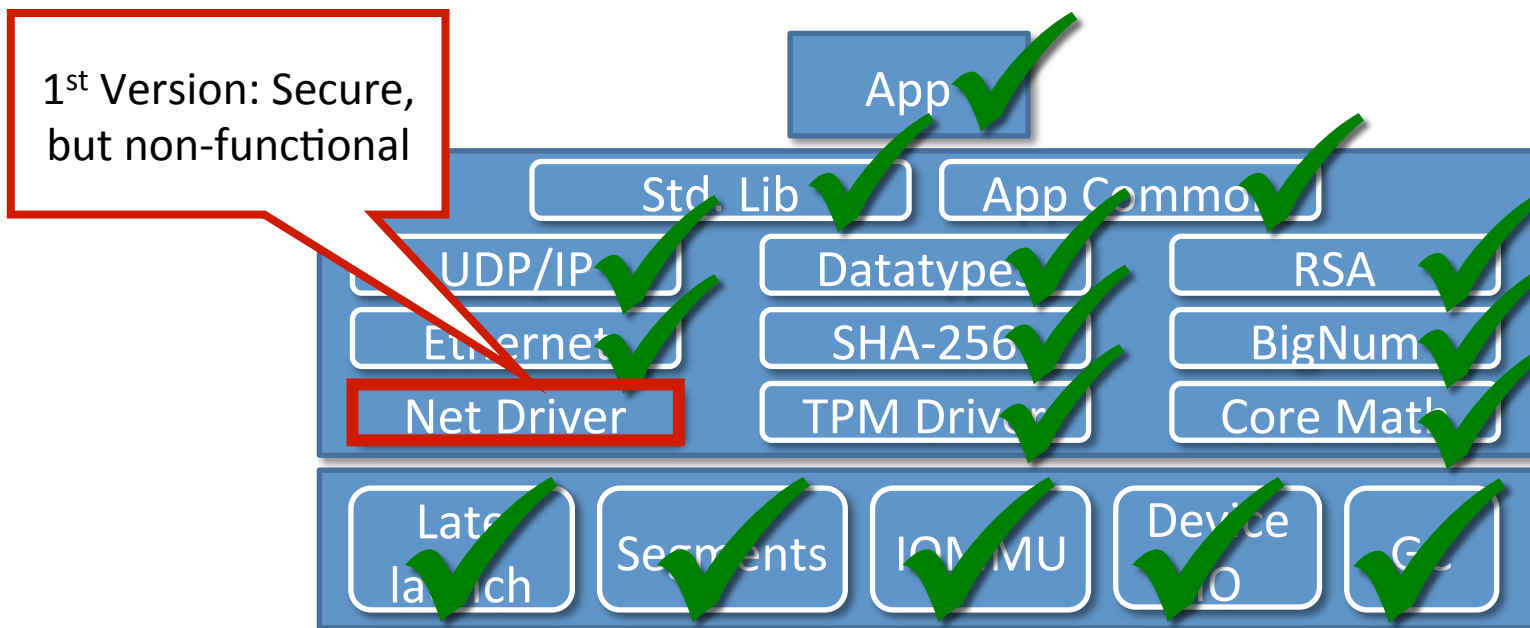
# Eval: Developer experience

- Rapid feedback via aggressive caching:
  - Method-level via Dafny's IDE plugins
  - File-level via new Dafny `include` feature
  - Cloud-aided verification



# Eval: Developer experience

Verified software works!




# Eval: Proof burden

~3 person-years


Previously 22+ pys



## Proof hints : Implementation LoC




# Eval: Performance



# The Ironclad Project


## Ironclad Apps

[OSDI 2014]




## IronFleet

[SOSP 2015]



## Everest HTTPS





# IronFleet: Proving Practical Distributed Systems Correct

[SOSP 2015]

## Buggy code update triggers Bing, Yahoo outage

Jan 5, 2015 9:00 AM  
Filed under [Search](#)

 Like 9



Tags

## Google blames software bug for Friday night Gmail outage

Promoted Content

Innovation at work



SHARE  TWITTER

WRITTEN BY  
Caroline  
Donnelly

Search users

JUST IN With the launch of its Edge browser, Microsoft ends an era

Topic: Amazon Follow via: 

## Amazon Web Services suffers outage, takes down Vine, Instagram, others with it

**Summary:** The cloud giant suffered an outage for about an hour on Sunday, showing once again the perils of an outsourced cloud service, as many AWS customers went down with it.

By  Zack Whittaker for [Between the Lines](#) | August 26, 2013 -- 13:22 GMT (06:22 PDT)

 Follow @zackwhittaker 10.6K followers [Get the ZDNet Product Watch newsletter now](#)

Comments 12  Share on Facebook 0  Tweet 1  Share more +

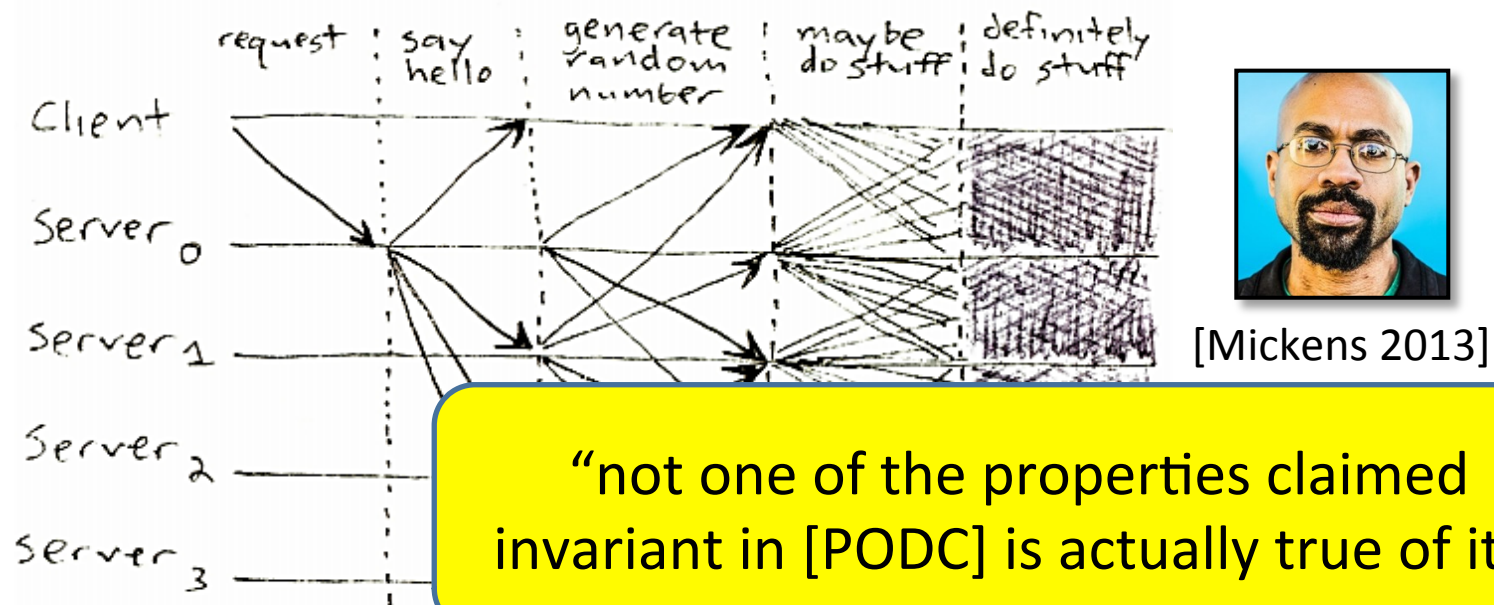


Figure 1: Typical Figure 2 from Byzantine fault paper: Our network protocol



[Zave 2015]

Under the same assumptions made in the Chord papers, the [SIGCOMM] version of the protocol is not correct, and not one of the properties claimed invariant in [PODC] is actually invariantly true of it. The [PODC] version satisfies one invariant, but is still not correct. The results are presented by means of counterexamples to the invariants in Section 4. In preparation for the results, Section 2 gives a brief summary of the protocol and failure assumptions, and Section 3 introduces the invariants.



# IronFleet




[SOSP

Implementations are correct,  
not just abstract protocols

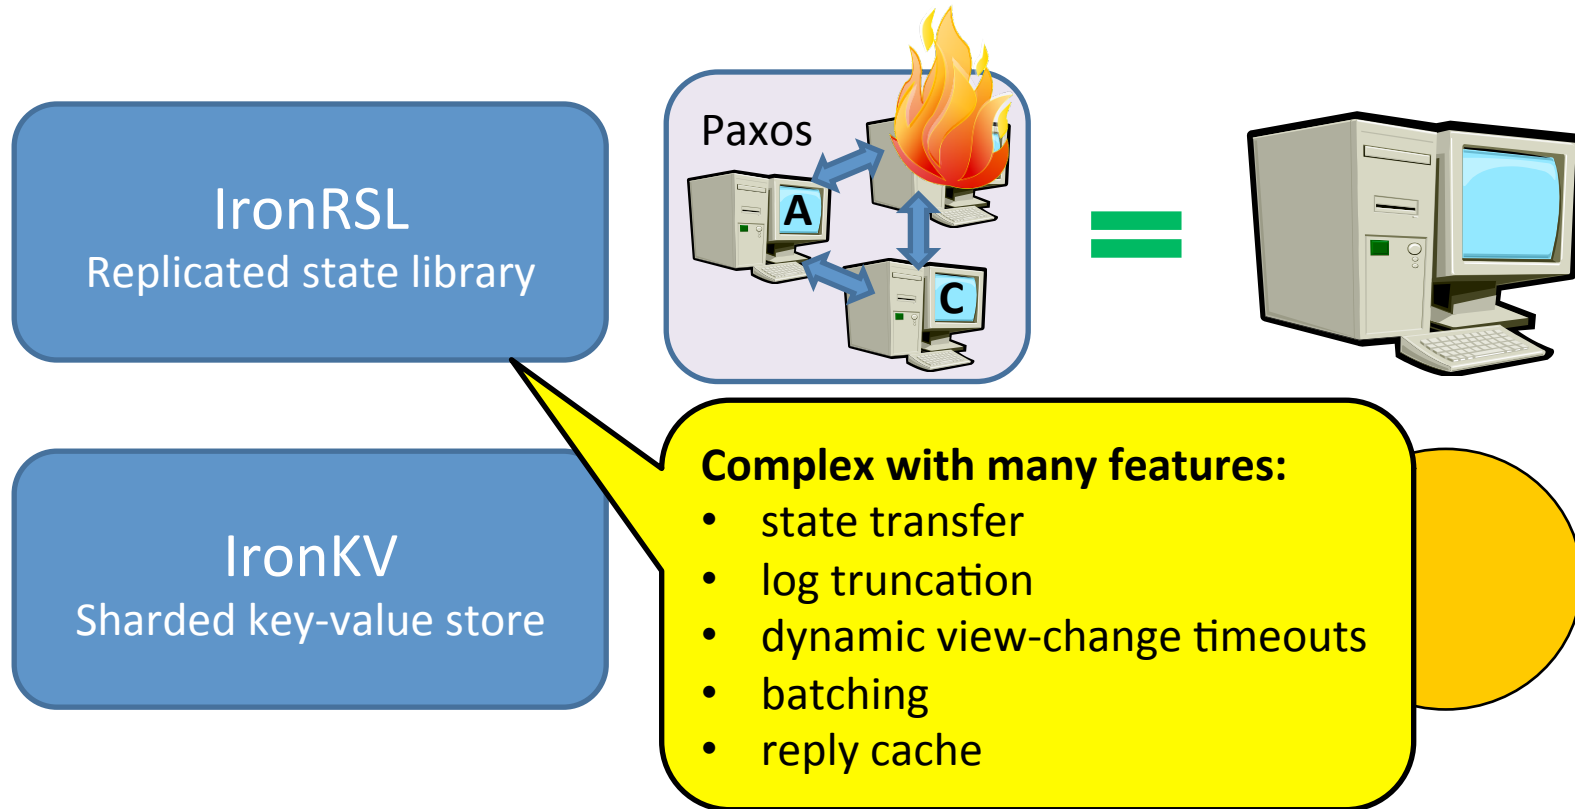
*We show how to build complex, efficient distributed systems whose implementations are provably safe and live.*

System will not take incorrect actions


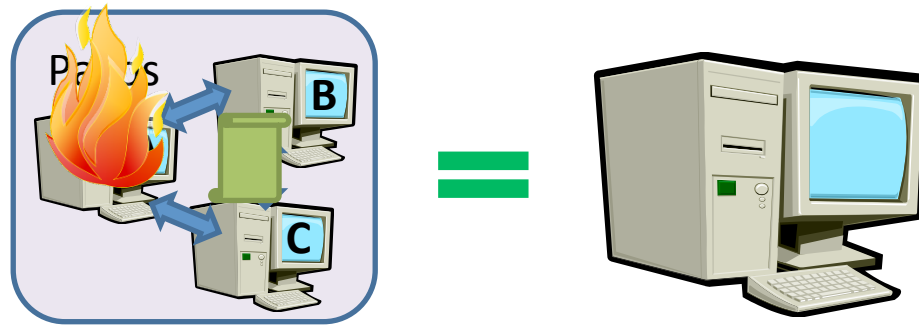
absolute



# Our verified distributed systems



# Running example: IronRSL replicated state library



Safety property:  
Equivalence to single machine

Liveness property:  
Clients eventually get replies

# Verification rules out *all* bugs by construction

Invariant violations

Race conditions

Integer overflow

Buffer overflow

Parsing errors

Marshalling errors

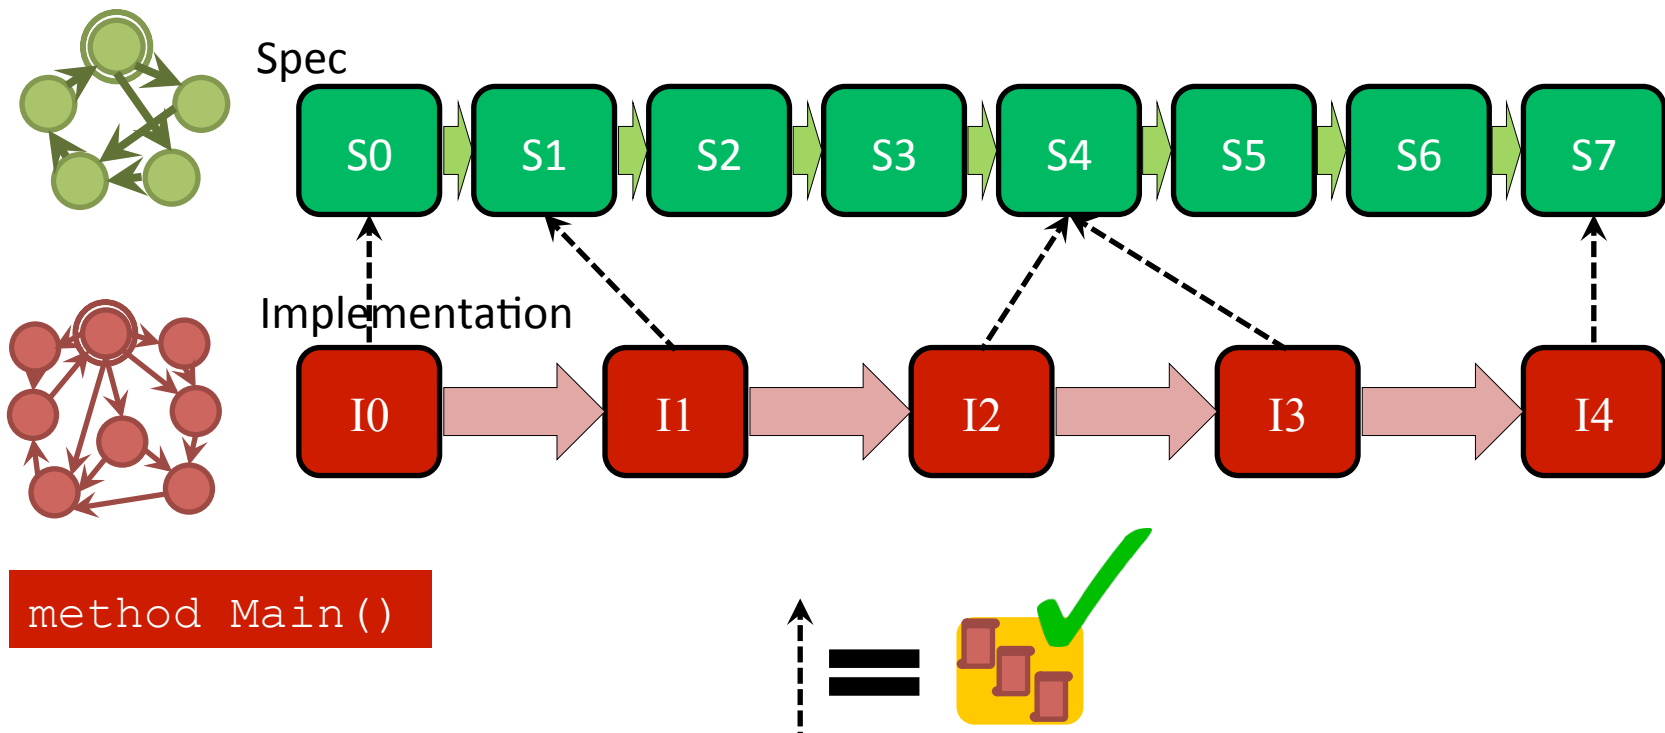
Deadlock

Livelock



# Background: Refinement

[Milner 1971, Park 1981, Lamport 1983, Lynch 1983, ...]



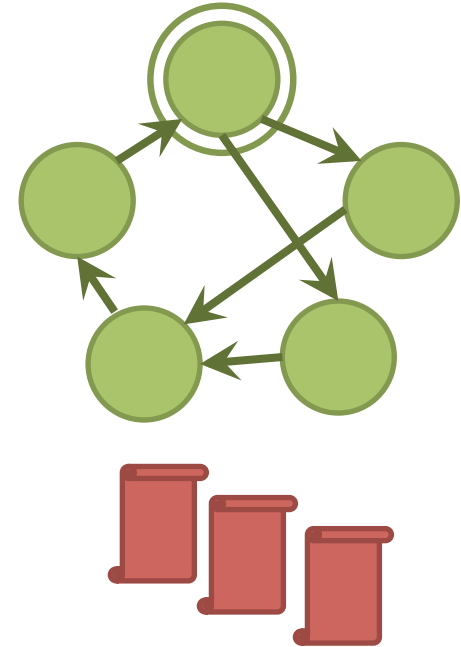
# Specification is a simple, logically centralized state machine

```
type SpecState = int
```





```
function SpecInit(s:SpecState) : bool  
{  
    s == 0  
}
```

```
function SpecNext(s:SpecState,  
                 s':SpecState) : bool  
{  
    s' == s + 1  
}
```

```
function SpecRelation(realPackets:set<Packet>,  
                    s:SpecState) : bool  
{  
    forall p, i :: p in realPackets &&  
        p.msg == MarshallCounterVal(i) ==> i <= s  
}
```



# Implementation



```
method Main()  
{  
    var s:ImplState;  
    s := ImplInit();  
    while (true) {  
        s := EventHandler(s);  
    }  
}
```

Host implementation  
is a single-threaded  
event-handler loop

# Proving correctness is hard

Subtleties of distributed protocols

Maintaining global invariants

Ensuring progress

Dealing with hosts acting concurrently

Complexities of implementation


Using efficient data structures

Memory management

Avoiding integer overflow



# Two-level refinement



# Distributed System Layer

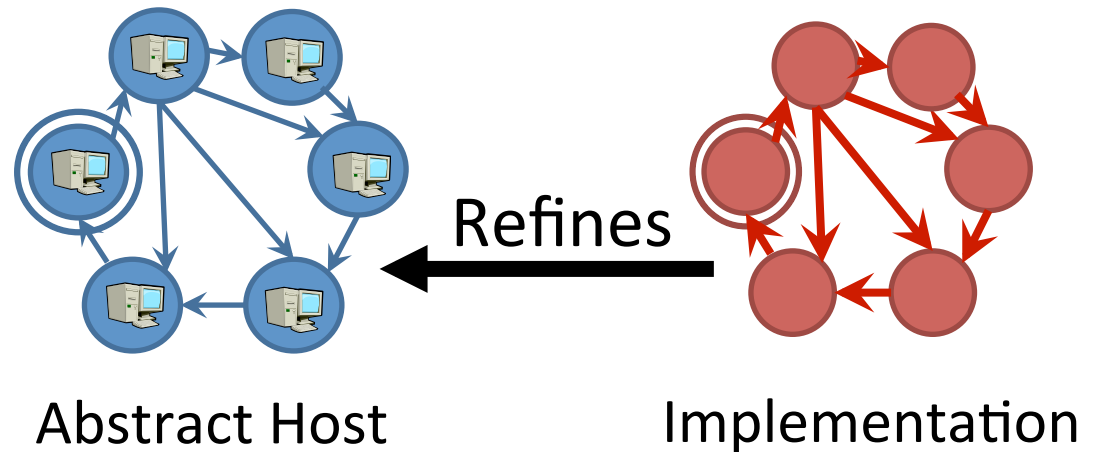
```
seq<int>
```

```
array<uint64>
```

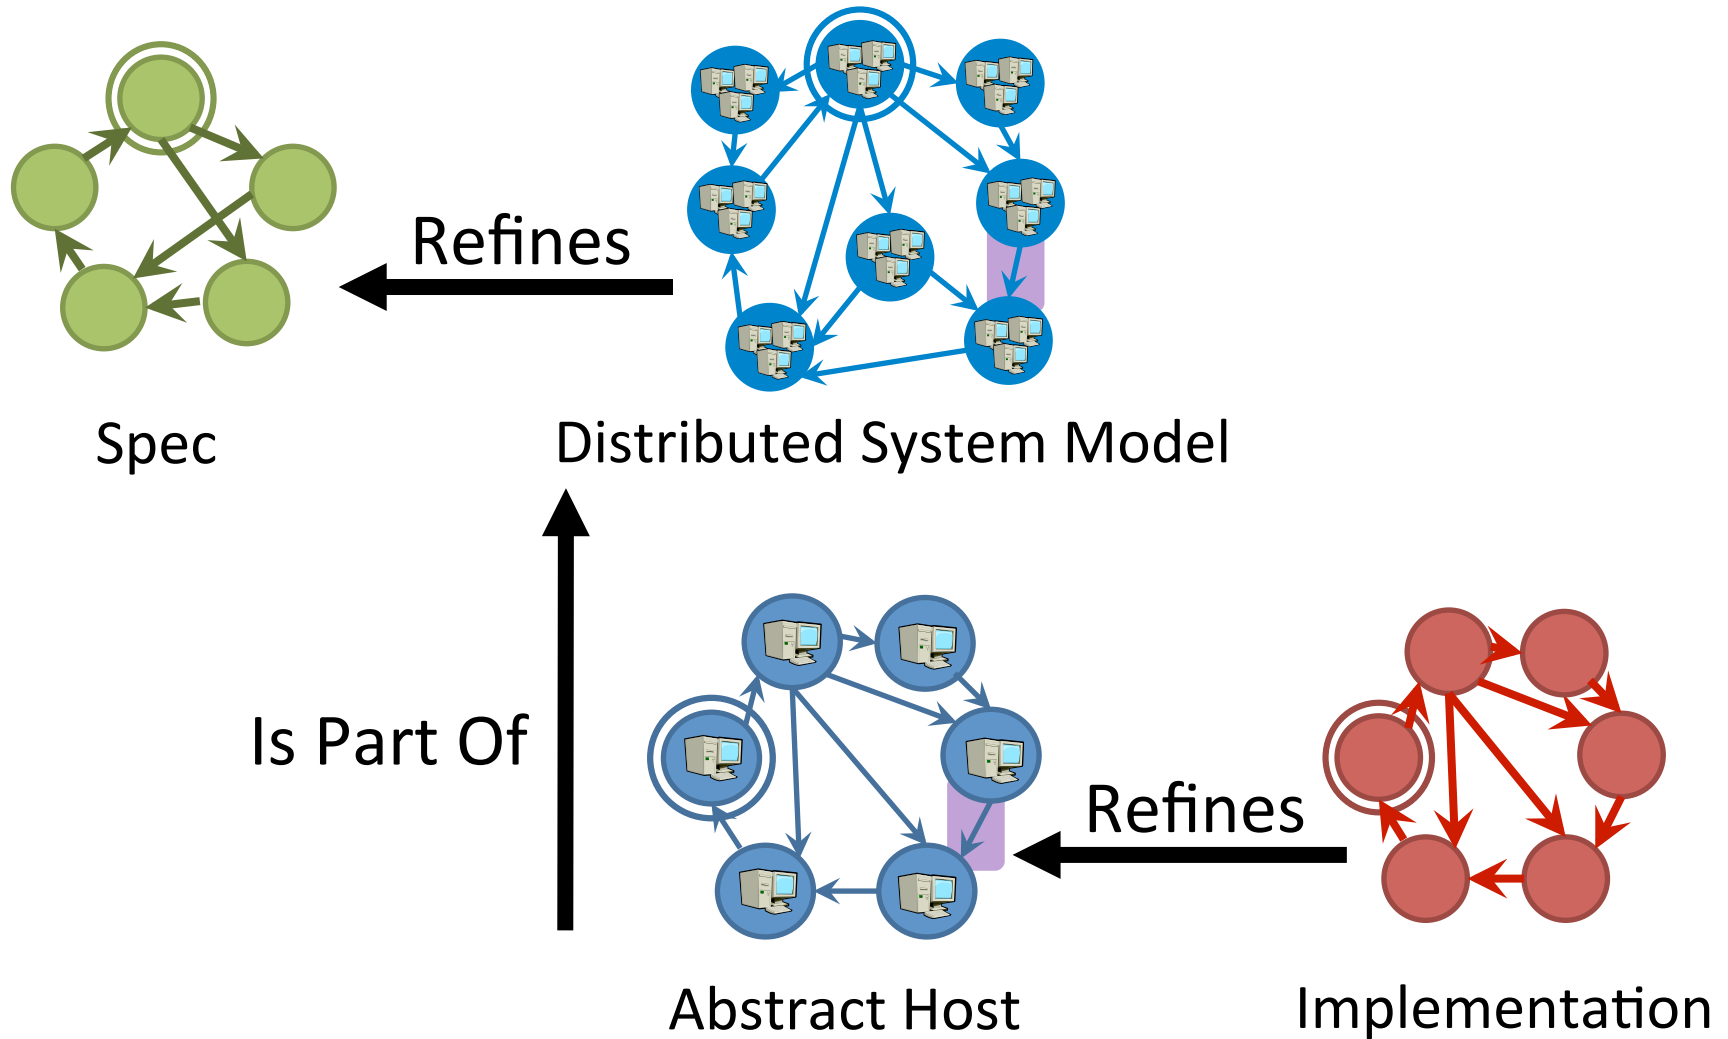
```
function ProtocolNext(s:HostState, s':HostState) : bool  
method EventHandler(s:HostState) returns (s':HostState)
```

```
type Message = MessageRequest() | MessageReply() | ...
```

```
type Packet = array<byte>
```



# Distributed System Layer



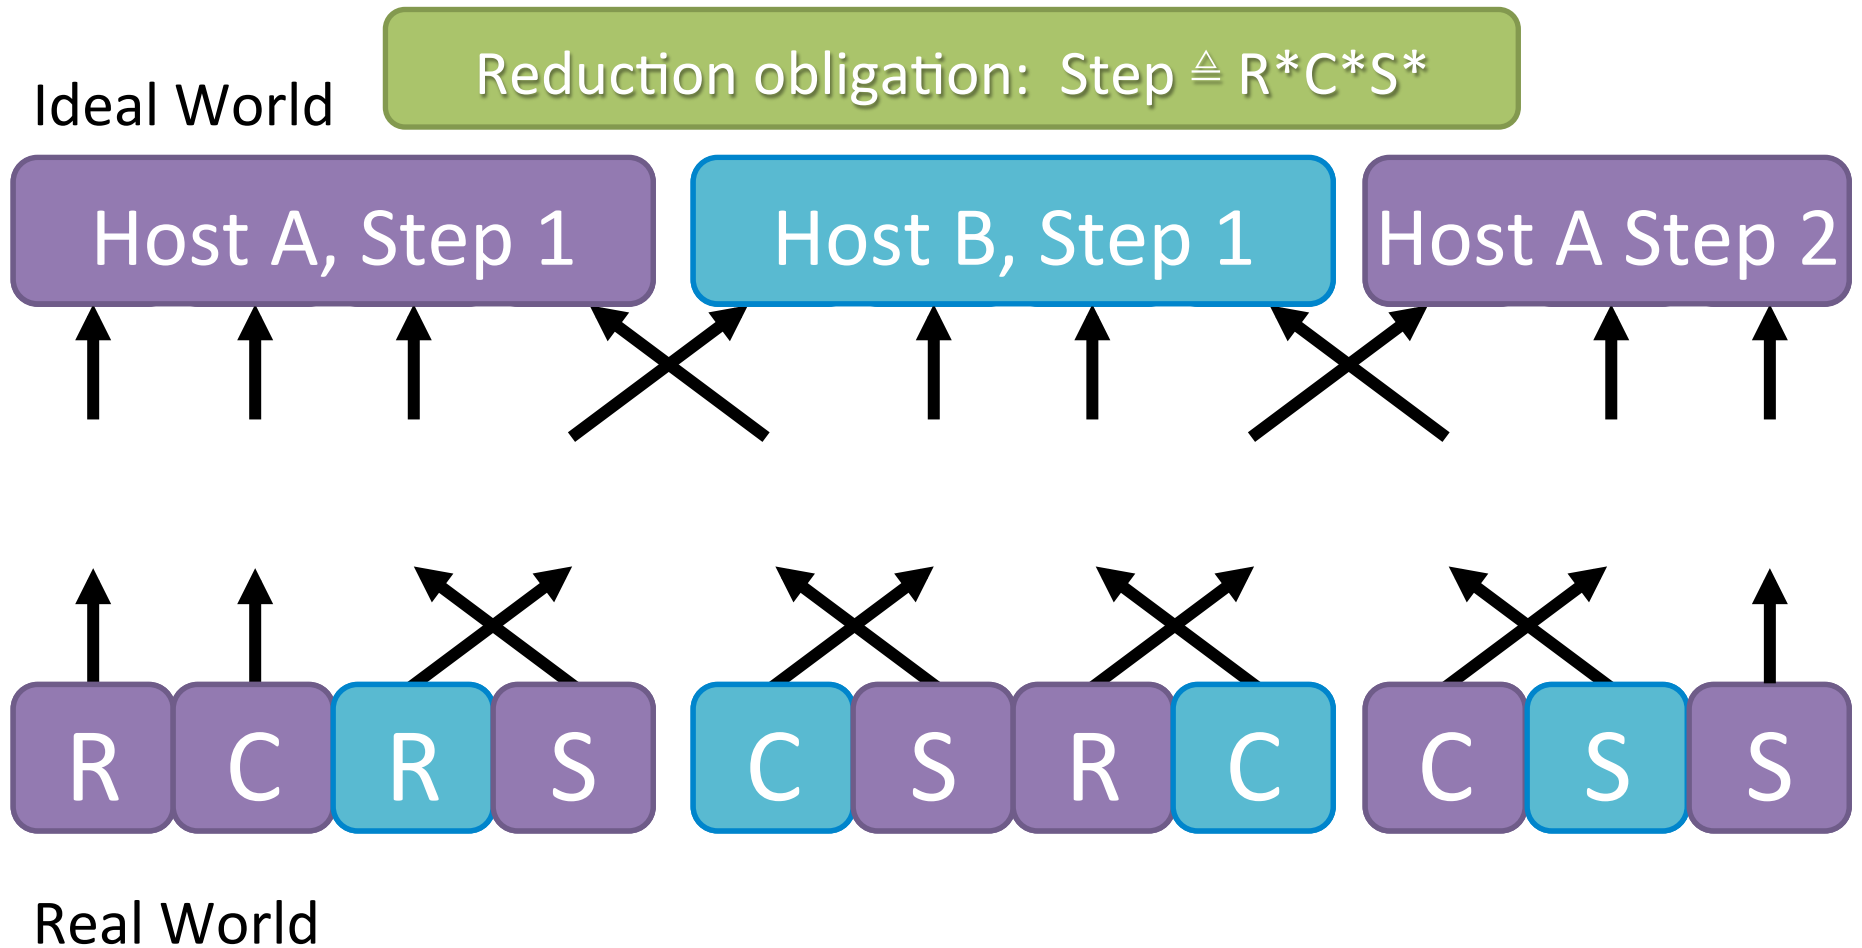
# Concurrency containment strategy

Ideal World



Real World

# Concurrency containment strategy




# Evaluation questions

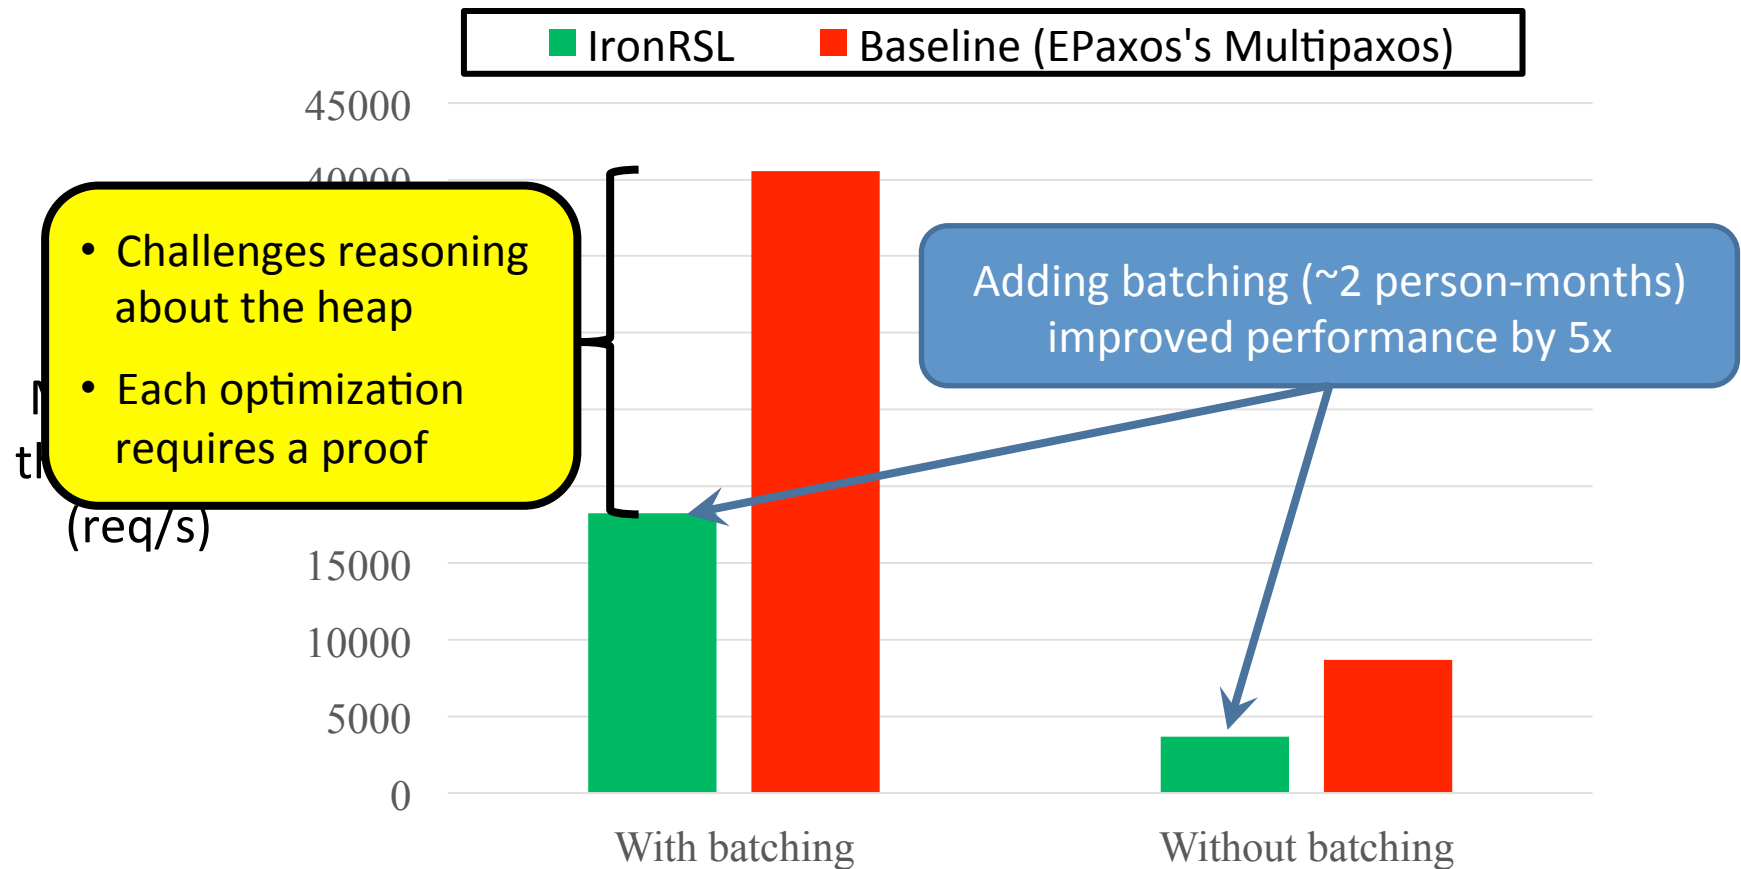
- How does verification affect development?
- How does verification affect performance?

# Developer experience: Proof burden

	<u>Functional</u>	<u>DS Safety</u>	<u>DS Live</u>
IronFleet:	3.5:1	5.4:1	7.8:1
Ironclad:	4.8:1	---	---
seL4:	25.7:1	---	---
Verdi:		94.0:1	



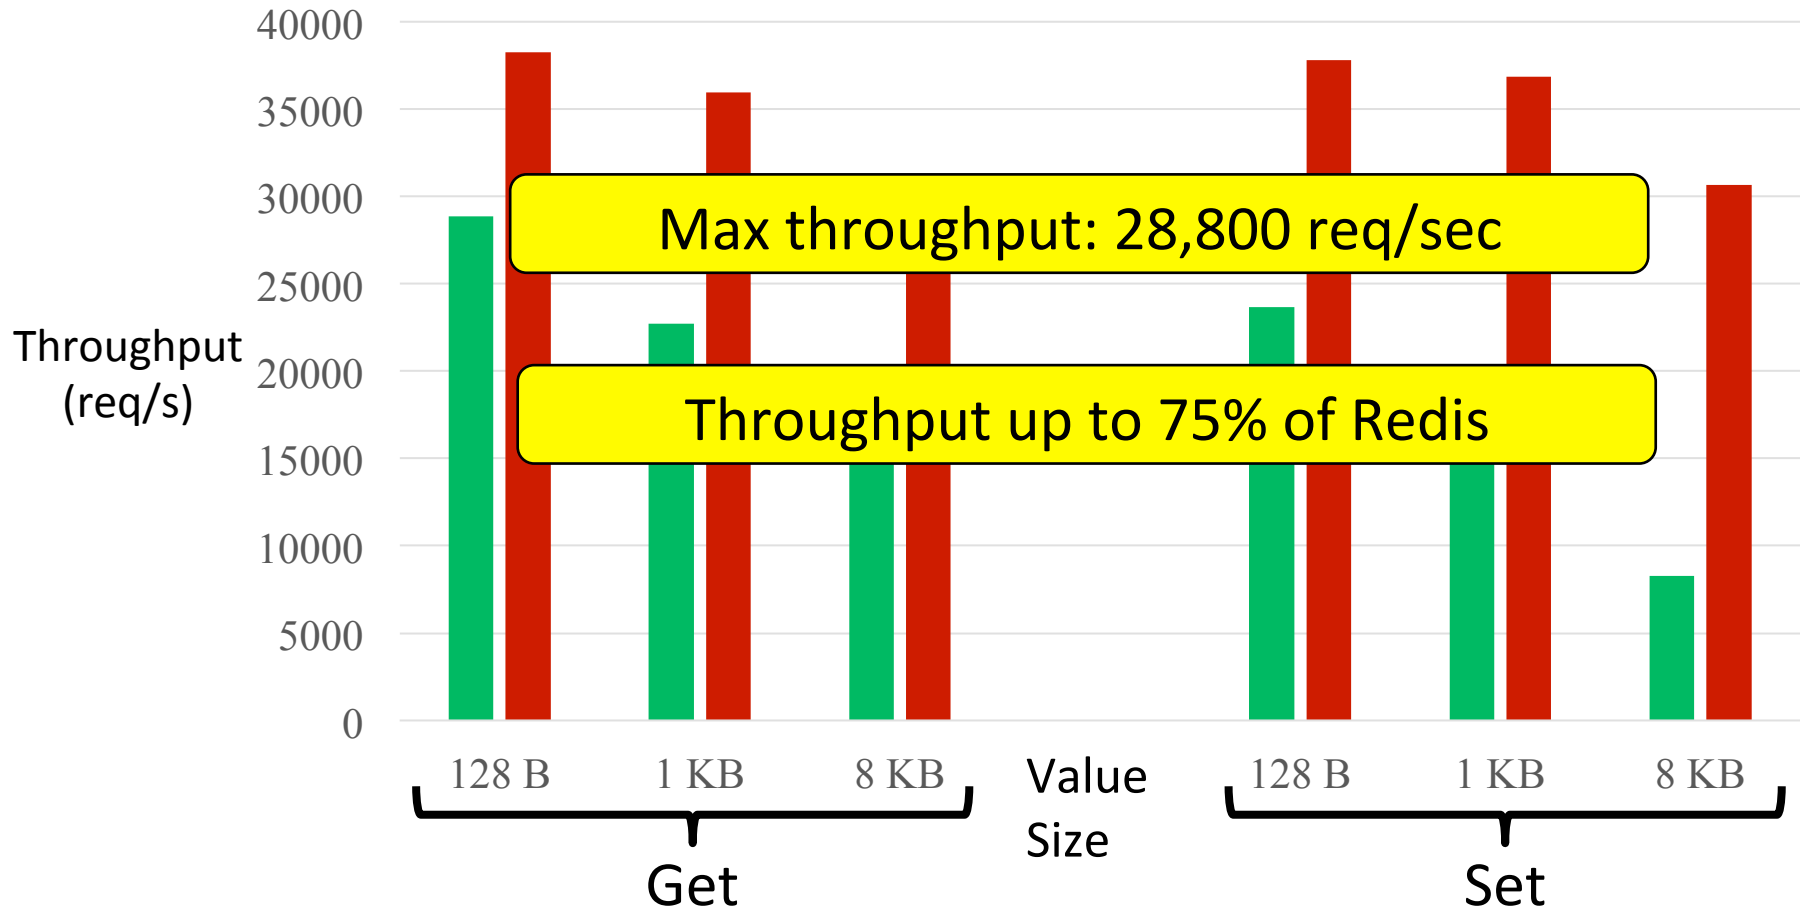
# IronRSL performance



We trade some performance for strong correctness, but no fundamental reason verified code should be slower




# IronKV performance



# The Ironclad Project


## Ironclad Apps

[OSDI 2014]




## IronFleet

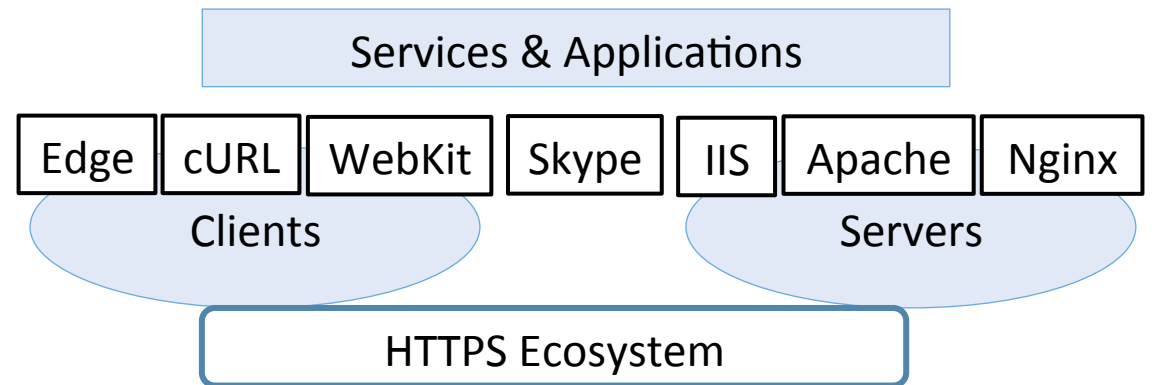
[SOSP 2015]



## Everest HTTPS

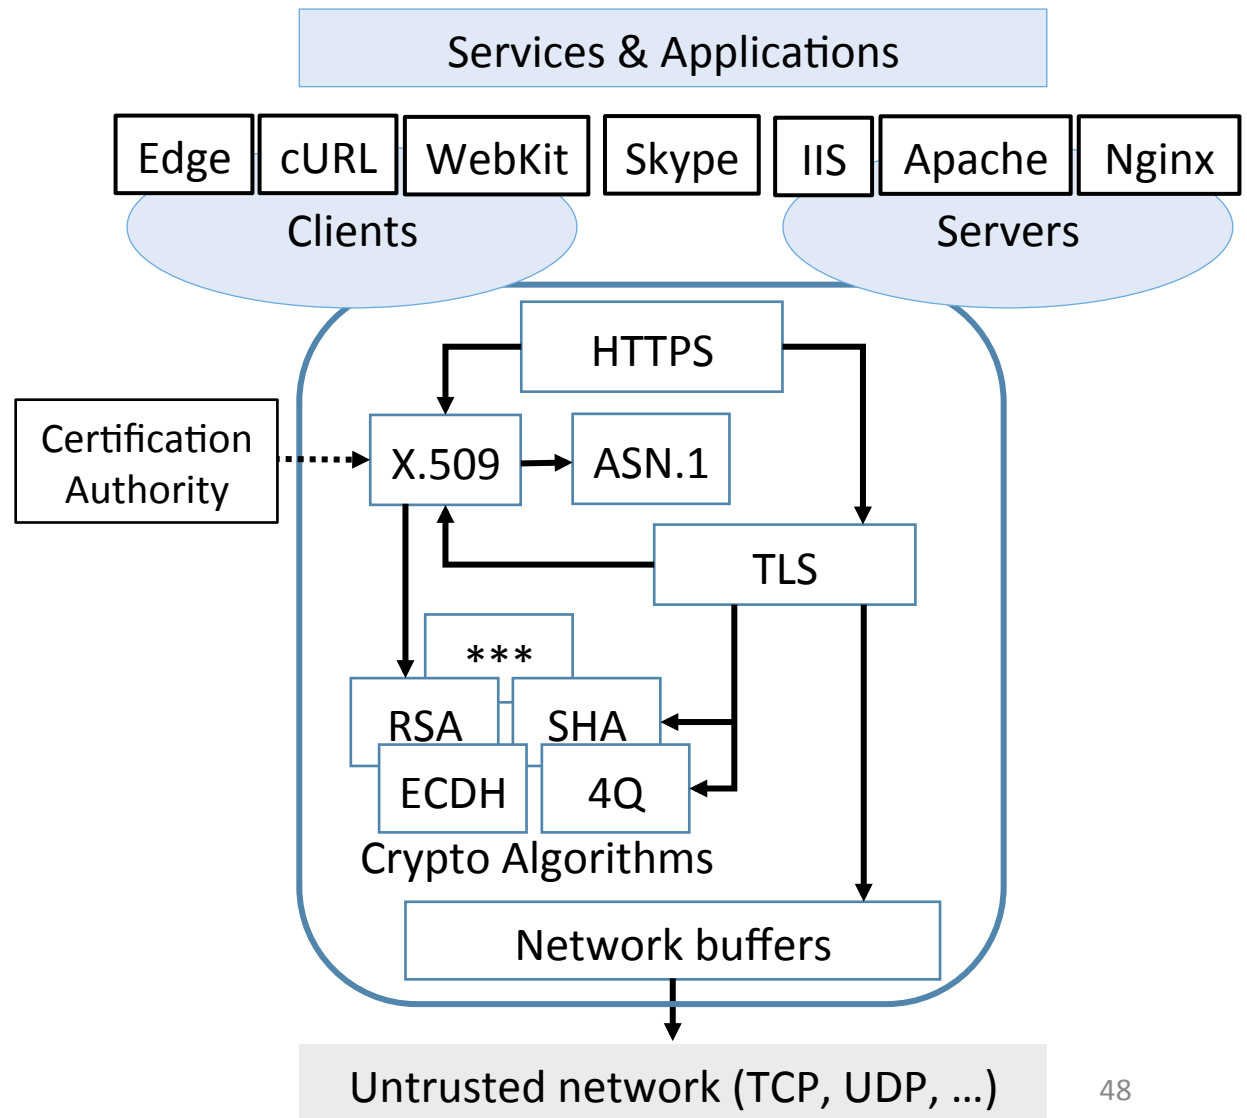


# The HTTPS Ecosystem is critical



- Most widely deployed security protocol?
  - 40% all Internet traffic (+40%/year)
- Web, cloud, email, VoIP, 802.1x, VPNs, ...

# The HTTPS Ecosystem is complex

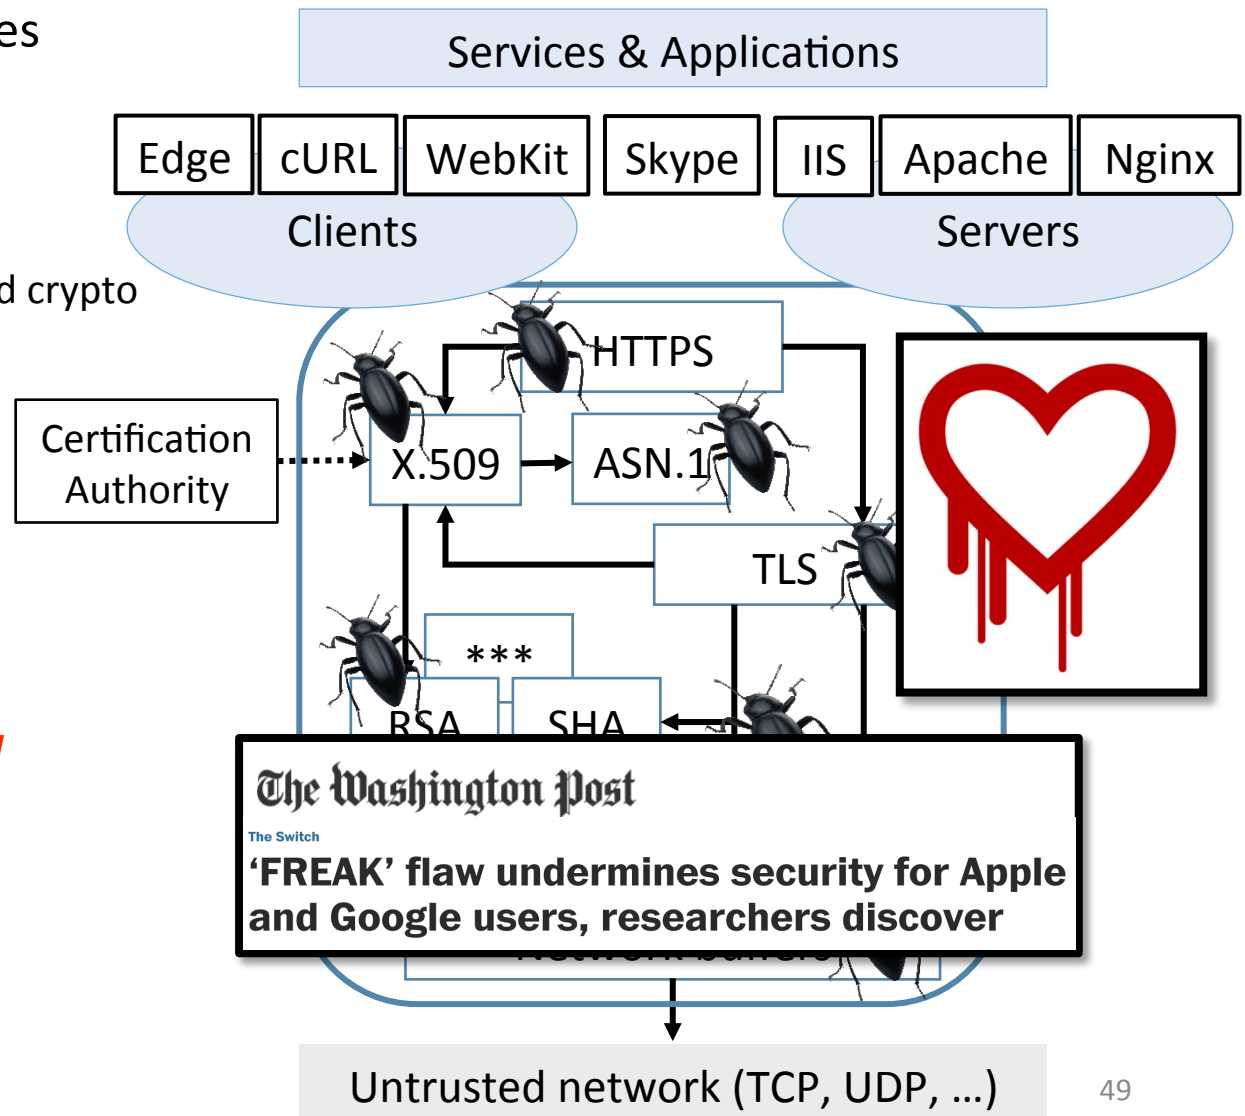


# The HTTPS Ecosystem is buggy

- 20 years of attacks & fixes
  - Buffer overflows
  - Memory management
  - Incorrect state machines
  - Lax certificate parsing
  - Weakly or badly implemented crypto
  - Side channels
  - Error-inducing APIs
  - Flawed standards
  - ...

- Many implementations
  - OpenSSL, Schannel, NSS, ...

*Still patched every month!*

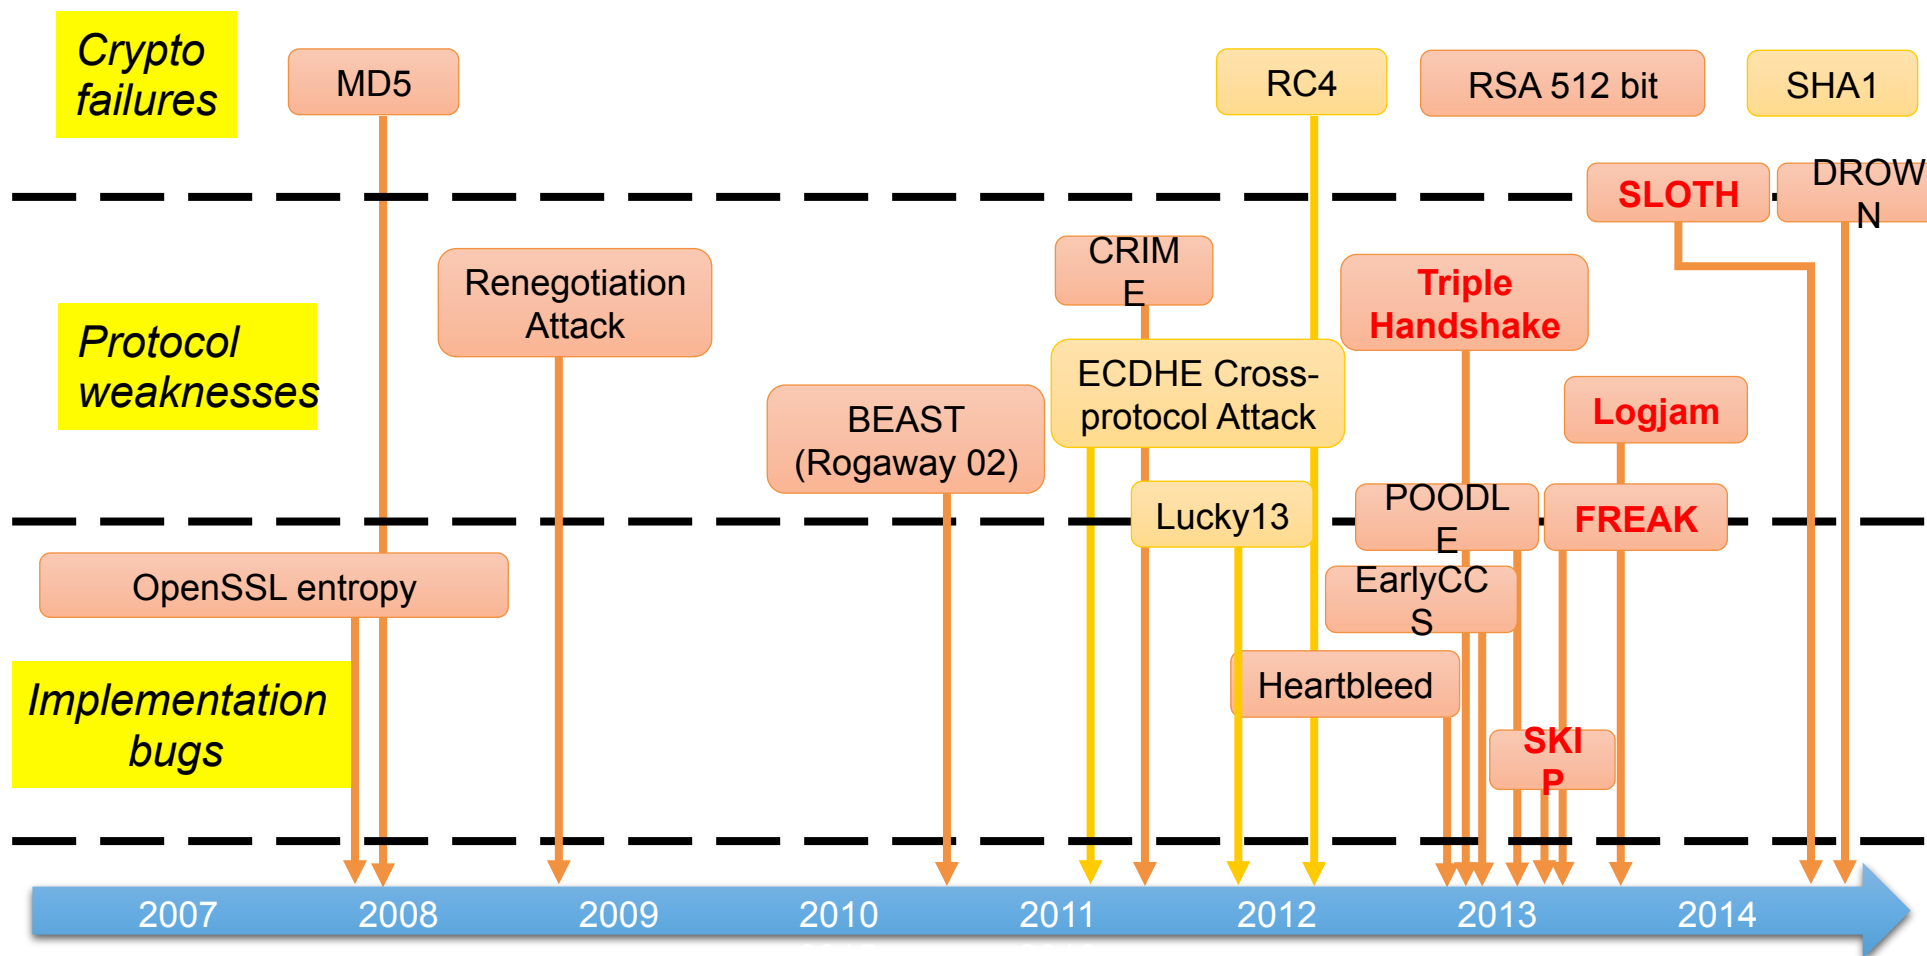


# Unsolved Attacks against HTTPS

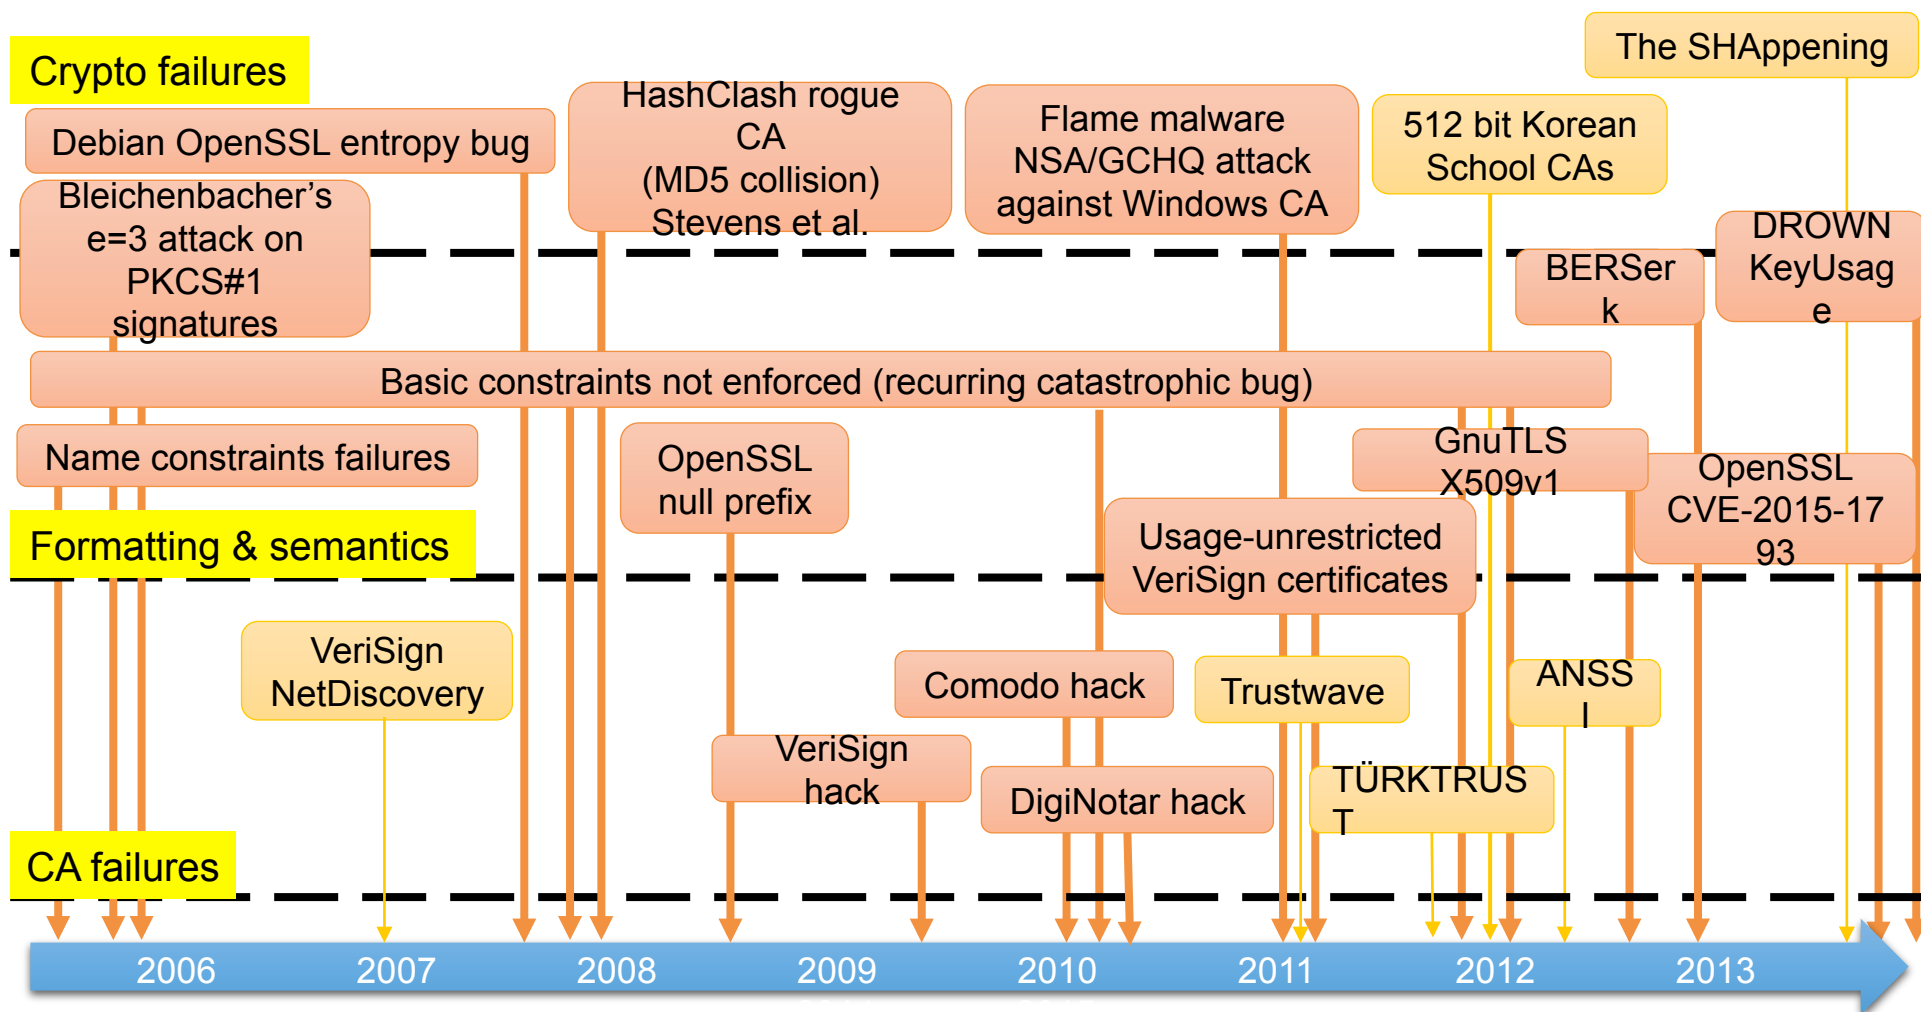
SSL Stripping (Marlinspike)	Cookie-based Attacks (various variants)	CRIME / BREACH (Rizzo, Duong et al.)	Virtual Host Confusion (Delignat-Lavaud)
TLS is optional in HTTP and can be disabled by an active attacker	Shared cookie database for HTTP and HTTPS can be used to mount various session fixation and login CSRF attacks.	Attackers can easily mount adaptive chosen-plaintext attacks. Encryption after compression can leak secrets through length.	HTTPS servers do not correlate transport-layer and HTTP identities, leading to origin confusion
Mitigated by correct use of HTTP Strict Transport Security (HSTS)	Mitigated by new binding proposals (ChannelID, Token Binding). Mitigation is not widely implemented.	Mitigated by refreshing secrets (e.g. CSRF tokens). Some protocol-specific mitigations (QUICK, HTTP2)	Mitigated by configuration of HTTPS servers with strict host rules
Mitigation is not widely used and vulnerability is still widespread in practice.	Extremely difficult to mitigate in all browsers with current technologies. Can be used to attack almost every website.	Ad-hoc mitigation; attack is still widespread in practice as HTTP compression remains popular.	Ad-hoc mitigation; attack is still widespread in practice.



# High-Profile TLS Attacks



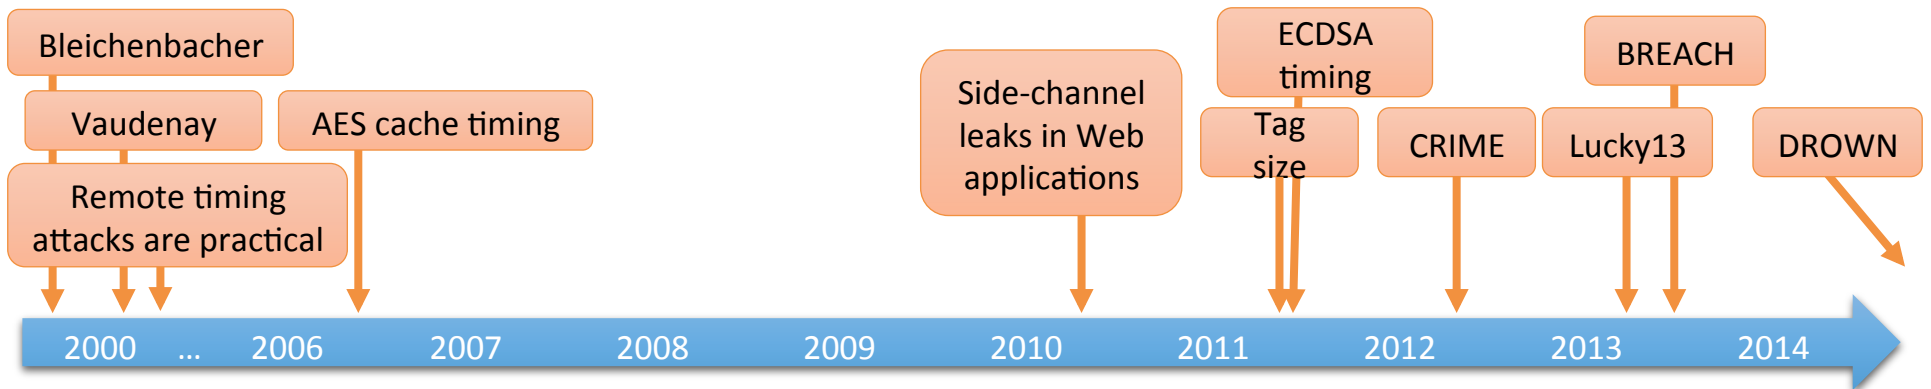
# A Timeline of Recent PKI Failures





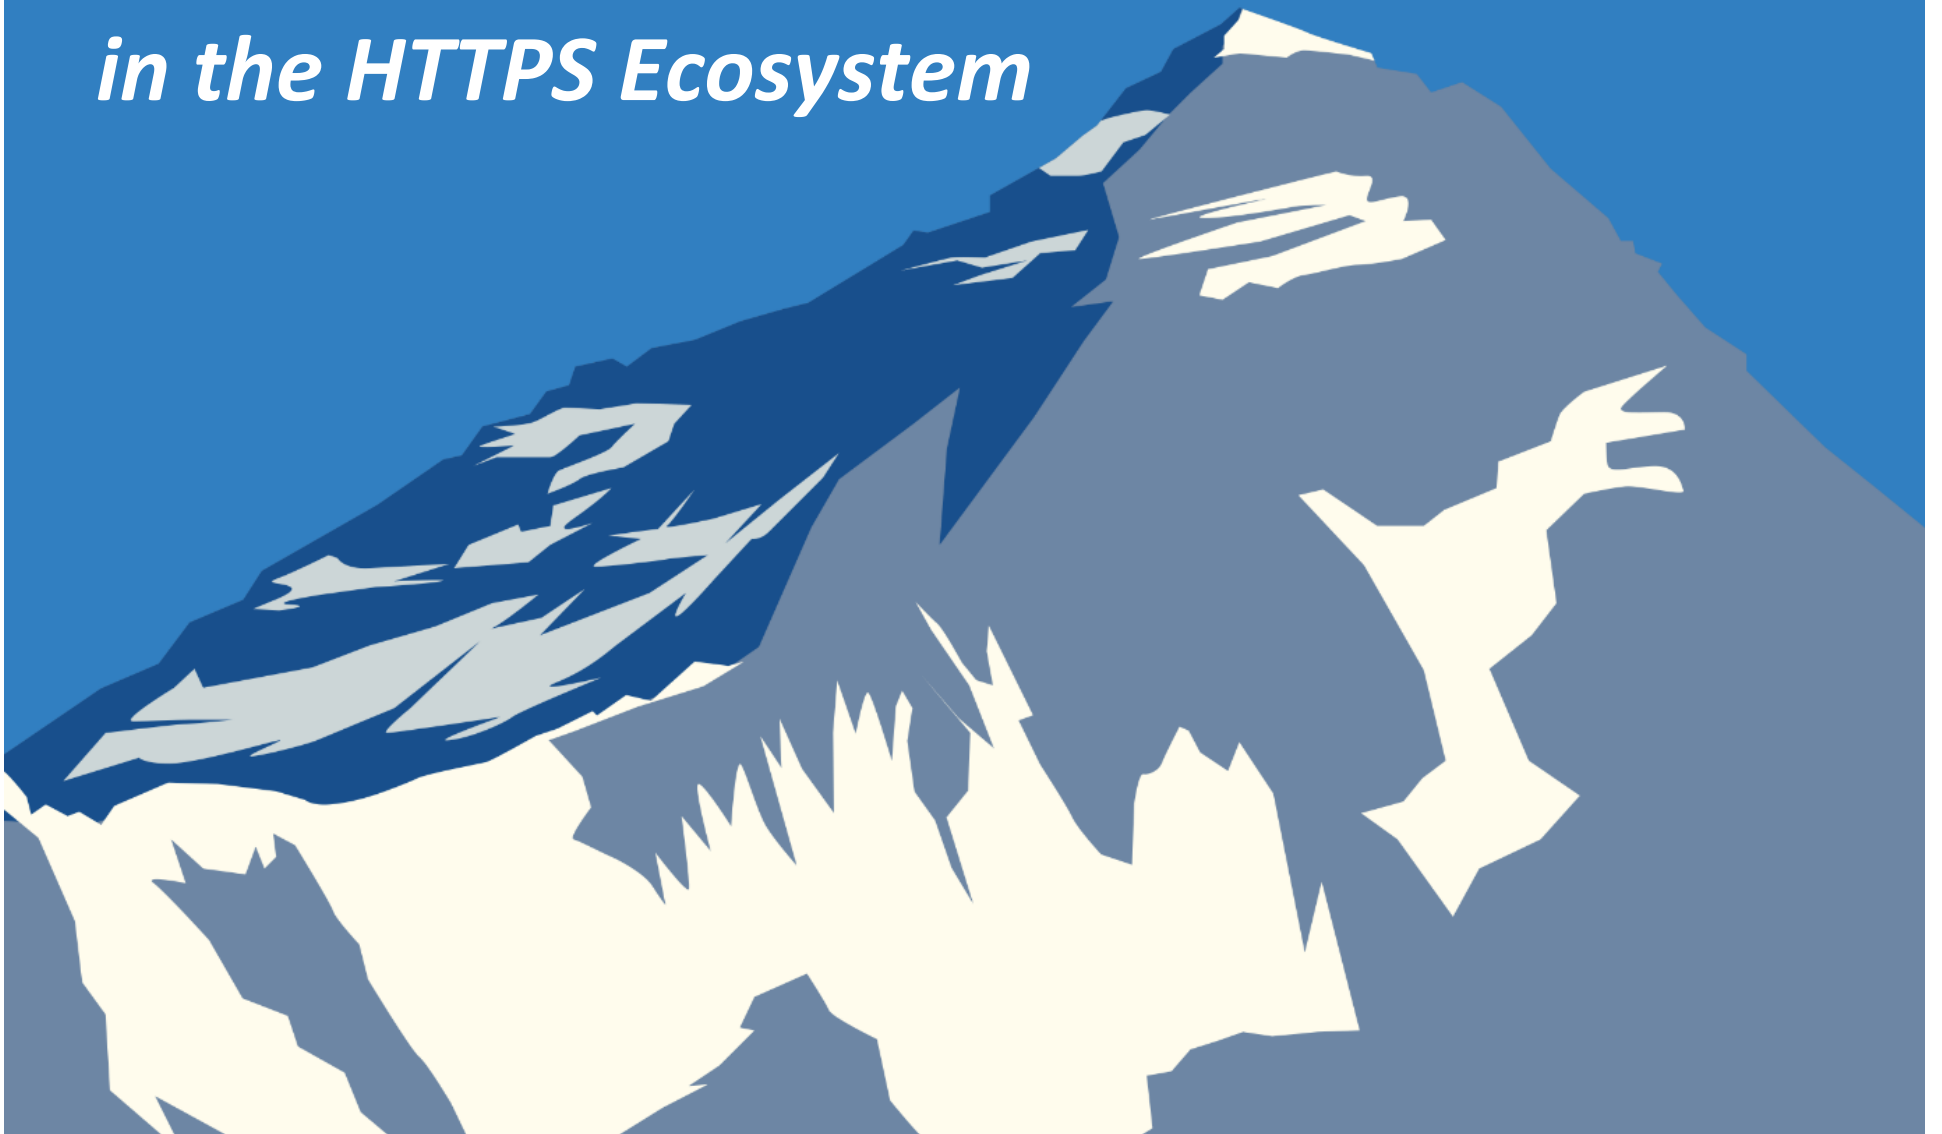
# Side Channel Challenge (Attacks)

Protocol-level side channels	Traffic analysis	Timing attacks against cryptographic primitives	Memory & Cache
TLS messages may reveal information about the internal protocol state or the application data	Combined analysis of the time and length distributions of packets leaks information about the application	A remote attacker may learn information about crypto secrets by timing execution time for various inputs	Memory access patterns may expose secrets, in particular because caching may expose sensitive data (e.g. by timing)
<ul style="list-style-type: none"> <li>• Hello message contents (e.g. time in nonces, SNI)</li> <li>• Alerts (e.g. decryption vs. padding alerts)</li> <li>• Record headers</li> </ul>	<ul style="list-style-type: none"> <li>• CRIME/BREACH (adaptive chosen plaintext attack)</li> <li>• User tracking</li> <li>• Auto-complete input theft</li> </ul>	<ul style="list-style-type: none"> <li>• Bleichenbacher attacks against PKCS#1 decryption and signatures</li> <li>• Timing attacks against RC4 (Lucky 13)</li> </ul>	<ul style="list-style-type: none"> <li>• OpenSSL key recovery in virtual machines</li> <li>• Cache timing attacks against AES</li> </ul>



*Everest:*


*Deploying Verified-Secure Implementations  
in the HTTPS Ecosystem*



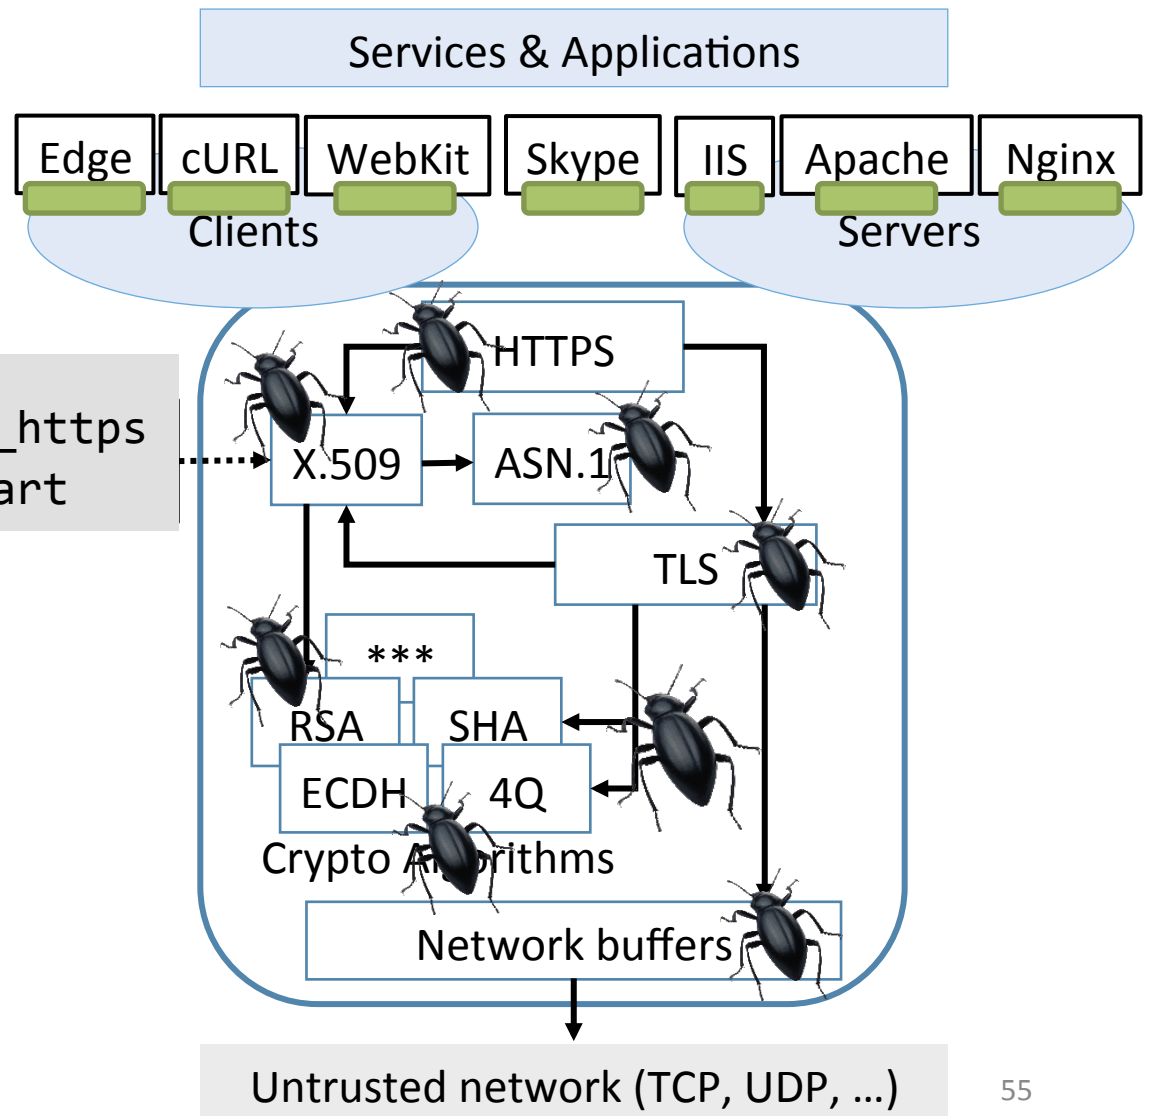
# Everest Goals

- Fully verified replacement
- Widespread deployment
- Trustworthy, usable tools

```
$ apt-get install verified_https
$ /
```



Logos for Dafny, BOOGIE, Z3, and EM THEOREM PROVER are displayed within a green-bordered box.



# Research Questions

- How do we decide whether new protocols are secure?
  - Especially when interoperating with insecure protocols
- Can we make verified systems as fast as unverified?
- How do we handle advanced threats?
  - Ex: Side channels
- Why should we trust automated verification tools?
- How can verification be more accessible?
  - Especially to non-experts in verification

# Summary

- Ironclad Apps guarantee end-to-end security to remote parties: Every instruction meets the app's security spec
- IronFleet extends these techniques to prove the safety and liveness of distributed systems
- Everest will showcase the power of verification and its applicability to real-world security problems
- Verification of systems code is possible, and we're scaling it to even larger more complex systems

<https://github.com/Microsoft/Ironclad>

Thank you!  
parno@cmu.edu

