

Compiler-instrumented, Dynamic Secret-Redaction of Legacy Processes for Attacker Deception

Frederico Araujo and Kevin Hamlen
The University of Texas at Dallas

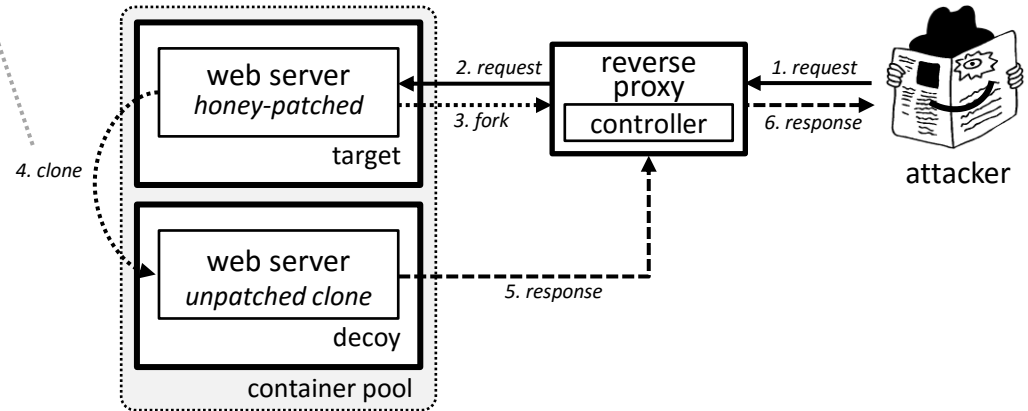
- Goal:
 - Remove or replace secrets in address spaces of *running programs*, yielding processes that **CONTINUE RUNNING** (but with no secrets)
- Potential Applications:
 - Debugging: Safely disclose redacted crash dumps to developers
 - Intrusion Response: Dynamic secret redaction without loss of service
 - Cyber Deception: Runtime replacement of secrets with honey-data
 - *Honey-patching [CCS'14]*

- Goal:
 - Remove or replace secrets in address spaces of *running programs*, yielding processes that **CONTINUE RUNNING** (but with no secrets)
- Potential Applications:
 - Debugging: Safely disclose redacted crash dumps to developers
 - Intrusion Response: Dynamic secret redaction without loss of service
 - **Cyber Deception: Runtime replacement of secrets with honey-data**
 - *Honey-patching [CCS'14]*

```
.....L.....
.@.. GET / HTTP/1.1 /browse/doc1.html
en_US xyz-198 8229788/6160/11/.....
.Accept-Encoding: gzip,deflate,sdch...
...Accept-Language: en-US,en;q=0.8...
Cookie: app.token= BACC-76GF-ABS3-ZOV2
74f89abc43de7.....*.....
SESSIONID=2321CFA5DA771A284D13DD67798A
557.....E.$3Z.l8.M..e5.....
...7ED1D554E.....*..?.e.b...L...
.....*.....App.token= BACC-65CH-
Accept:text/html,application/xhtml+xml
,application/xml;q=0.9,*/*;q=0.8.....
.....*.....
Linux x86_64; rv:32.0) Gecko/20100101
Firefox/32.0.....@.....
.....2d-4f59f9ff30097.....*.....
.GET / HTTP/1.1.....
```



- user 1
- user 2
- attacker



Main idea: *Instrument programs with operations that track (explicit) dataflows of secrets.*

- Program vulnerability detection
 - TaintCheck, LIFT, Mimemu, Argos, ...
- Information leak detection
 - TaintDroid, TaintEraser, AndroidLeaks, Spandex, D2Taint, ...
- Study of sensitive data lifetime
 - TaintBochs
- Analysis of spyware behavior
 - Panorama, Hookfinder, PHP Aspis, ...
- Test set generation
 - Memsherlock, ConfAid, ...

```
/* first colon delimits username:password */
s1 = strchr(hostinfo, ':', s - hostinfo);
if (s1) {
    uptr->user = memdup(hostinfo, s1 - hostinfo);
    ++s1;
    uptr->password = memdup(s1, s - s1);
}
```

```
void safe_free(char *s) {
    if (s is a secret)           // how to test whether s is secret?
        slow_secure_free(s);
    else
        free(s);
}
```

Taint Introduction:

```
/* first colon delimits username:password */
s1 = strchr(hostinfo, ':', s - hostinfo);
if (s1) {
    uptr->user = memdup(hostinfo, s1 - hostinfo);
    ++s1;
    uptr->password = memdup(s1, s - s1);
    dfsan_set_label(SECRET, uptr->password, sizeof(s - s1))
}
```

```
void safe_free(char *s) {
    if (s is a secret)           // how to test whether s is secret?
        slow_secure_free(s);
    else
        free(s);
}
```

Taint Introduction:

```
/* first colon delimits username:password */
s1 = strchr(hostinfo, ':', s - hostinfo);
if (s1) {
    uptr->user = memdup(hostinfo, s1 - hostinfo);
    ++s1;
    uptr->password = memdup(s1, s - s1);
    dfsan_set_label(SECRET, uptr->password, sizeof(s - s1))
}
```

Taint Check:

```
void safe_free(char *s) {
    if (dfsan_get_label(s) == SECRET)
        slow_secure_free(s);
    else
        free(s);
}
```


Taint Introduction:

```
...
/* first colon delimits username:password */
s1 = strchr(hostinfo, ':', s - hostinfo);
if (s1) {
    uptr->user = memdup(hostinfo, s1 - hostinfo);
    ++s1;
    uptr->password = memdup(s1, s - s1);
    dfsan_set_label(SECRET, uptr->password, sizeof(s - s1))
}
...
```

Taint Introduction:

```
p = &(uptr->password);
*p = malloc(...);
...
/* first colon delimits username:password */
s1 = strchr(hostinfo, ':', s - hostinfo);
if (s1) {
    uptr->user = memdup(hostinfo, s1 - hostinfo);
    ++s1;
    *p = memdup(s1, s - s1);
    dfsan_set_label(SECRET, *p, sizeof(s - s1))
}
...
```



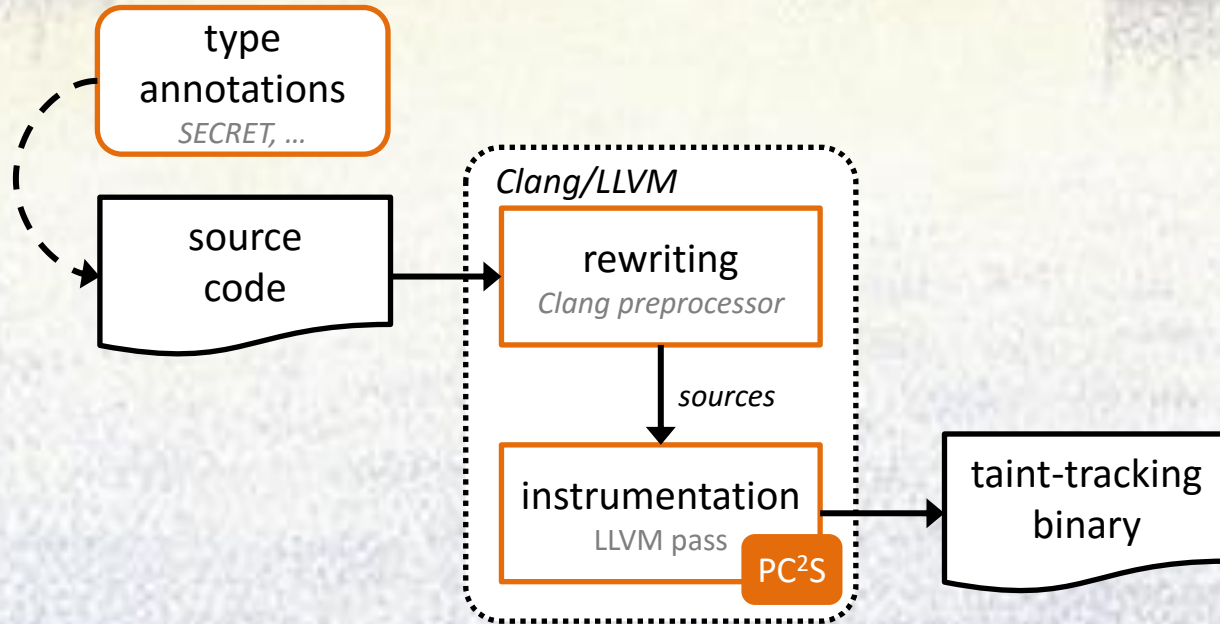
*Apache is
2.2M SLOC!*



*type
qualifiers!*

```
struct apr_uri_t {  
    NONSECRET char *user;  
    SECRET_STR char *password;  
    ...  
} SECRET;
```

- *Declarative vs. Operational* Secret Annotations
 - Fewer declarations than operations for user to annotate
 - Compiler infers and implements operations from declarations
 - Compiler optimizes operational implementation
- **SECRET** = struct contains secrets
- **SECRET_STR** = field is a pointer to a null-terminated sequence of secret chars
- **NONSECRET** = field is a non-secret within a SECRET struct



1. Overview

2. Challenges & Approach

3. Implementation

4. Evaluation

5. Application Study

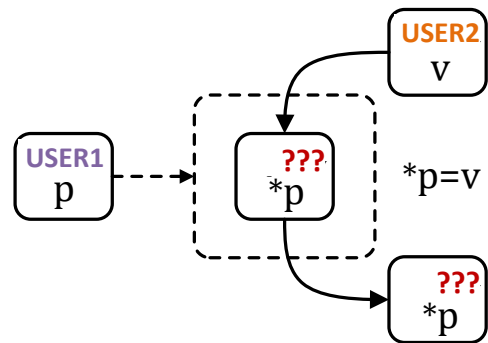
6. Conclusions

```
p = &(uptr->password);
*p = malloc(...);
...
/* first colon delimits username:password */
s1 = strchr(hostinfo, ':', s - hostinfo);
if (s1) {
    uptr->user = memdup(hostinfo, s1 - hostinfo);
    ++s1;
    memcpy(*p, s1, s - s1);
}
...
```

Need taint propagation semantics for...

- field access operator (->)
- address-of (&) operator
- assignments (=)
- dereferencing assignments (*p = ...)
- dynamic memory allocations (malloc)


```
mytype *p = ...;
mytype v = ...;
dfsan_set_label(USER1, p, sizeof(p));
dfsan_set_label(USER2, v, sizeof(v));
*p = v;
```




What should be the resulting label of ***p** ?

- Two standard answers:
 - *No-Combine Semantics*: label of ***p** is just **USER2**
 - Rationale: **v** was copied; its ownership didn't change.
 - *Combine Semantics*: label of ***p** is **USER1** \sqcup **USER2** (joint ownership)
 - Rationale: Failing to redact value at ***p** now possibly divulges value of pointer **p**.
 - Conclusion: Value ***p** is now one of **USER1**'s secrets (as well as continuing to be one of **USER2**'s secrets).

propagation semantics:



$p = \&(uptr \rightarrow password);$
 $*p = \text{malloc}(\dots);$
 $*p[i] = v \text{ (for all } i)$



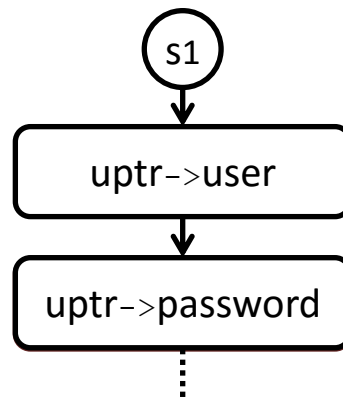
$x = *p[i] \text{ (for all } i)$

revisiting our initial example...

```

struct apr_uri_t {
  NONSECRET char *user;
  SECRET_STR char *password;
  ...
} SECRET;
    
```

$uptr \rightarrow user = v1;$
 $uptr \rightarrow password = v2;$



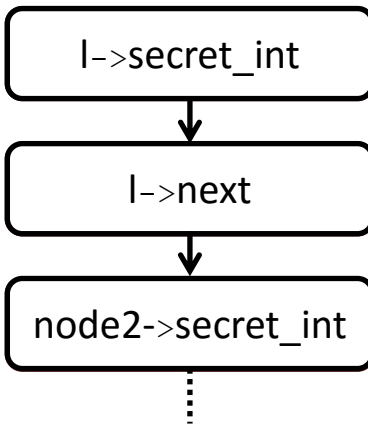
but...

```
l->secret_int = 1234;  
l->next = node2;  
node2->secret_int;
```

l->secret_int = 1234

l->next = node2

node2->secret_int




very common...

```
while (freelist != NULL) {  
    node = freelist;  
    freelist = node->next;  
    free(node);  
}
```

```
while (prev) {  
    prev->eos_sent = 1;  
    prev = prev->prev;  
}
```

...

```
mytype *p = ...;
mytype v = ...;
dfsan_set_label(USER1, p, sizeof(p));
dfsan_set_label(USER2, v, sizeof(v));
*p = v;
```


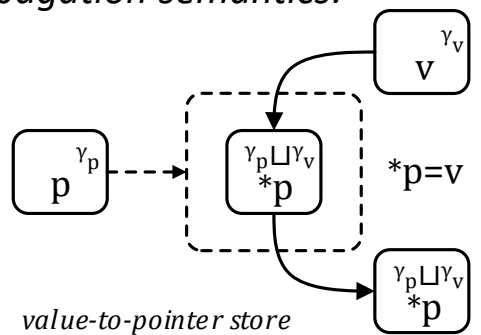


What should be the resulting label of `*p` ?

A New Third Answer:

- *Conditional-Combine Semantics*: label of `*p` depends upon the *static type* of `v`!
 - if `v` has *pointer* type, then use *No-Combine Semantics* (`USER2`).
 - if `v` has *non-pointer* type, then use *Combine Semantics* (`USER1` \sqcup `USER2`).

propagation semantics:

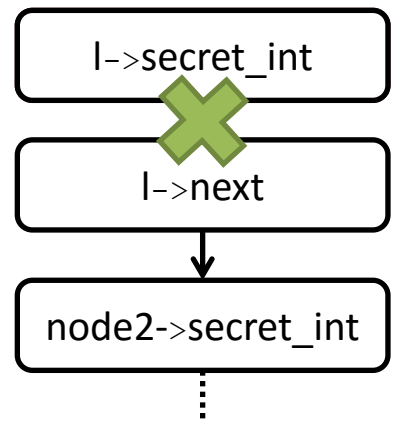


taint policy: “do not combine when pointee has pointer type”

let's try again...

```
l->secret_int = 1234;
l->next = node2;
node2->secret_int;
```

```
l->secret_int = 1234
l->next = node2
node2->secret_int
```



1. Overview
2. Challenges & Approach
3. Implementation
4. Evaluation
5. Application Study
6. Conclusions

Annotated Types

```
struct request_rec {  
    NONSECRET ... *pool;  
    apr_uri_t parsed_uri;  
    ...  
} SECRET;
```

Rewriting

clang transformation

```
new = (request_rec *) apr_palloc(r->pool, ...);
```

```
new = (request_rec *) signac_alloc(apr_palloc, r->pool, ...);
```

Instrumentation

clang/LLVM
-dfsan-pc²s

instrumented
binary

libsignaC

```
#define SECRET __attribute__((annotate("secret")))
```

clang -Xclang -ast-dump

```
| -RecordDecl 0x8943a40 <line:15:9, line:19:1> line:15:16 struct request_rec definition  
| | -AnnotateAttr 0x8943d60 <line:3:31, col:48> "secret"  
| | -FieldDecl 0x8943c20 <line:16:5, col:29> col:29 referenced pool 'apr_pool_t * ...  
| ` -FieldDecl 0x8943ca0 <line:17:5, col:15> col:15 parsed_uri 'apr_uri_t': 'struct apr_uri_t'
```

Annotated Types

```
struct request_rec {
  NONSECRET ... *pool;
  apr_uri_t parsed_uri;
  ...
} SECRET;
```

Rewriting

clang transformation

```
new = (request_rec *) apr_palloc(r->pool, ...);
```

```
new = (request_rec *) signac_alloc(apr_palloc, r->pool, ...);
```

Instrumentation

clang/LLVM
-dfsan -pc²s

instrumented
binary

libsignaC

```
#define NONSECRET __attribute__((type_annotate("nonsecret")))
#define SECRET_STR __attribute__((type_annotate("secret_str"))) } Quala type qualifiers
```

```
rec->pool = pool;
```



clang -S -emit-llvm

```
%3 = load %struct.apr_pool_t** %pool, align 8
```

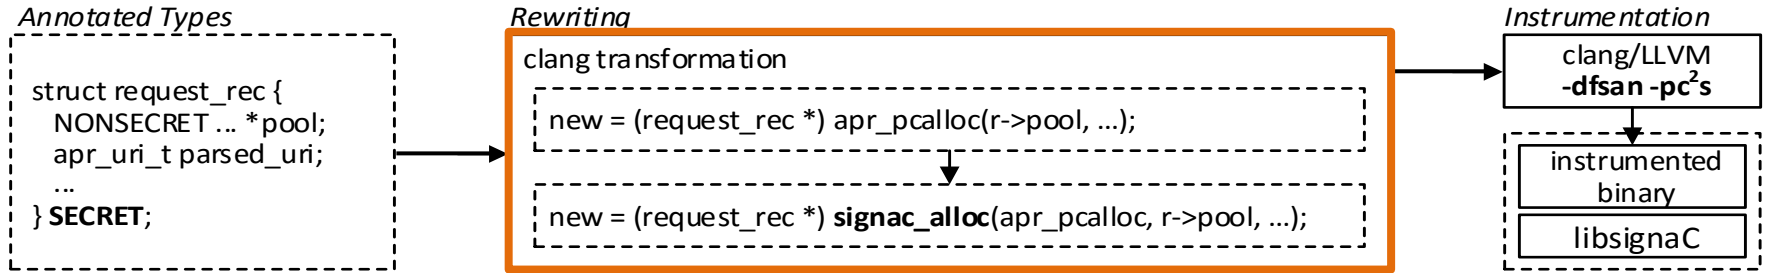
```
%4 = load %struct.request_rec** %rec, align 8
```

```
%pool2 = getelementptr inbounds %struct.request_rec* %4, i32 0, i32 0
```

```
store %struct.apr_pool_t* %3, %struct.apr_pool_t** %pool2, align 8, !tyann !1
```

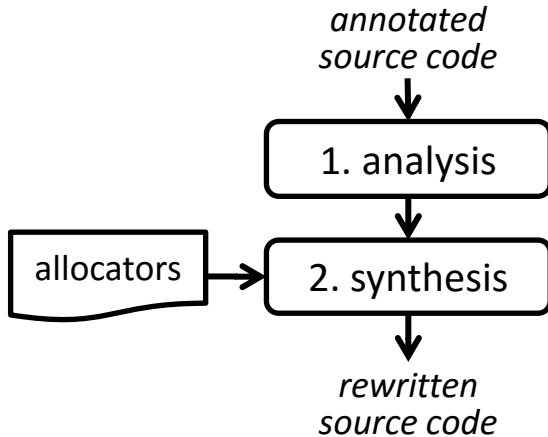
```
...
```

```
!1 = !{"nonsecret", i8 0}
```

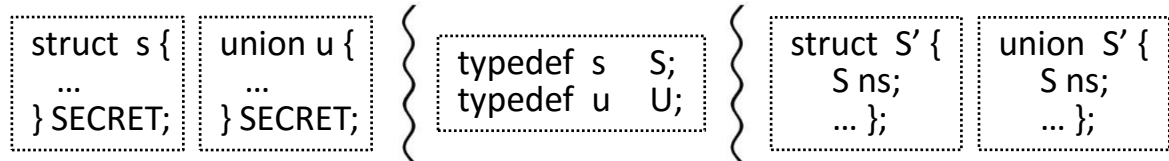


→ Clang tooling API: AST Matchers + Rewriting API

→ Allocators list: {malloc, calloc, apr_palloc, apr_pcalloc, ...}




1. Compute the list of all **security-relevant datatypes**



2. Wrap all security-relevant datatypes instantiations with **signac_alloc**


```
#define signac_alloc(alloc, args...) ({ \  
  void * __p = alloc ( args ); \  
  signac_taint(&__p, sizeof(void*)); \  
  __p; })
```

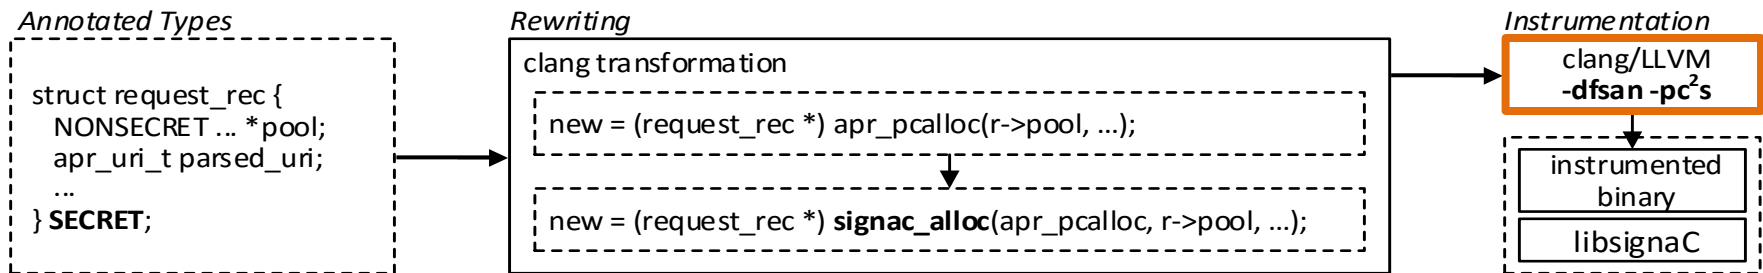


- implemented as an extension to **DFSan**
- low-overhead representation of labels: **16-bit integers** allocated sequentially
- maps without reserving the lower 32TB of the process address space for **shadow memory**
- *union labels* organized as a dynamically growing binary (DAG) – the **union table**

Start	End	Memory Region
0x700000008000	0x800000000000	application memory
0x200000000000	0x700000008000	union table
0x000000010000	0x200000000000	shadow memory
0x000000000000	0x000000010000	reserved by kernel

example:
Z = X + Y





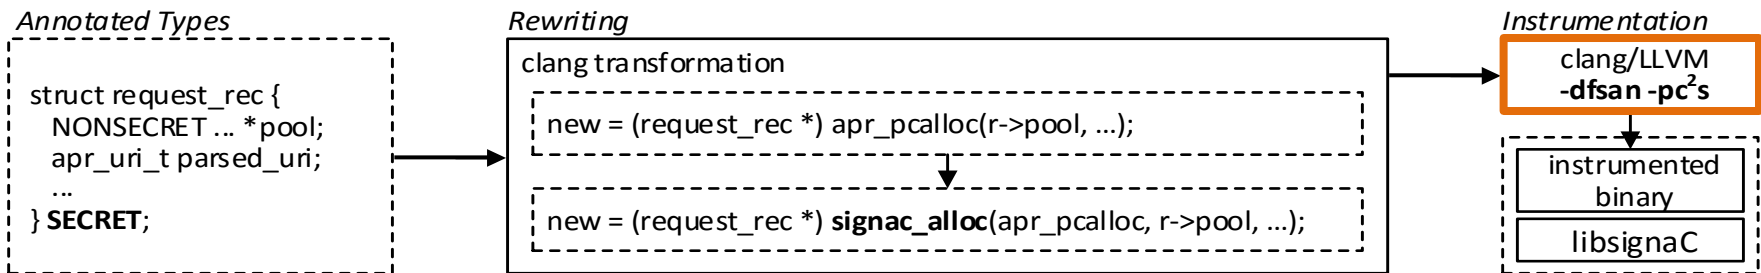
- label propagation across external library interfaces expressed as an **ABI list**
- DFSan predefines an ABI list that covers glibc

```

fun:malloc=custom
fun:realloc=discard
fun:free=discard
...
fun:isalpha=functional
fun:isdigit=functional
...
fun:memcpy=custom
fun:memset=custom
fun:strcpy=custom
  
```

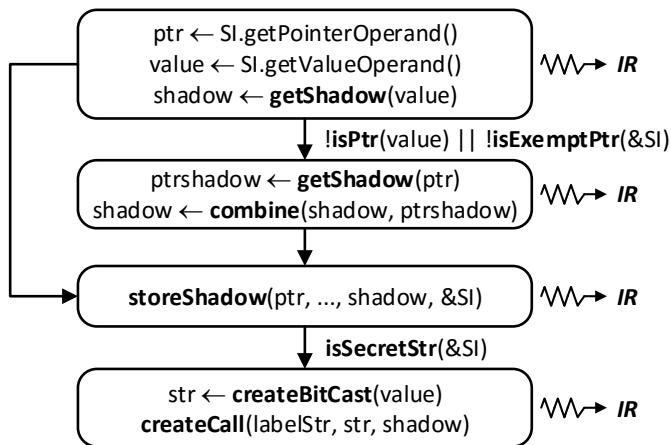
discard $\rho_{dis}(\bar{\gamma}) := \perp$
functional $\rho_{fun}(\bar{\gamma}) := \sqcup \bar{\gamma}$
custom *custom-defined label propagation wrapper*

- other libraries mapped to the ABI: OpenSSL, PCRE, APR, ...
- *memory transfer functions* (e.g., *strcpy*, *strdup*) and *input functions* (e.g., *read*, *pread*) ABI extensions for PC²S



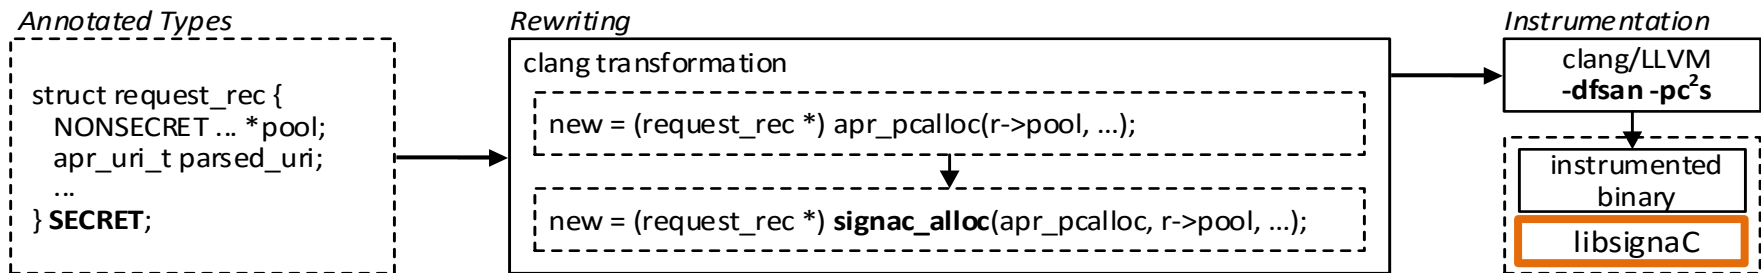
- instrumentation operates on **LLVM IR**, inserting label propagation code
- propagation policy parametrized at the compiler's front-end: **pc2s-on-store**, **pc2s-on-load**

*example:
store instruction*



*check for **NON_SECRET**
annotated pointers*

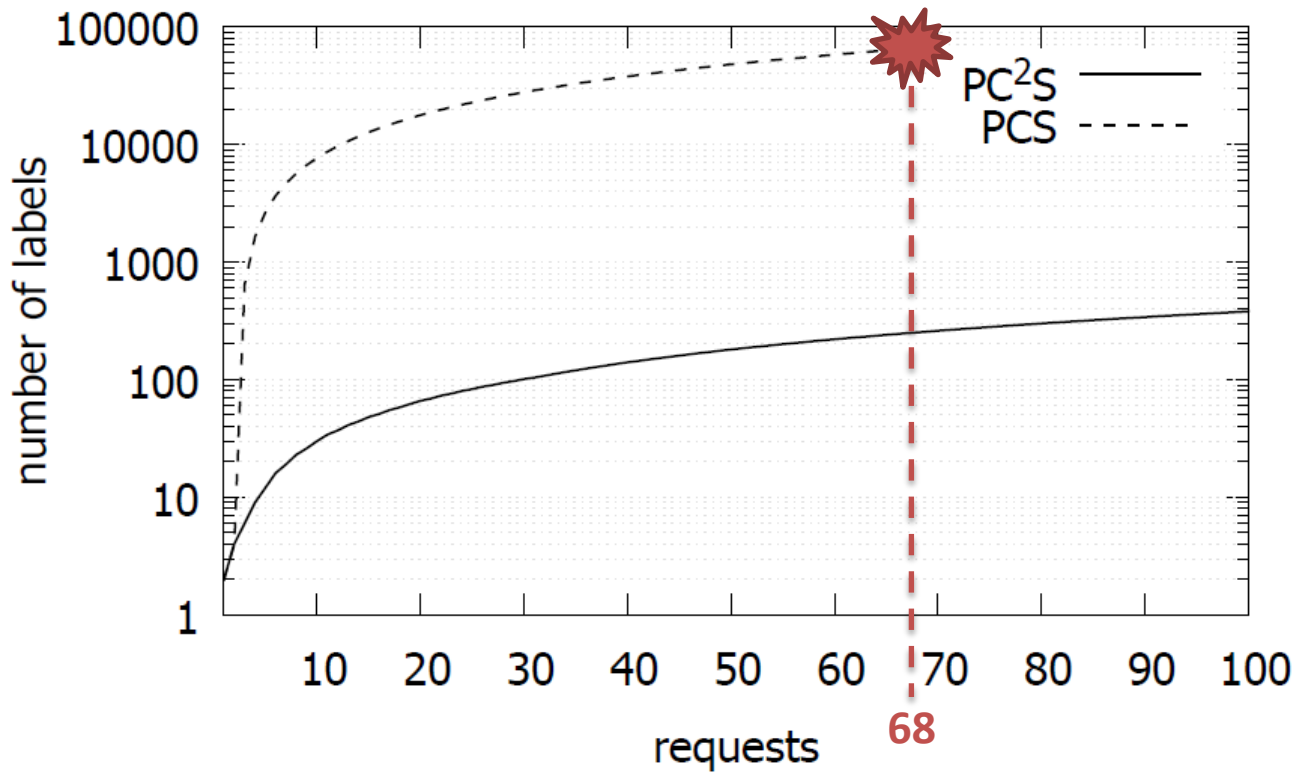
*special handling of
SECRET_STR annotated types*



→ libsignaC : tiny C library that encapsulates runtime support for the type annotation mechanism

- signac_init(*pl*)** initialize a tainting context with a fresh label instantiation *pl* for the current principal.
- signac_taint(*addr, size*)** taint each address in interval [*addr; addr+size*) with *pl*.
- signac_alloc(*alloc, ...*)** wrap allocator *alloc* and taint the address of its returned pointer with *pl*.

1. Overview
2. Challenges & Approach
3. Implementation
4. Evaluation
5. Application Study
6. Conclusions




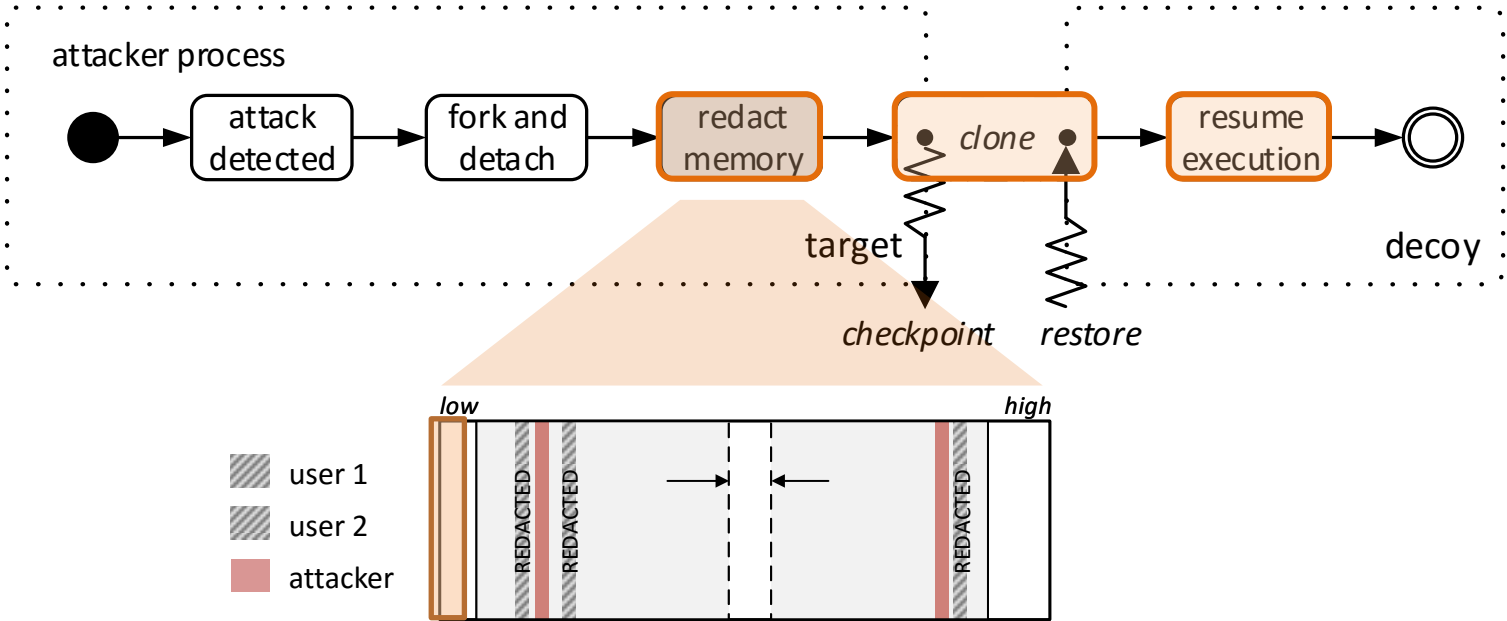



Table 2: Average overhead of instrumentation

Benchmark	C=1	C=10	C=50	C=100
Static	2.50	2.34	2.56	2.32
CGI Bash	1.29	0.98	1.00	0.97
PHP	0.41	0.37	0.30	0.31

1. Overview
2. Challenges & Approach
3. Implementation
4. Evaluation
5. Application Study
6. Conclusions





Deception Strategy: Artificially delay non-forking responses to match the forking delay.

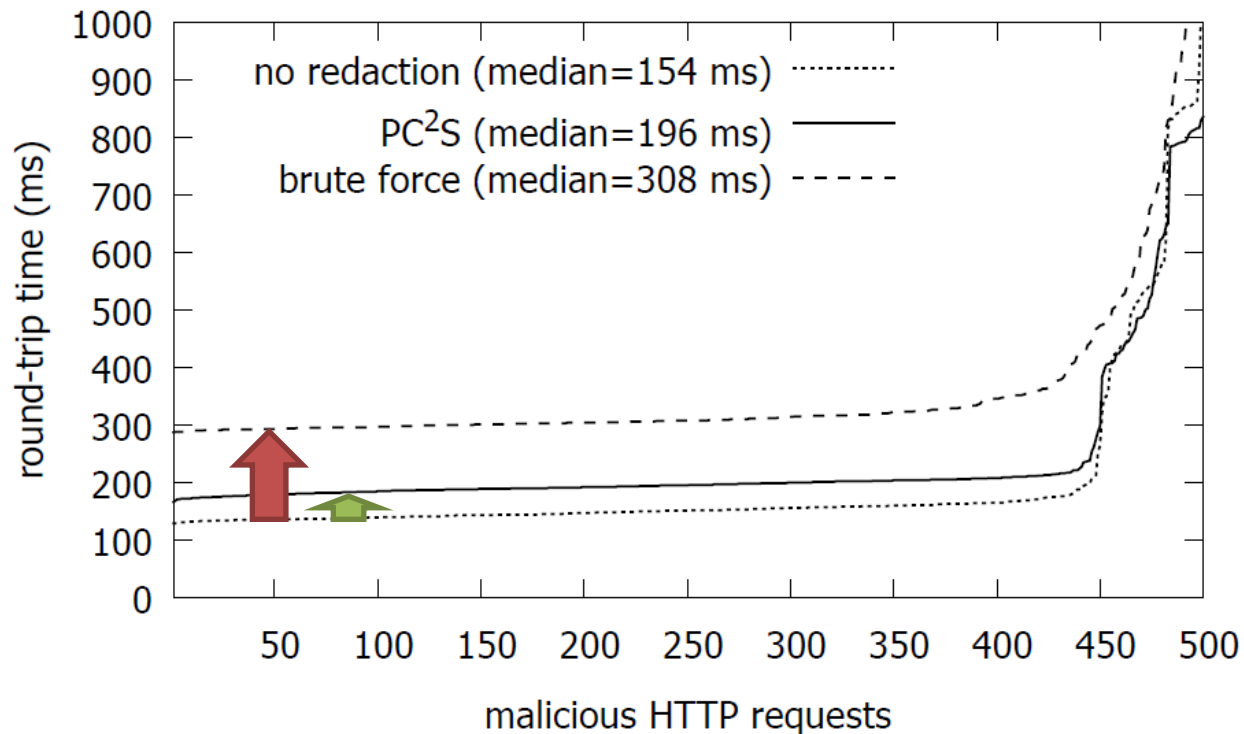


Table 1: Honey-patched security vulnerabilities

Software	Version	CVE-ID	Description
Apache	2.2.21	CVE-2011-3368	Improper URL Validation
	2.2.9	CVE-2010-2791	Improper timeouts of keep-alive connections
	2.2.15	CVE-2010-1452	Bad request handling
	2.2.11	CVE-2009-1890	Request content length out of bounds
	2.0.55	CVE-2005-3357	Bad SSL protocol check
OpenSSL	1.0.1f	CVE-2014-0160	Buffer over-read in heartbeat extension
Bash	4.3	CVE-2014-6271	Improper parsing of environment variables

1. Overview
2. Challenges & Approach
3. Implementation
4. Evaluation
5. Application Study
6. Conclusions

- ***Declarative annotation of secrets***
 - New pointer tainting methodology
 - Reduced secret annotation burden
- **New taint propagation semantics**
 - Accurate while containing taint spread and label creep
 - Implemented in LLVM
- Implemented a **memory redactor** for secure honey-patching
- **Tested** on three production web servers

Thank you!

Questions?

Frederico Araujo

(frederico.araujo@utdallas.edu)

<i>programs</i>	$\mathcal{P} ::= \bar{e}$
<i>commands</i>	$c ::= v := e \mid \text{store}(\tau, e_1, e_2) \mid \text{ret}(\tau, e)$ $\quad \mid \text{call}(\tau, e, \overline{args}) \mid \text{br}(e, e_1, e_0)$
<i>expressions</i>	$e ::= v \mid \langle u, \gamma \rangle \mid \Diamond_b(\tau, e_1, e_2) \mid \text{load}(\tau, e)$
<i>binary ops</i>	$\Diamond_b ::=$ typical binary operators
<i>variables</i>	v
<i>values</i>	$u ::=$ values of underlying IR language
<i>types</i>	$\tau ::= ptr \tau \mid \tau \top \mid$ primitive types
<i>taini labels</i>	$\gamma \in (\Gamma, \sqsubseteq)$ (label lattice)
<i>locations</i>	$\ell ::=$ memory addresses
<i>environment</i>	$\Delta : v \rightarrow u$
<i>prog counter</i>	pc
<i>stores</i>	$\sigma : (\ell \rightarrow u) \sqcup (v \rightarrow \ell)$
<i>functions</i>	f
<i>function table</i>	$\phi : f \rightarrow \ell$
<i>taini contexts</i>	$\lambda : (\ell \sqcup v) \rightarrow \gamma$
<i>propagation</i>	$\rho : \top \rightarrow \gamma$
<i>prop contexts</i>	$\mathcal{A} : f \rightarrow \rho$
<i>call stack</i>	$\Xi ::= nil \mid \langle f, pc, \Delta, \top \rangle :: \Xi$

Figure 2: Intermediate representation syntax.

NCS	$\rho_{\{\text{load}, \text{store}\}}(\tau, \gamma_1, \gamma_2) := \gamma_2$
PCS	$\rho_{\{\text{load}, \text{store}\}}(\tau, \gamma_1, \gamma_2) := \gamma_1 \sqcup \gamma_2$
PC ² S	$\rho_{\{\text{load}, \text{store}\}}(\tau, \gamma_1, \gamma_2) := (\tau \text{ is } ptr) ? \gamma_2 : (\gamma_1 \sqcup \gamma_2)$

Figure 4: Polymorphic functions modeling no-combine, pointer-combine, and PC²S label propagation policies.

$$\begin{array}{c}
\frac{}{\sigma, \Delta, \lambda \vdash u \Downarrow \langle u, \perp \rangle} \text{VAL} \quad \frac{}{\sigma, \Delta, \lambda \vdash v \Downarrow \langle \Delta(v), \lambda(v) \rangle} \text{VAR} \\
\frac{\sigma, \Delta, \lambda \vdash e_1 \Downarrow \langle u_1, \gamma_1 \rangle \quad \sigma, \Delta, \lambda \vdash e_2 \Downarrow \langle u_2, \gamma_2 \rangle}{\sigma, \Delta, \lambda \vdash \hat{\diamond}_b(\tau, e_1, e_2) \Downarrow \langle u_1 \hat{\diamond}_b u_2, \gamma_1 \sqcup \gamma_2 \rangle} \text{BINOP} \quad \frac{\sigma, \Delta, \lambda \vdash e \Downarrow \langle u, \gamma \rangle}{\sigma, \Delta, \lambda \vdash \text{load}(\tau, e) \Downarrow \langle \sigma(u), \rho_{\text{load}}(\tau, \gamma, \lambda(u)) \rangle} \text{LOAD} \\
\frac{\sigma, \Delta, \lambda \vdash e \Downarrow \langle u, \gamma \rangle \quad \Delta' = \Delta[v \mapsto u] \quad \lambda' = \lambda[v \mapsto \gamma]}{\langle \sigma, \Delta, \lambda, \Xi, pc, v := e \rangle \rightarrow_1 \langle \sigma, \Delta', \lambda', \Xi, pc + 1, \mathcal{P}[pc + 1] \rangle} \text{ASSIGN} \\
\frac{\sigma, \Delta, \lambda \vdash e_1 \Downarrow \langle u_1, \gamma_1 \rangle \quad \sigma, \Delta, \lambda \vdash e_2 \Downarrow \langle u_2, \gamma_2 \rangle \quad \sigma' = \sigma[u_1 \mapsto u_2] \quad \lambda' = \lambda[u_1 \mapsto \rho_{\text{store}}(\tau, \gamma_1, \gamma_2)]}{\langle \sigma, \Delta, \lambda, \Xi, pc, \text{store}(\tau, e_1, e_2) \rangle \rightarrow_1 \langle \sigma', \Delta, \lambda', \Xi, pc + 1, \mathcal{P}[pc + 1] \rangle} \text{STORE} \\
\frac{\sigma, \Delta, \lambda \vdash e \Downarrow \langle u, \gamma \rangle \quad \sigma, \Delta, \lambda \vdash e_{(u?1:0)} \Downarrow \langle u', \gamma' \rangle}{\langle \sigma, \Delta, \lambda, \Xi, pc, \text{br}(e, e_1, e_0) \rangle \rightarrow_1 \langle \sigma, \Delta, \lambda, \Xi, u', \mathcal{P}[u'] \rangle} \text{COND} \\
\frac{\sigma, \Delta, \lambda \vdash e_1 \Downarrow \langle u_1, \gamma_1 \rangle \quad \dots \quad \sigma, \Delta, \lambda \vdash e_n \Downarrow \langle u_n, \gamma_n \rangle \quad \Delta' = \Delta[\overline{\text{params}}_f \mapsto \overline{u_1 \dots u_n}] \quad \lambda' = \lambda[\overline{\text{params}}_f \mapsto \overline{\gamma_1 \dots \gamma_n}] \quad fr = \langle f, pc + 1, \Delta, \overline{\gamma_1 \dots \gamma_n} \rangle}{\langle \sigma, \Delta, \lambda, \Xi, pc, \text{call}(\tau, f, \overline{e_1 \dots e_n}) \rangle \rightarrow_1 \langle \sigma, \Delta', \lambda', fr :: \Xi, \phi(f), \mathcal{P}[\phi(f)] \rangle} \text{CALL} \\
\frac{\sigma, \Delta, \lambda \vdash e \Downarrow \langle u, \gamma \rangle \quad fr = \langle f, pc', \Delta', \bar{\gamma} \rangle \quad \lambda' = \lambda[v_{\text{ret}} \mapsto \mathcal{A} f \bar{\gamma}]}{\langle \sigma, \Delta, \lambda, fr :: \Xi, pc, \text{ret}(\tau, e) \rangle \rightarrow_1 \langle \sigma, \Delta'[v_{\text{ret}} \mapsto u], \lambda', \Xi, pc', \mathcal{P}[pc'] \rangle} \text{RET}
\end{array}$$

Figure 3: Operational semantics of a generalized label propagation semantics.