


Whole-Program Privilege and Compartmentalization Analysis with the Object-Encapsulation Model

Yudi Yang, Weijie Huang, Kelly Kaoudis, Nathan Dautenhahn
LangSec 2023




Computing systems contain over-privileged components



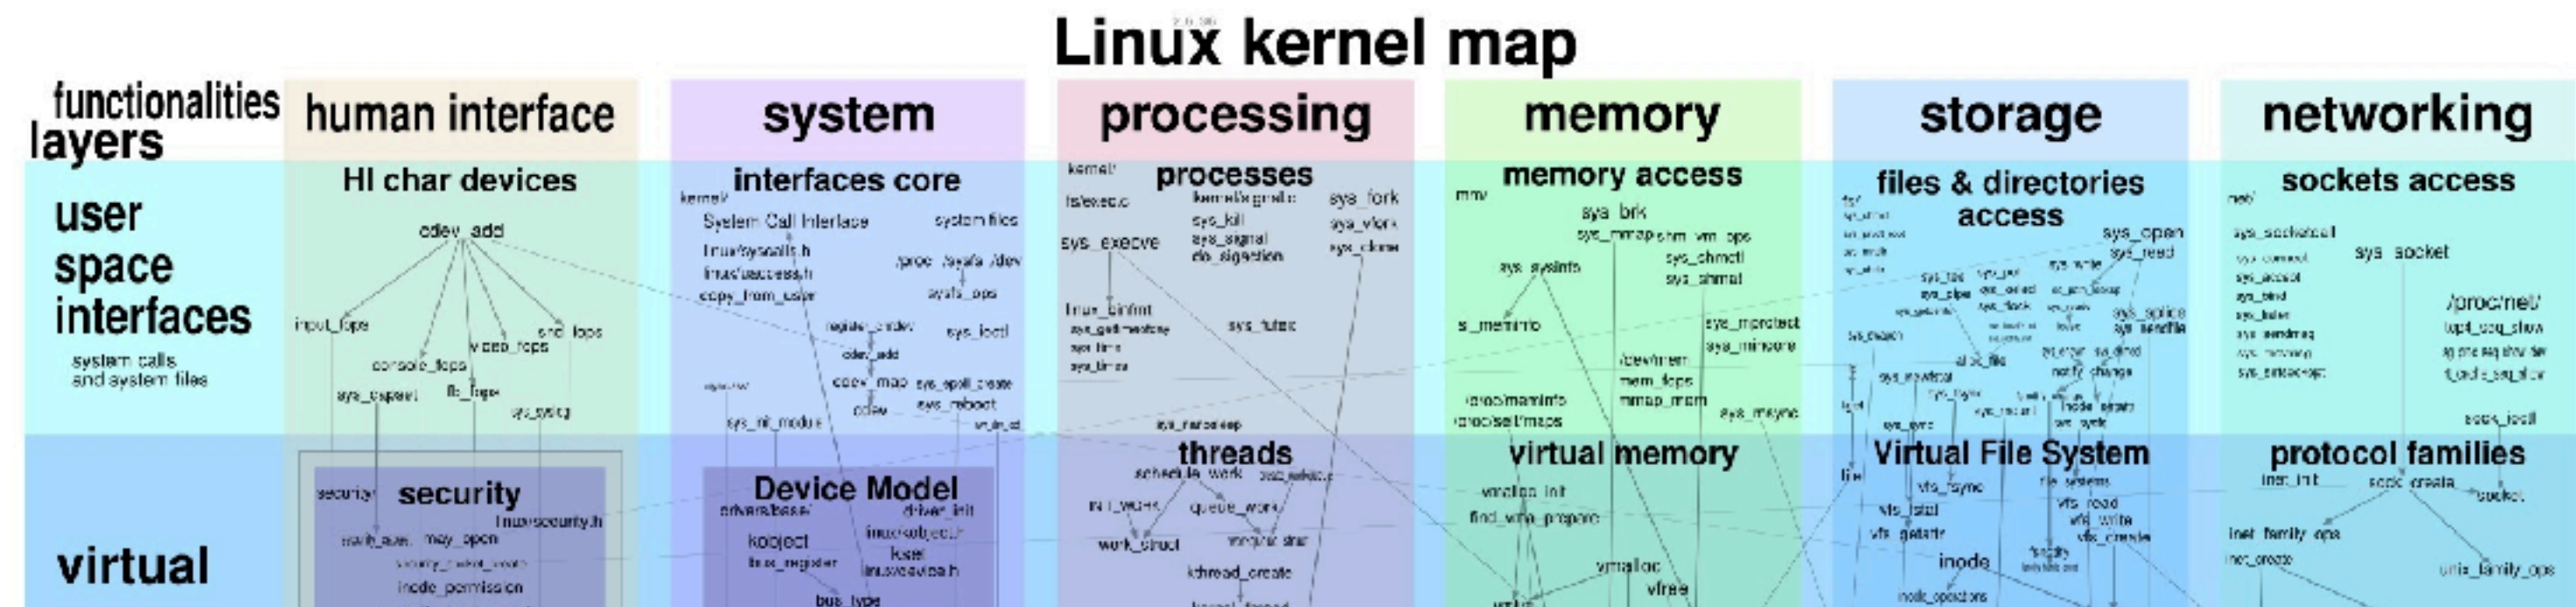
cr: <http://www2.ece.ohio-state.edu/citr/Demo97/>

Ideally, restrict access to what it needs to do its job through compartmentalization

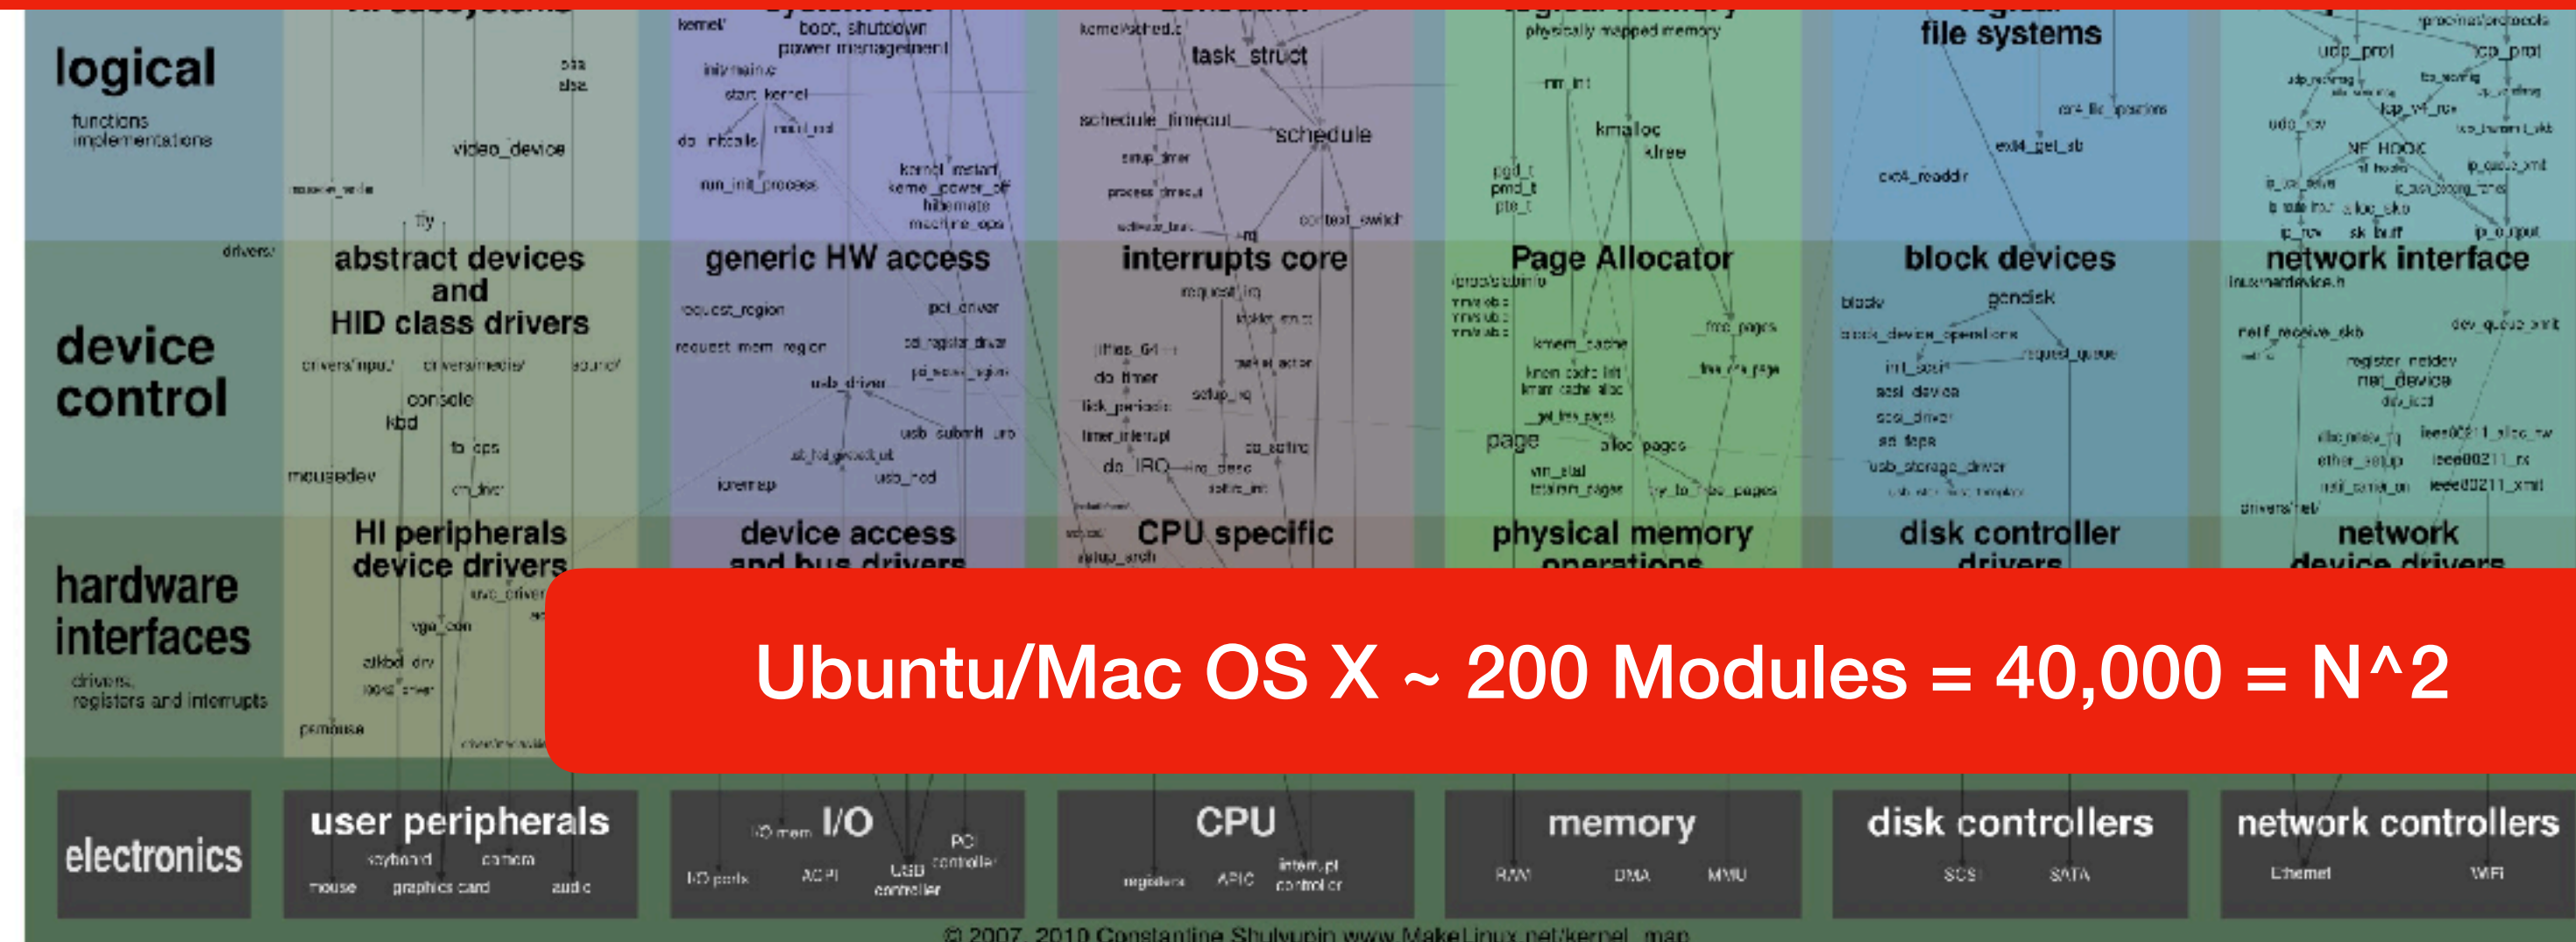


Stop and contain attacks through boundaries

But how do we retrofit into a system that was not built for compartmentalization?




Mapping such a policy to a whole-system is not scalable: what's a privilege and where to put boundaries?





Ubuntu/Mac OS X ~ 200 Modules = 40,000 = N^2

Objective: automatically derive and systematically evaluate policy



We don't know labels on the sources and destinations of messages and operations

Unknown and complex




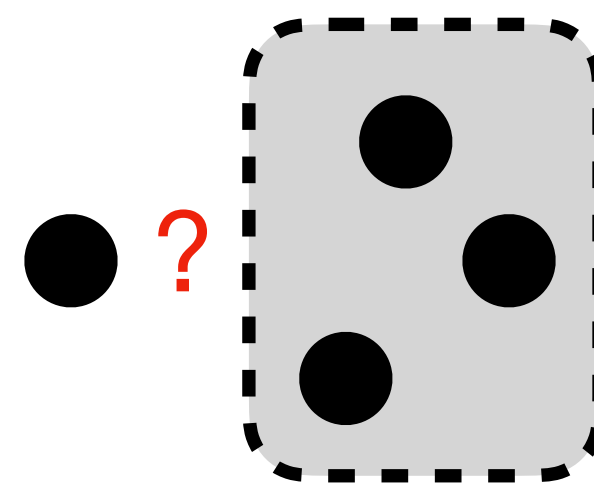

Decomposed and organized

A dramatic sunset over a field with a line of trees. The sky is filled with dark, heavy clouds that are illuminated from below by the setting sun, creating a vibrant orange and yellow glow. The sun is positioned just above the horizon, which is marked by a dense line of dark trees. The foreground consists of a field of tall grasses, some of which are lit by the low sun, creating a golden hue. The overall scene is one of natural beauty and tranquility.

Hypothesis: emergent properties of modularity provide baseline policy

Key questions and approach overview

- What is a privilege?
- How to select boundaries?
- How to systematically evaluate?
- Program Capability Graph
- Object-Encapsulation Model
- Separability Analysis




What is a privilege?

Insight #1: code as privilege


- Interaction between functions and modules: subjects
- Message sending through read / write / execute: operations
- Communication is a privilege: capabilities
- Derive through static or dynamic analysis



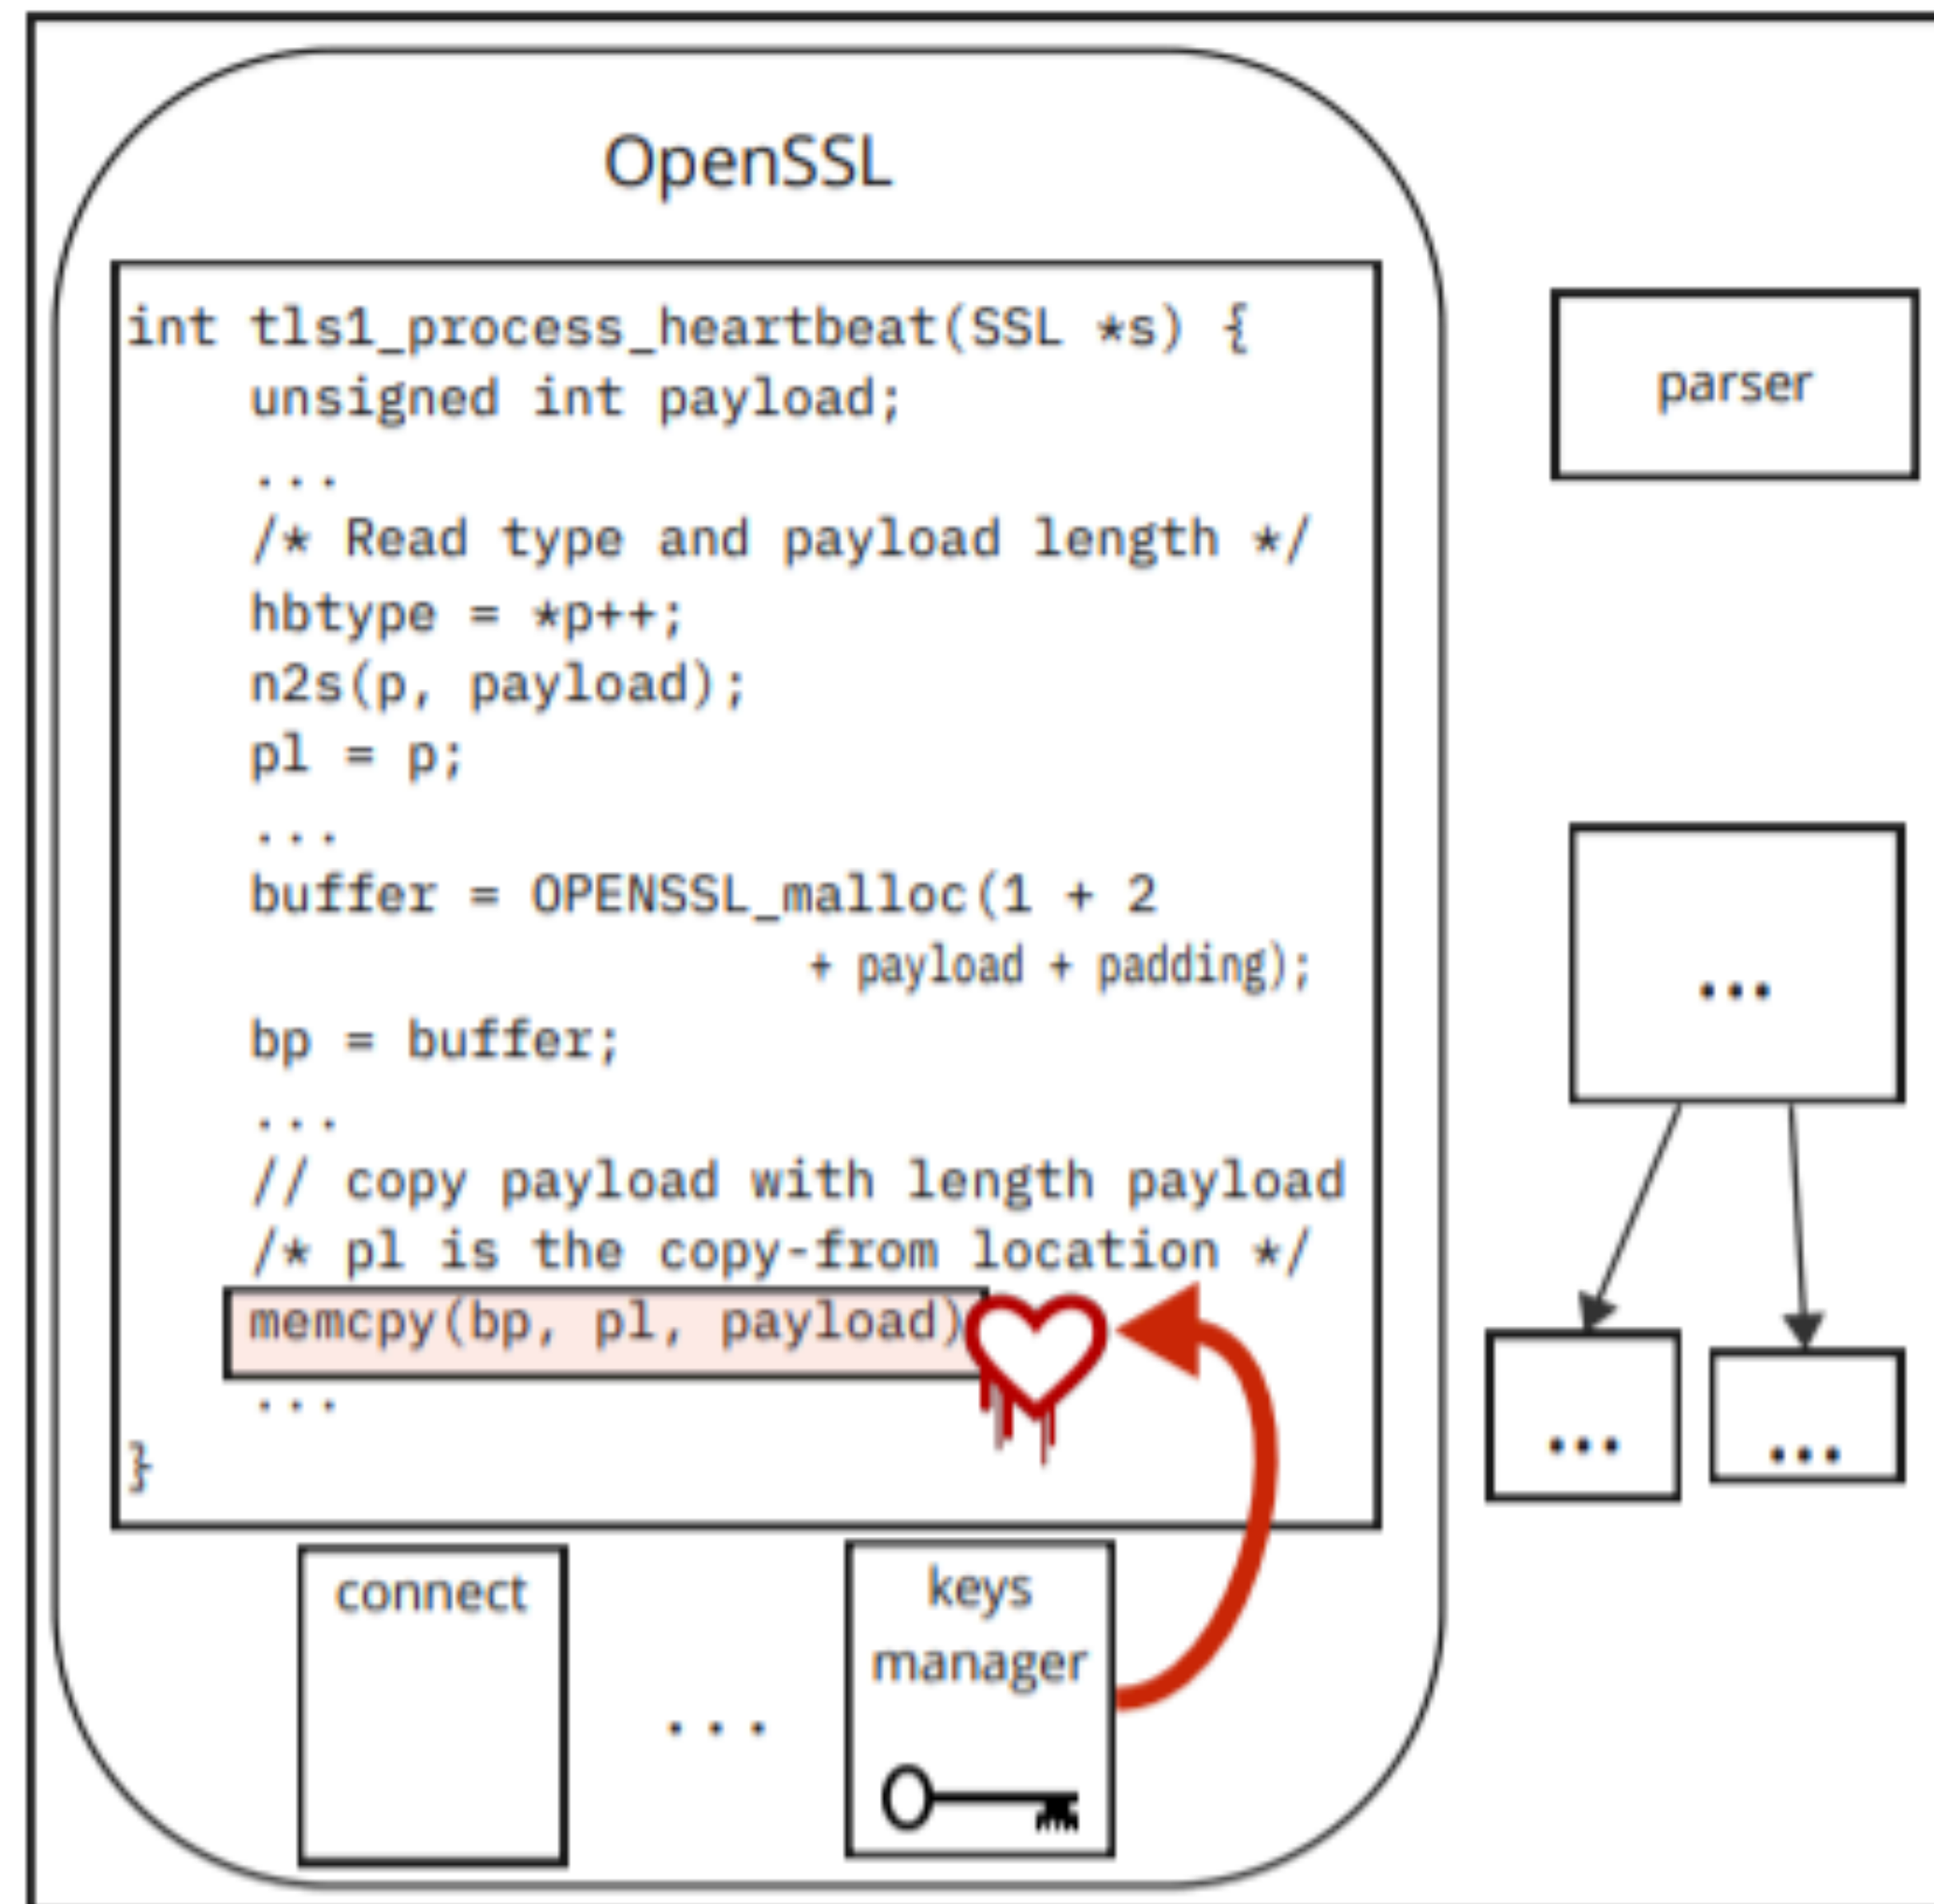
Human Level




Human Level



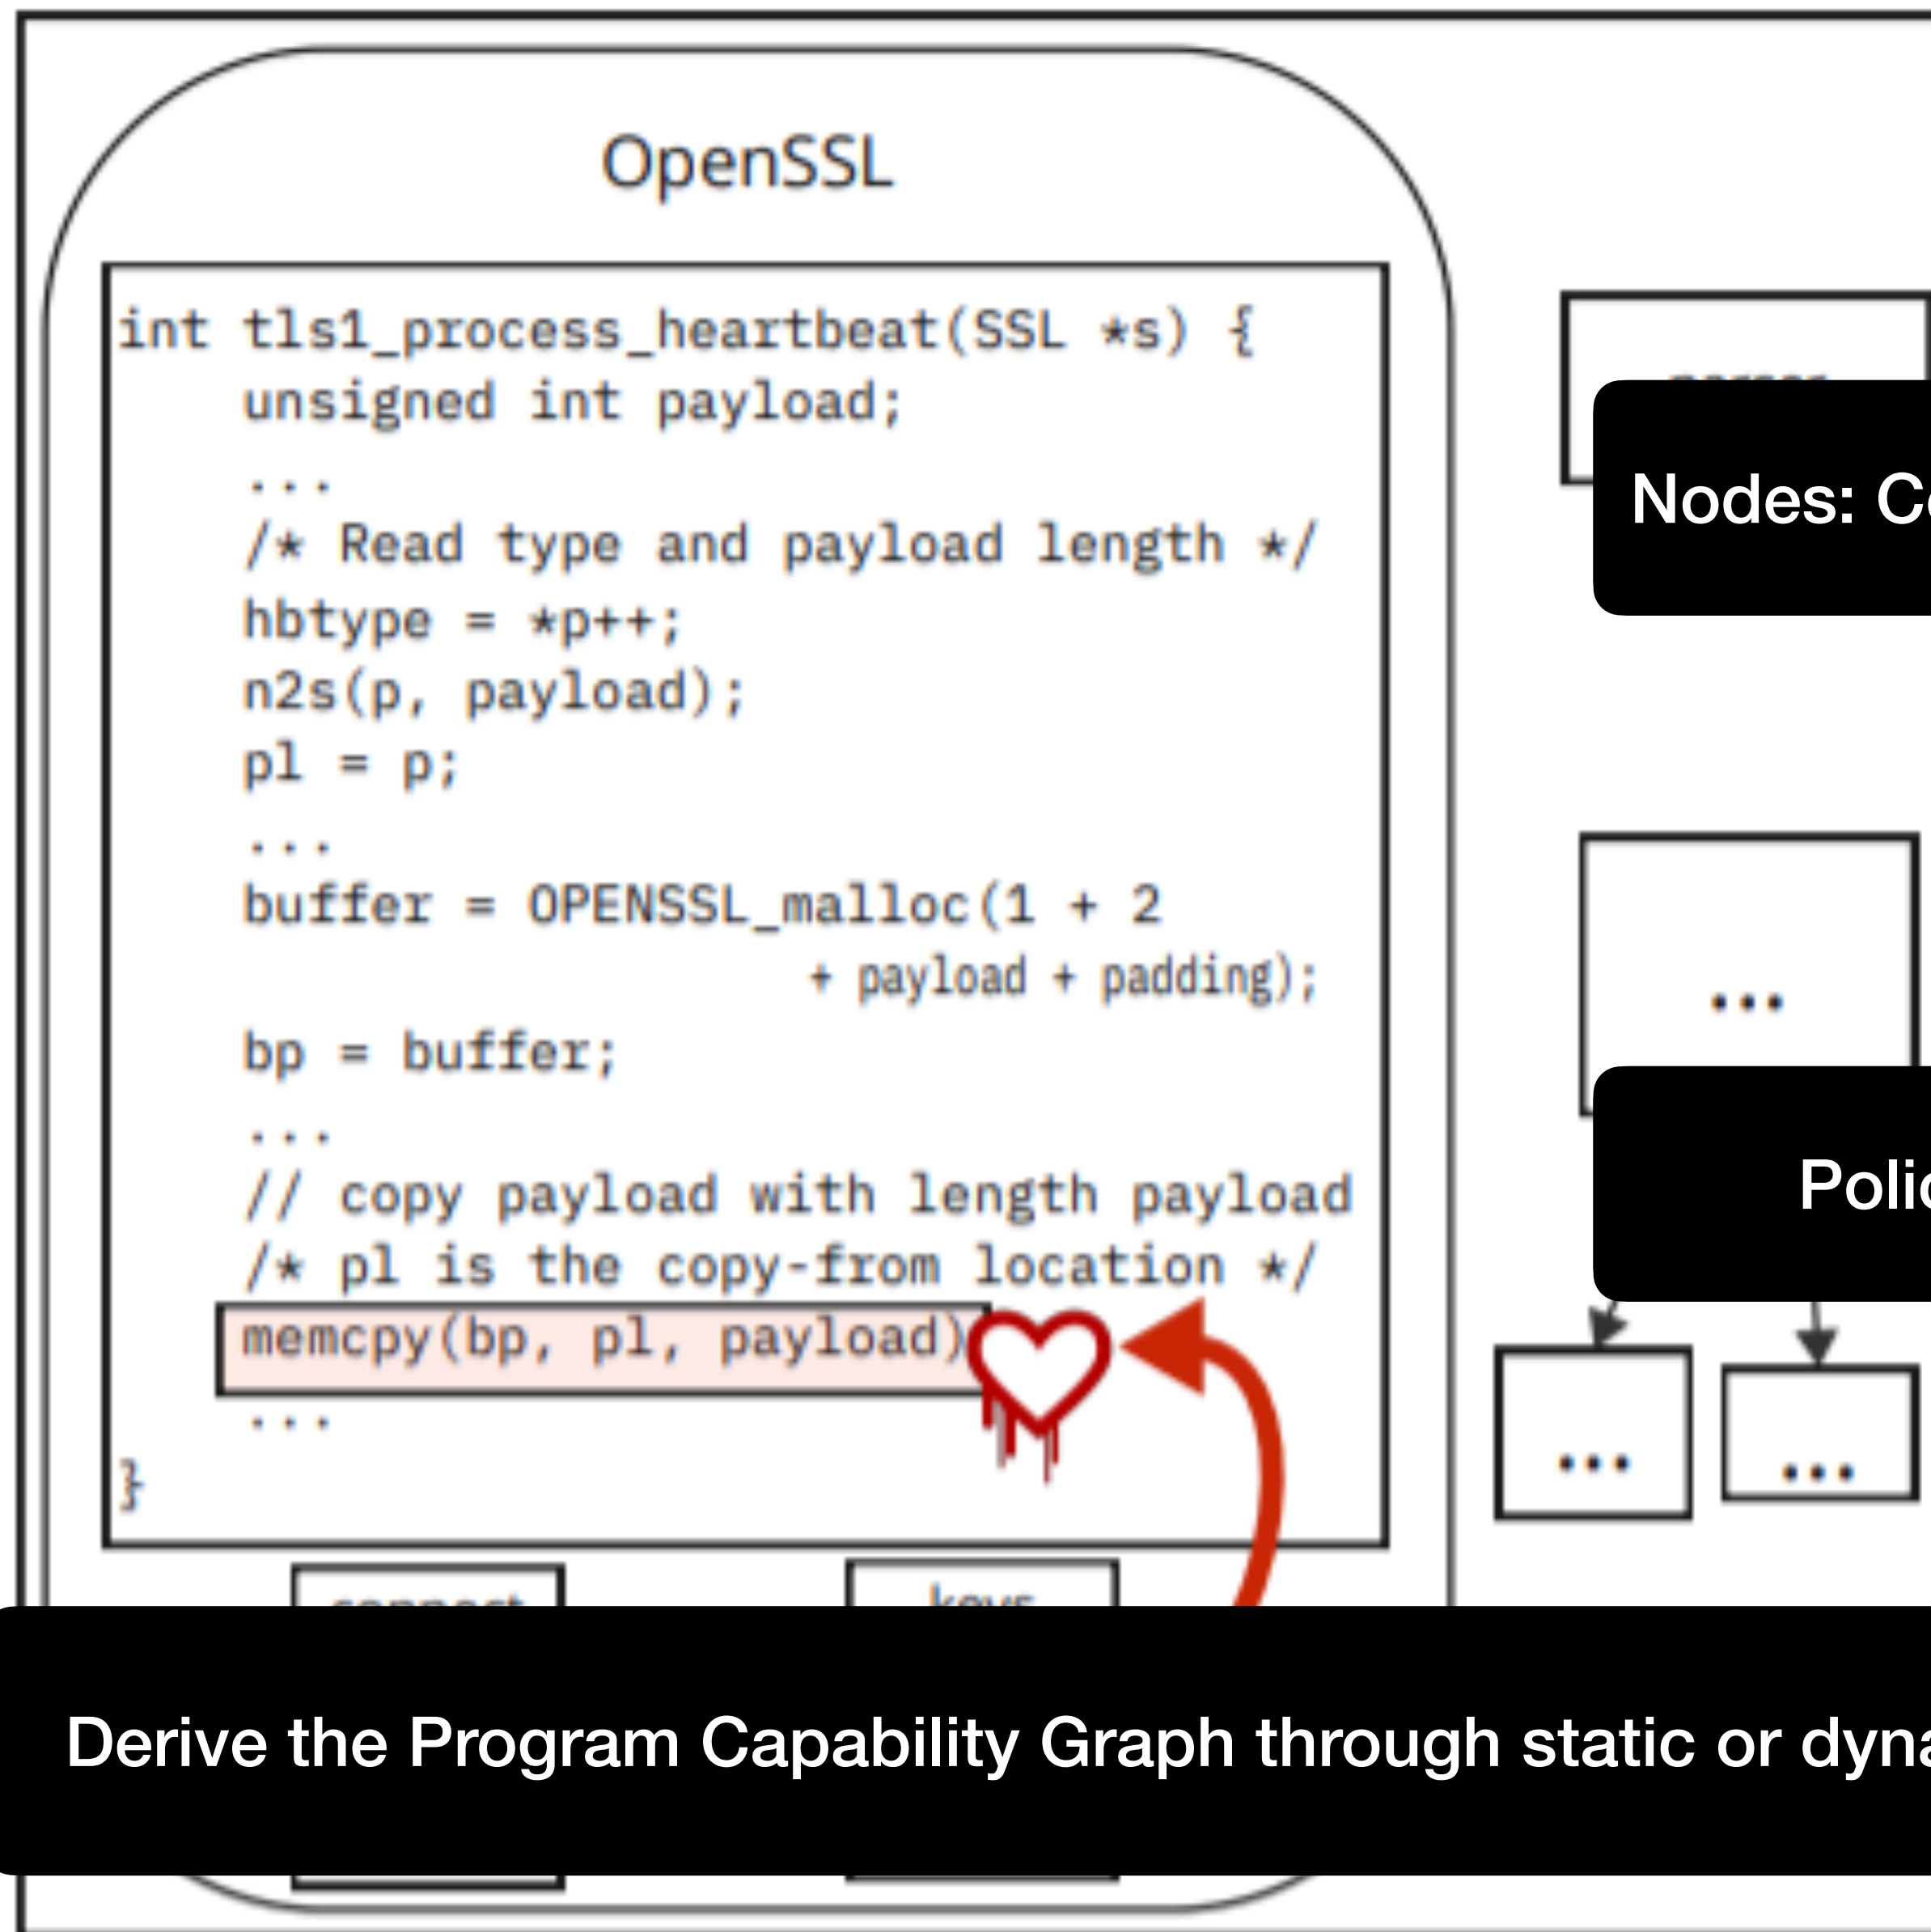
Server Source Code



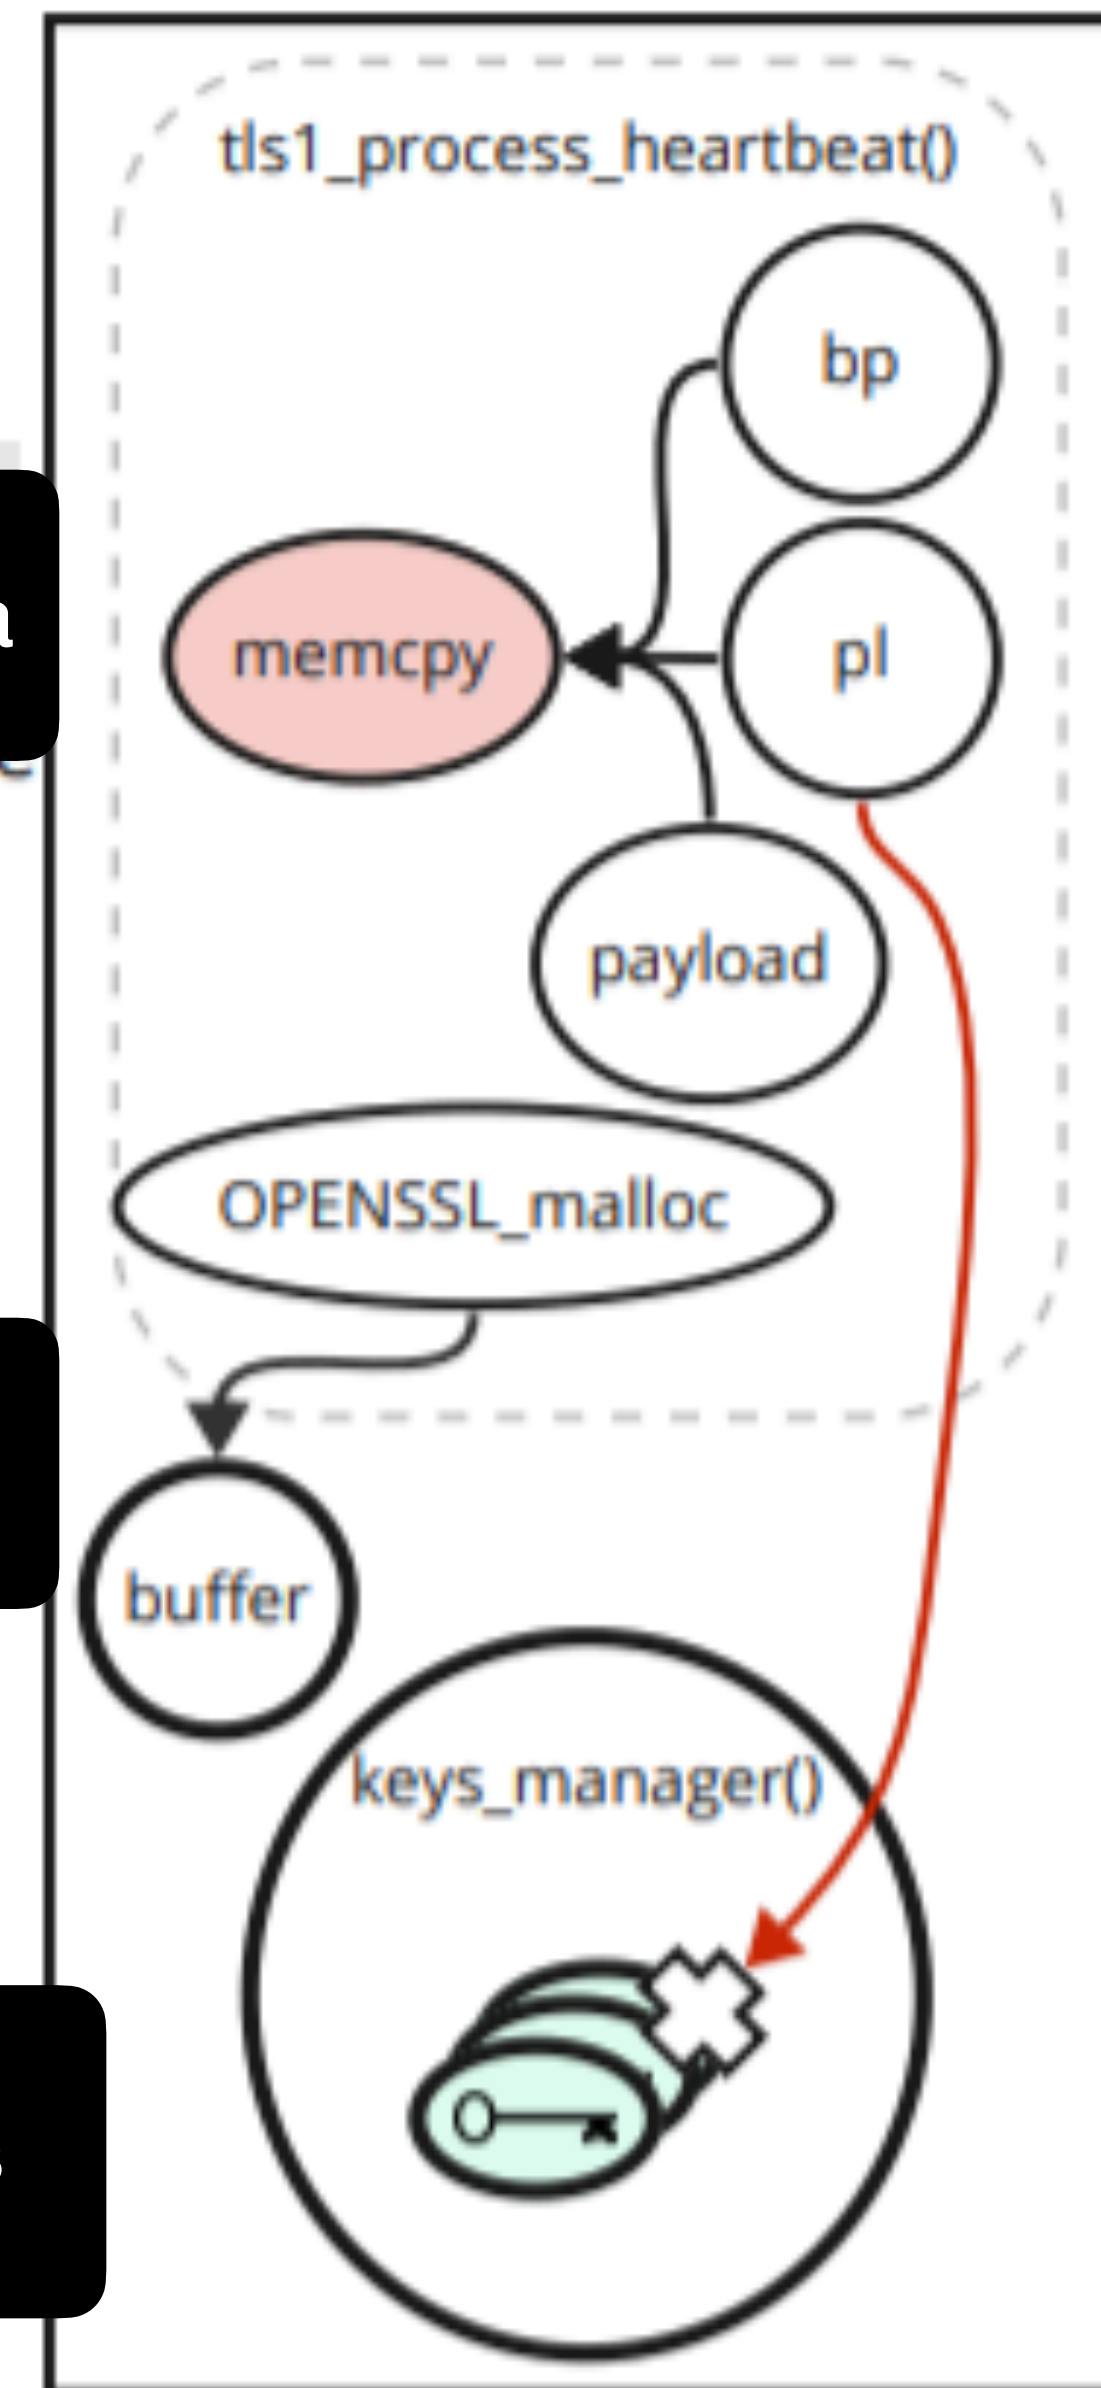
Human Level



Server Source Code



Program Capability Graph



Nodes: Code and Data


Policy: RWX

Derive the Program Capability Graph through static or dynamic analysis


What is a meaningful compartmentalization?

Insight #2: Emergent modularity as compartments

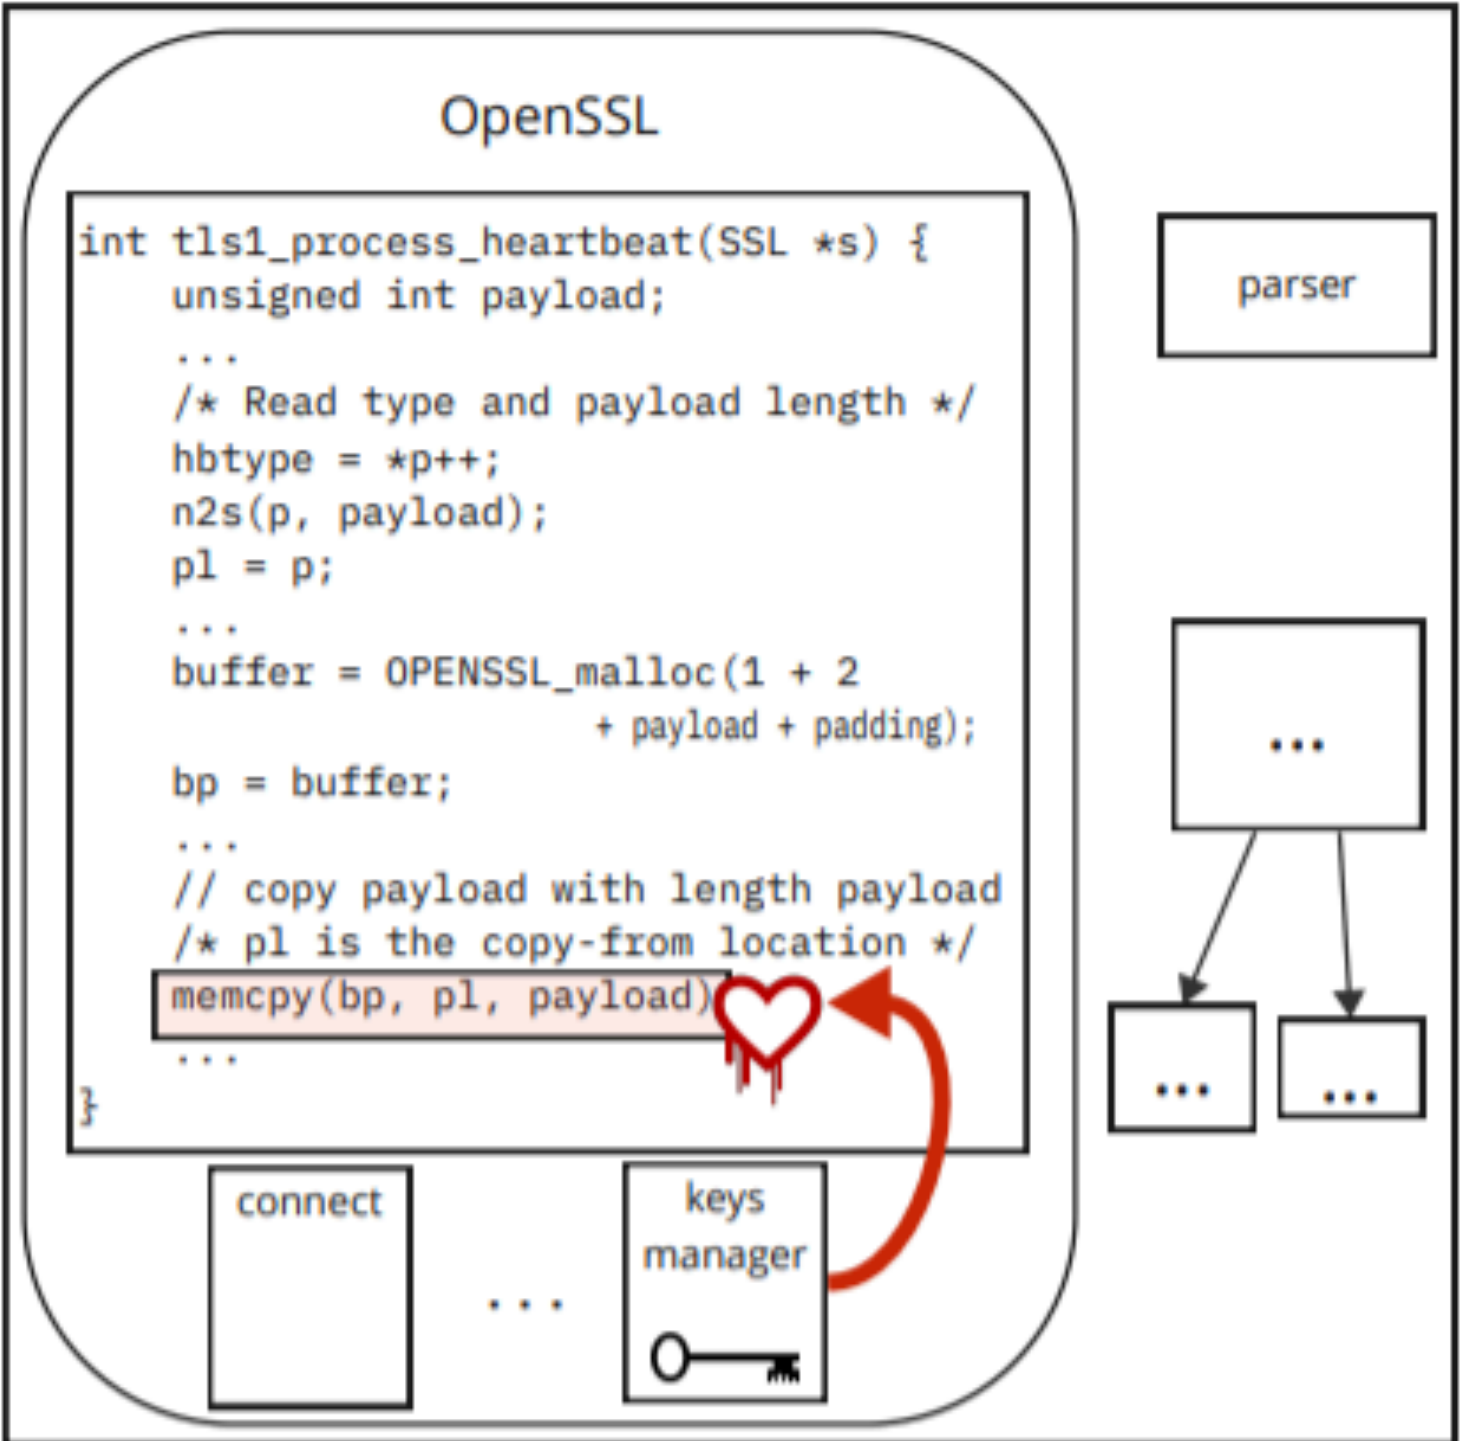
Developers use encapsulation to manage complexity



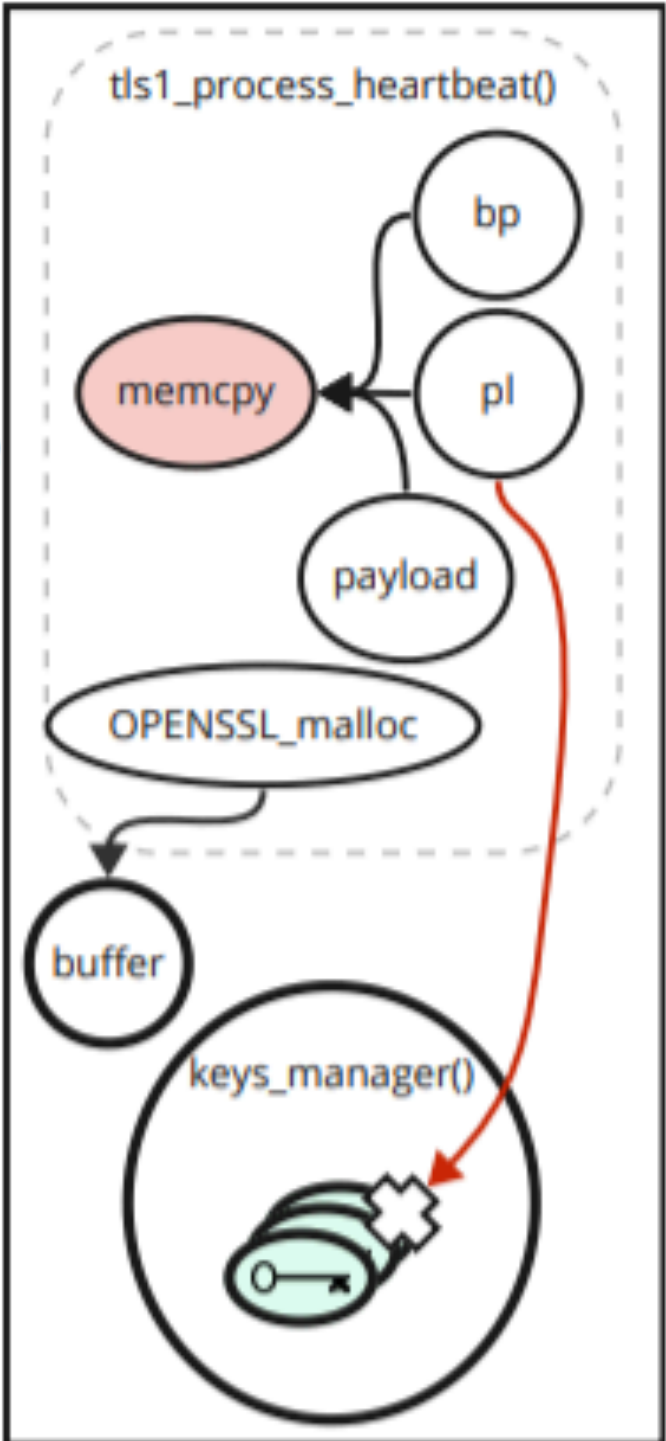
Human Level



Server Source Code




Program Capability Graph

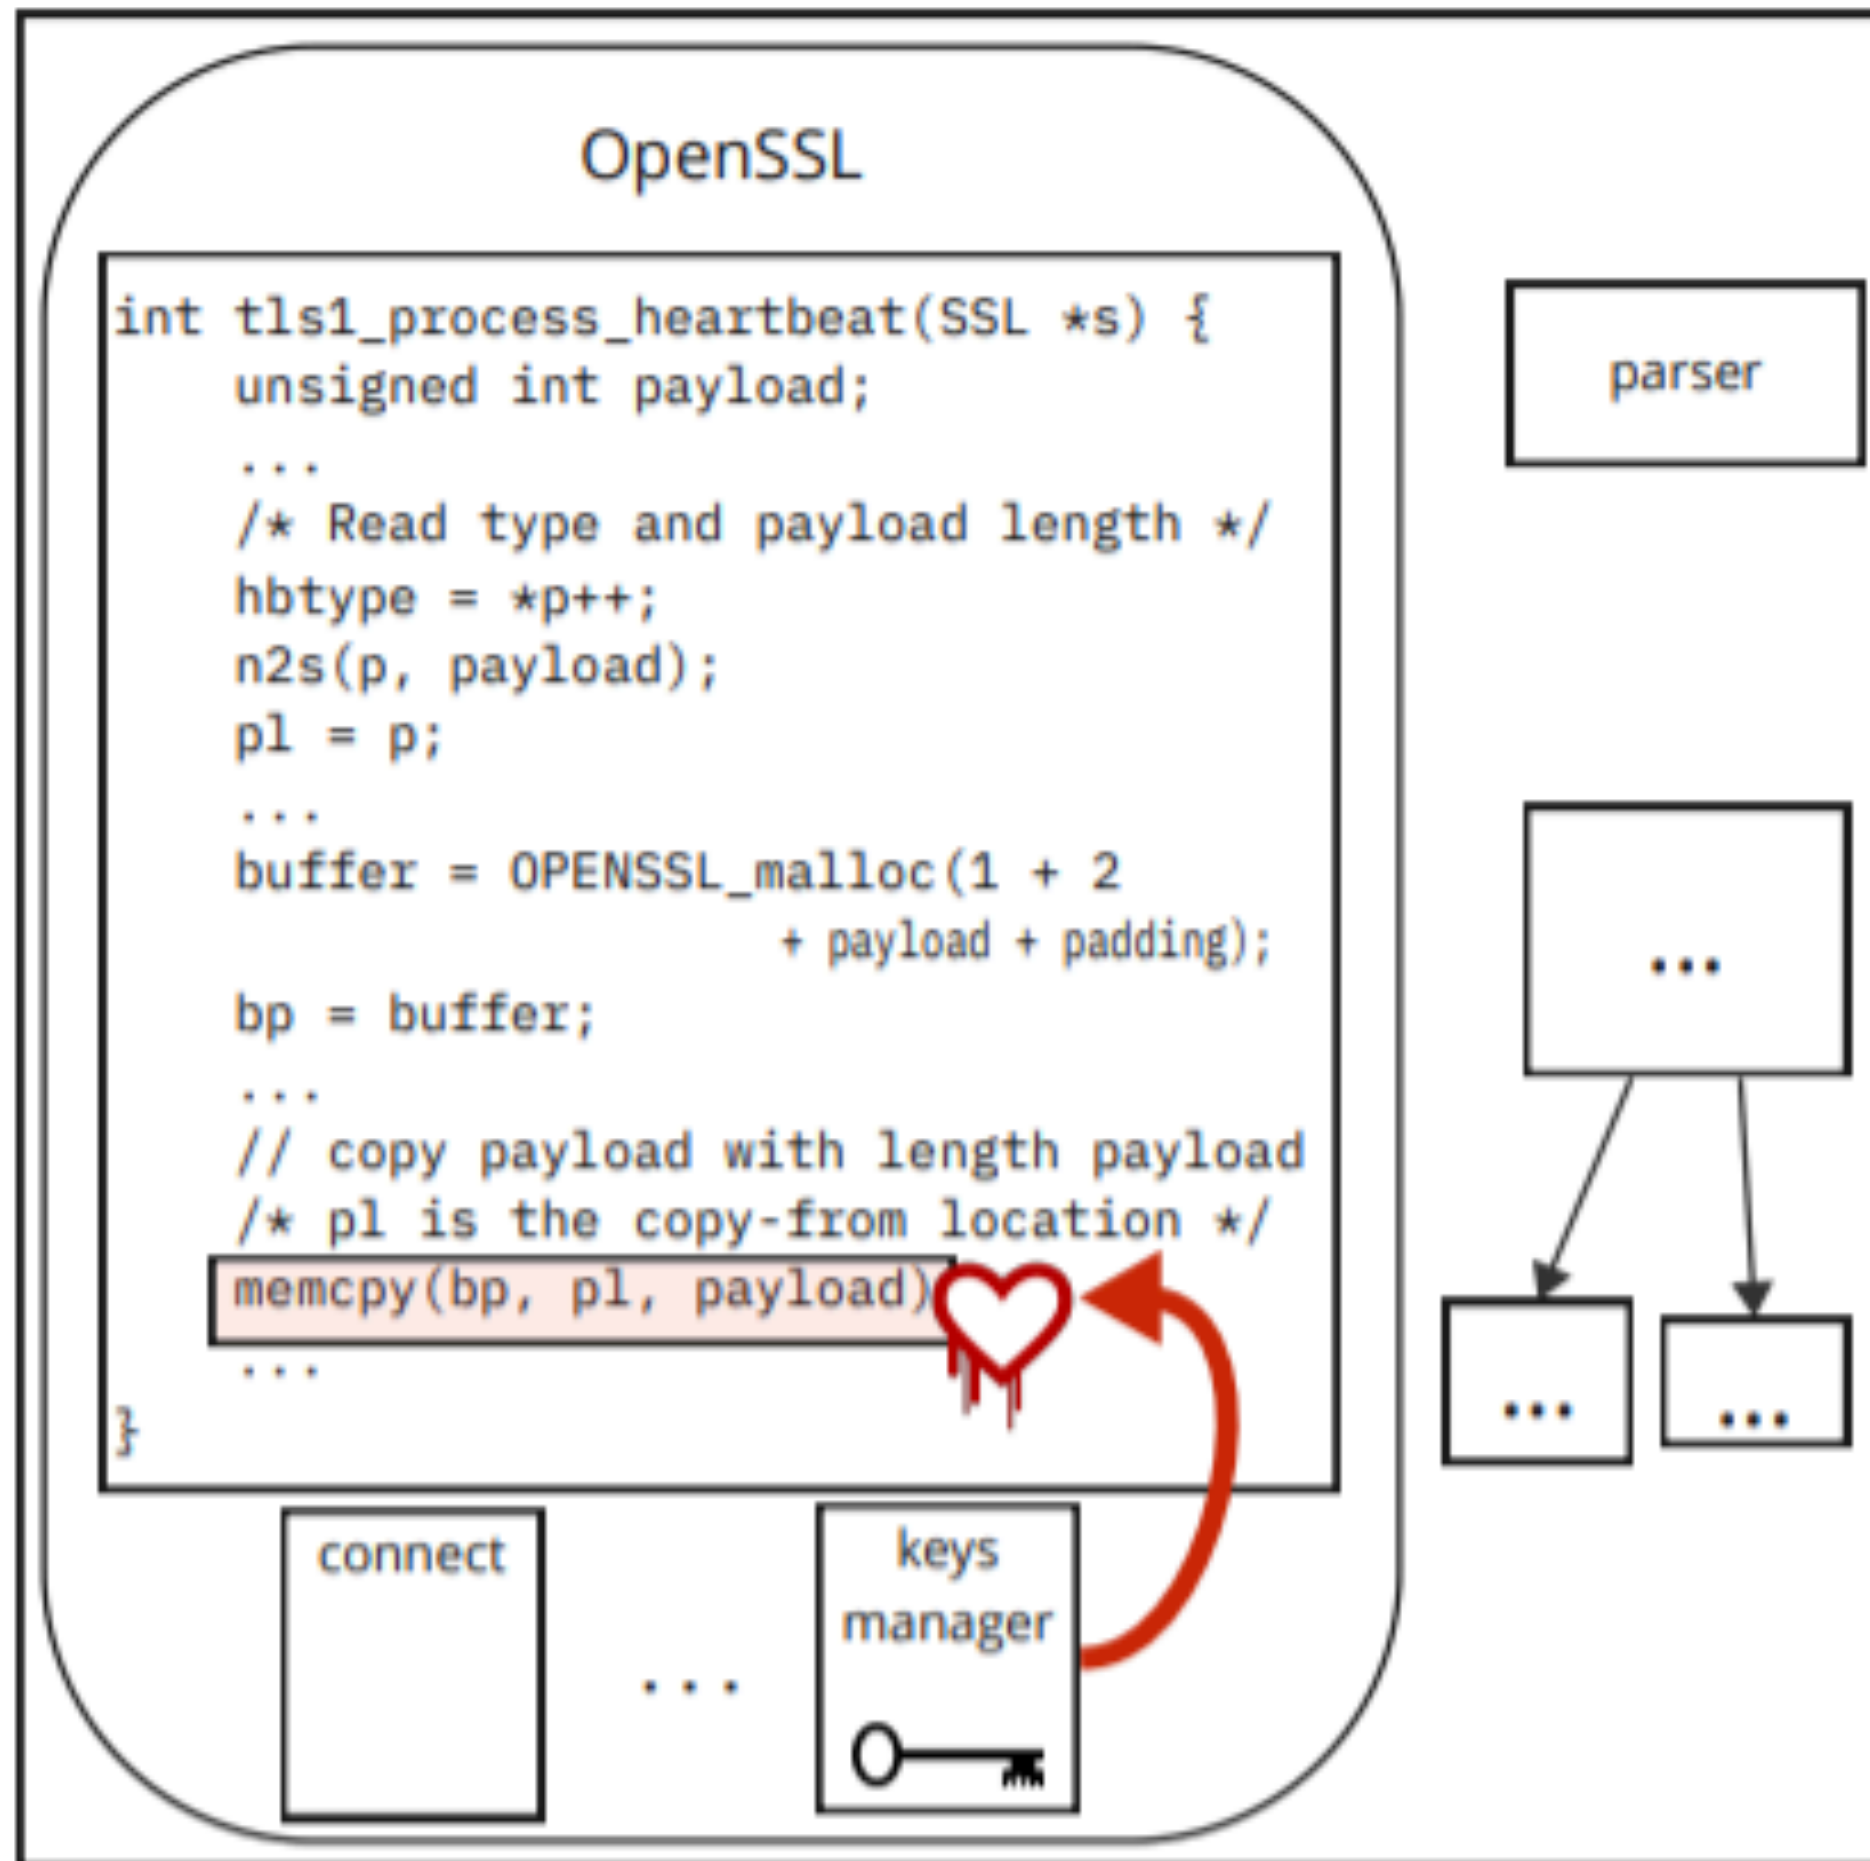


Threat Tagging
 ○ Sensitive
 ○ Suspect

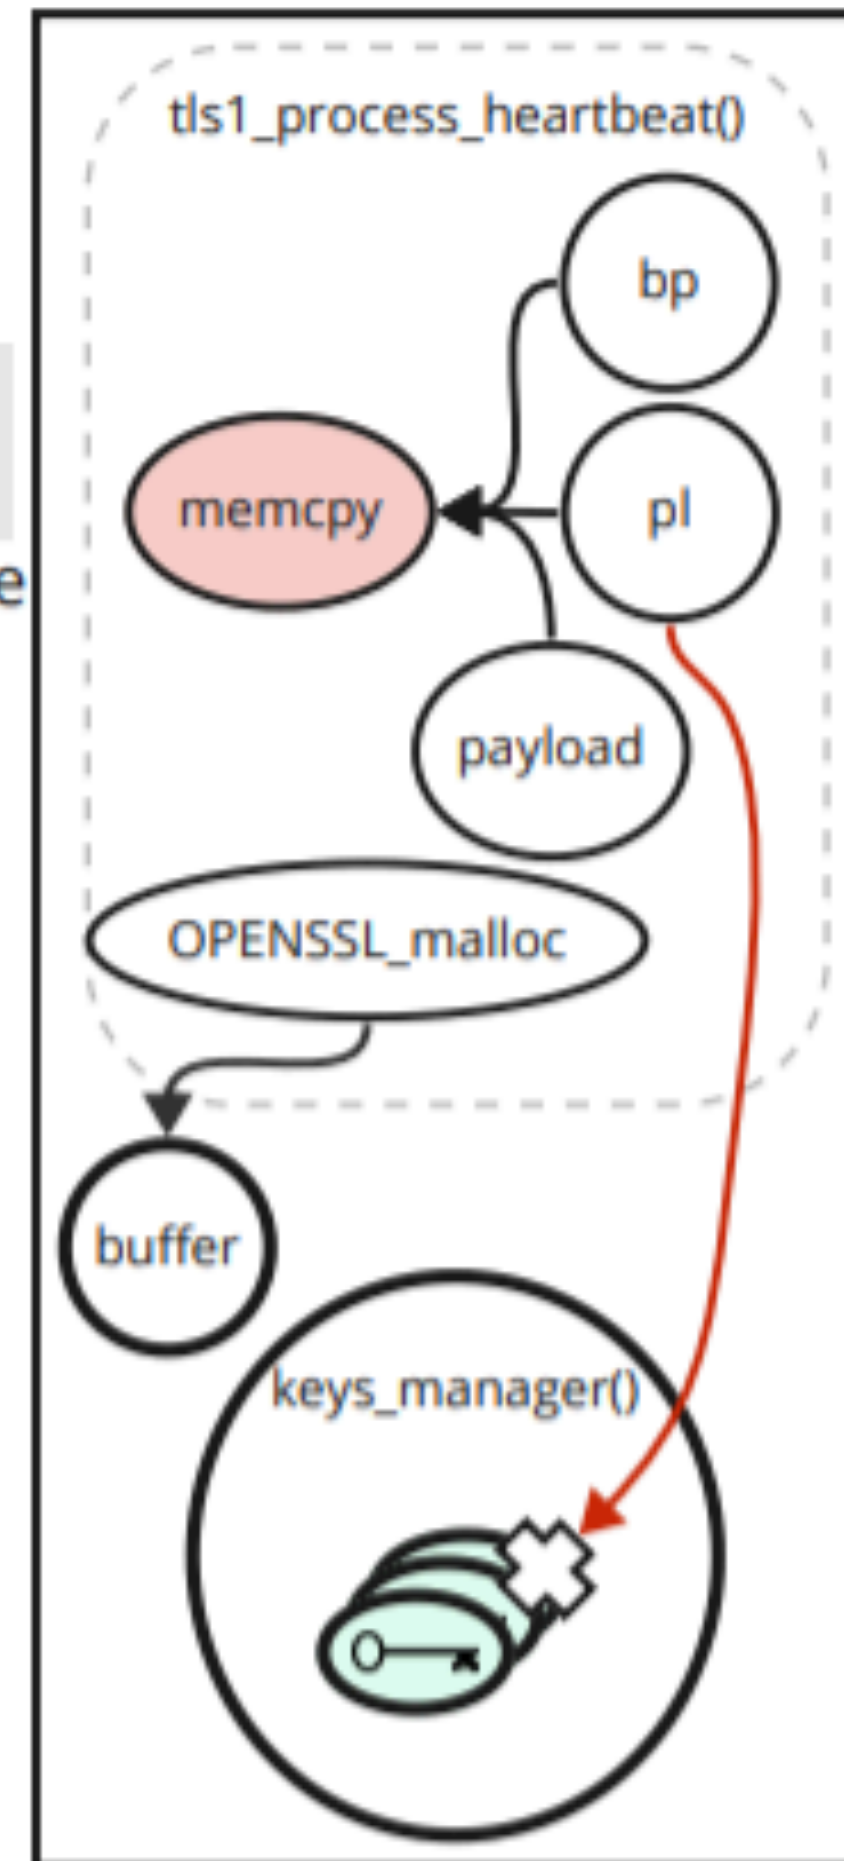
Human Level



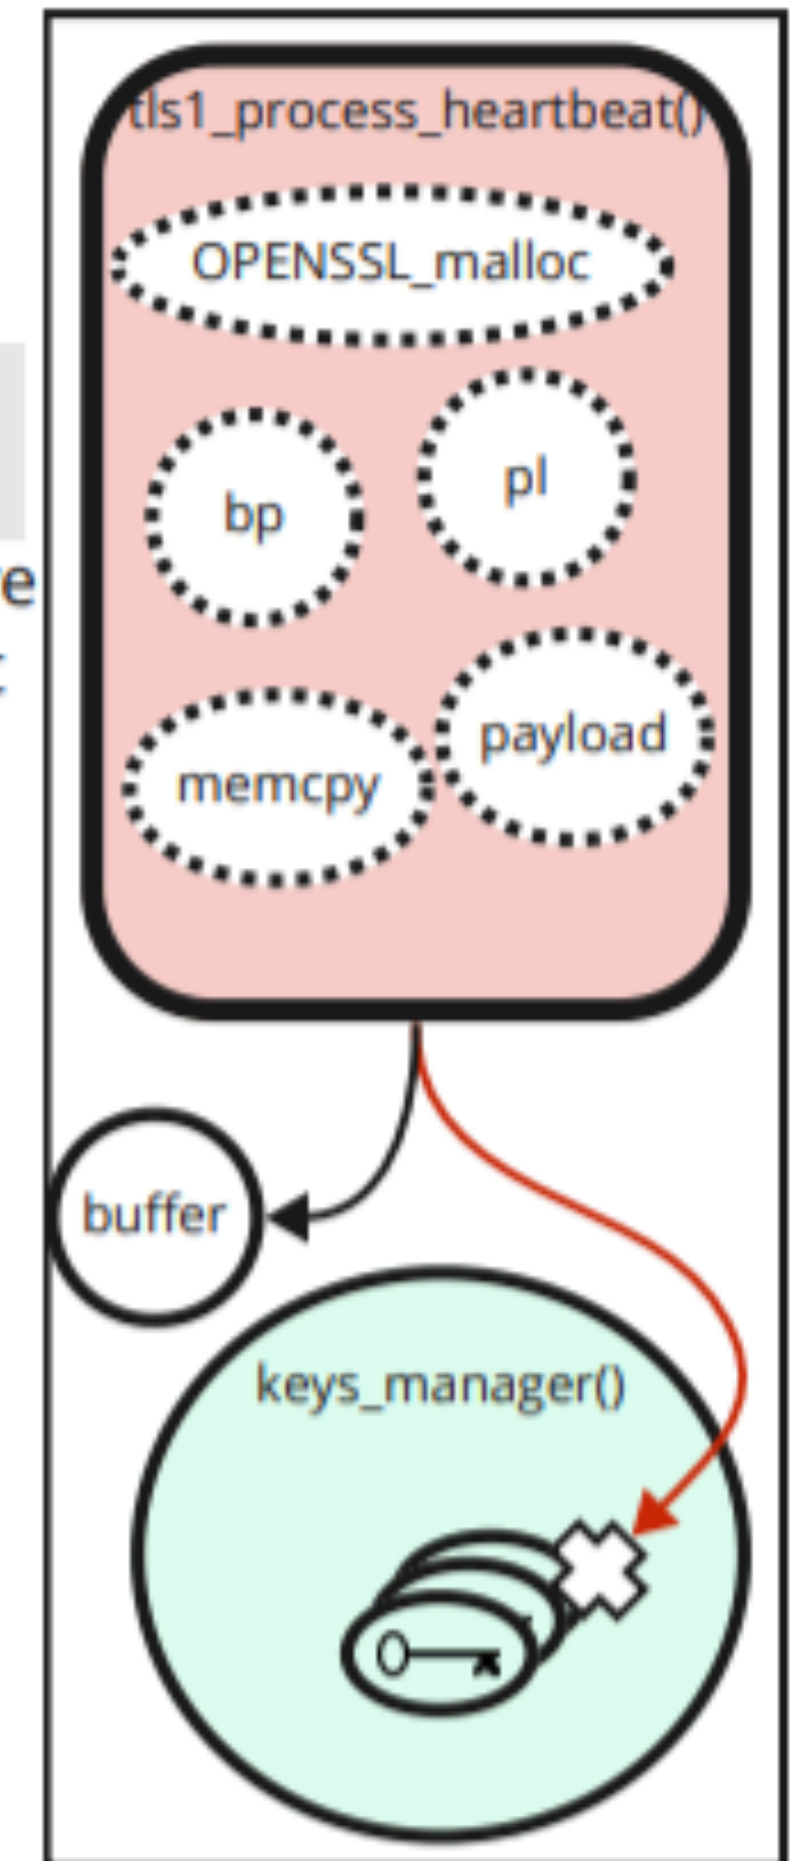
Server Source Code



Program Capability Graph



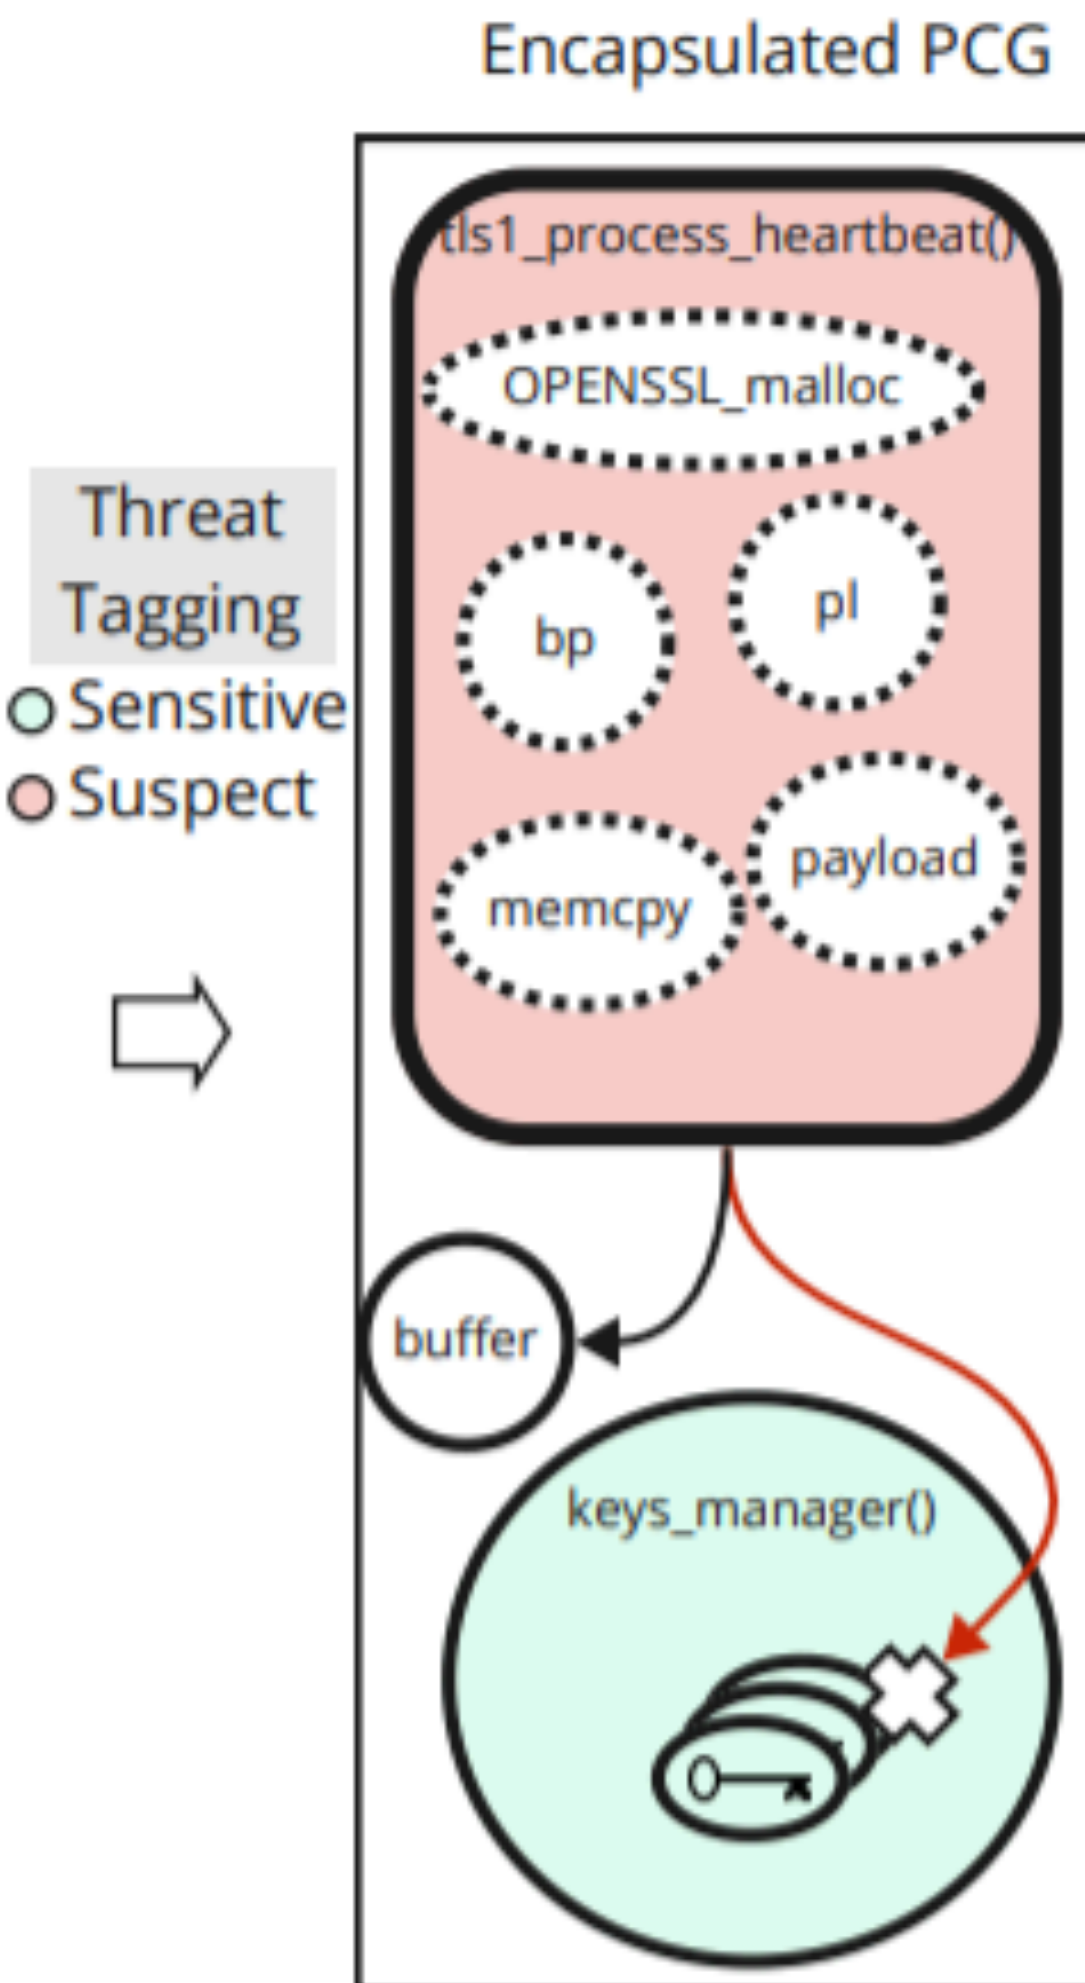
Encapsulated PCG



Group operations into encapsulated components with private and public data and interfaces

What is a compartment?

Answer: Object-Encapsulation Model



Partition based on belongs-to: tag each object and code as belonging to a specific compartment

Where do we start?

Solution: Lexical Scopes as Objects

Server Source Code


OpenSSL

```
int tls1_process_heartbeat(SSL *s) {  
    unsigned int payload;  
    ...  
    /* Read type and payload length */  
    hbtype = *p++;  
    n2s(p, payload);  
    pl = p;  
    ...  
    buffer = OPENSSL_malloc(1 + 2  
        + payload + padding);  
    bp = buffer;  
    ...  
    // copy payload with length payload  
    /* pl is the copy-from location */  
    memcpy(bp, pl, payload);  
    ...  
}
```

connect


keys

parser



Program Capability Graph

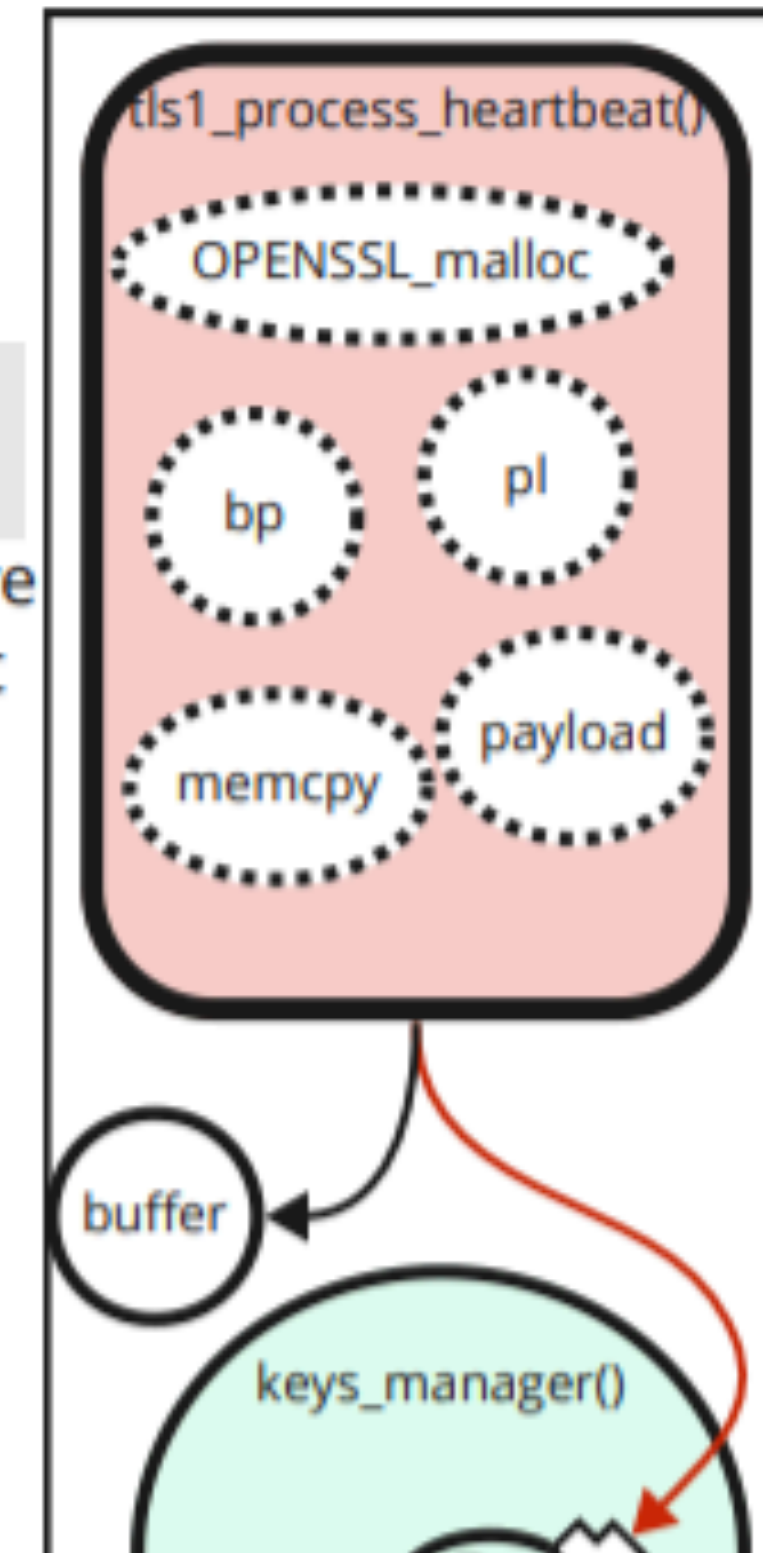
Threat
Tagging
○ Sensitive
○ Suspect



Threat
Tagging
○ Sensitive
○ Suspect



Encapsulated PCG



Any label will do, but can use file scope as base compartment labels

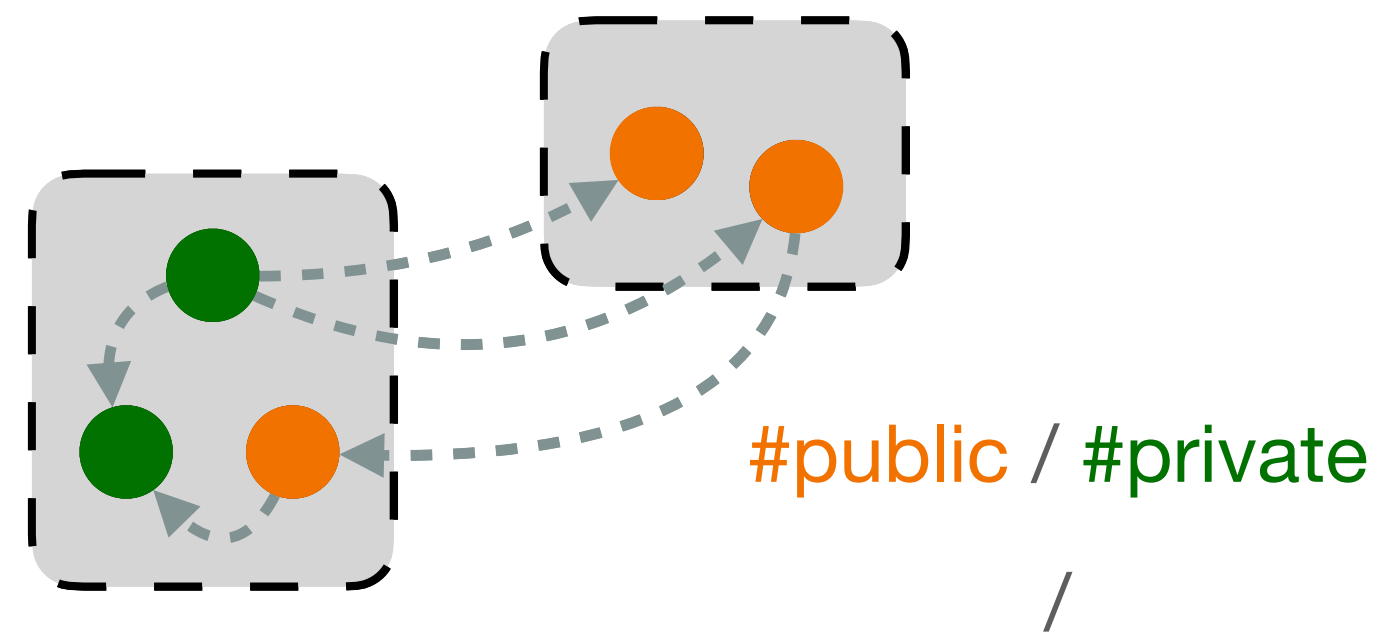
How to evaluate it?

Solution: Separability Analysis

- Privilege reduced: encapsulation analysis
- Security improved: threat modeling and measurement
- Compatibility reduced: 100%
- Transformation effort: Ongoing work
- Enforcement cost: Ongoing work

Encapsulation Analysis


- The External Access Ratio metric



- The PS-From Ratio metric



External Access Ratio Metric



75% of objects in the Linux Kernel are only ever accessible to the file that allocated the object!




File-based Compartmentalization Accessibility

System	# of Files	CLOC (in C)	Encobj _s .*	Objects*	Scalars*	% Reducible
libxml2	43	215,796	22	176	31,292	48.12%
jansson	14	7,529	4	15	607	70.88%
libroxml	12	7,205	3	10	1,199	74.24%
libexpat	8	25,452	2	13	1,259	76.79%
facil.io	30	22,602	6	39	851	78.28%
json-c	14	8,501	2	11	127	82.97%
Nginx	129	138,467	61	484	96,464	52.60%
CPython	262	530,504	106	6,327	394,867	59.53%
MuPDF	542	860,824	175	1,045	62,128	67.62%
PHP	479	1,162,182	144	2,167	236,744	69.82%
gimp	1,103	894,939	8	21	138	99.29%

* Accessibility per encapsulated-object in Average

PS-From Ratio Metric


Comparative Analysis of Parsers



Comparative study of JSON libraries using PSFR


Name	Parser Data	Parser Code	External Data	External Code	Inaccessible Data	Inaccessible Code
json-c	67 (7.70%)	16 (1.84%)	8 (0.92%)	19 (2.18%)	580 (66.67%)	180 (20.69%)
facil.io	3284 (73.63%)	1078 (24.17%)	45 (1.01%)	0 (0.00%)	53 (1.19%)	0 (0.00%)
jansson	138 (13.5%)	32 (3.14%)	14 (1.37%)	26 (2.55%)	650 (63.79%)	159 (15.60%)
cJSON	325 (74.37%)	112 (25.63%)	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)

External Access Ratio Metric (Cont.)




Threat model labeling and analysis

- Threat Labeling: Sensitive/Suspicious
- Access Distance Metric
- Exposure Reduction Metric



Access Distance Metric



execve

Distance: 1
#Files: 1 24

Distance: 2
#Files: 10


Distance: 3
#Files: 47

Distance: 4
#Files: 50

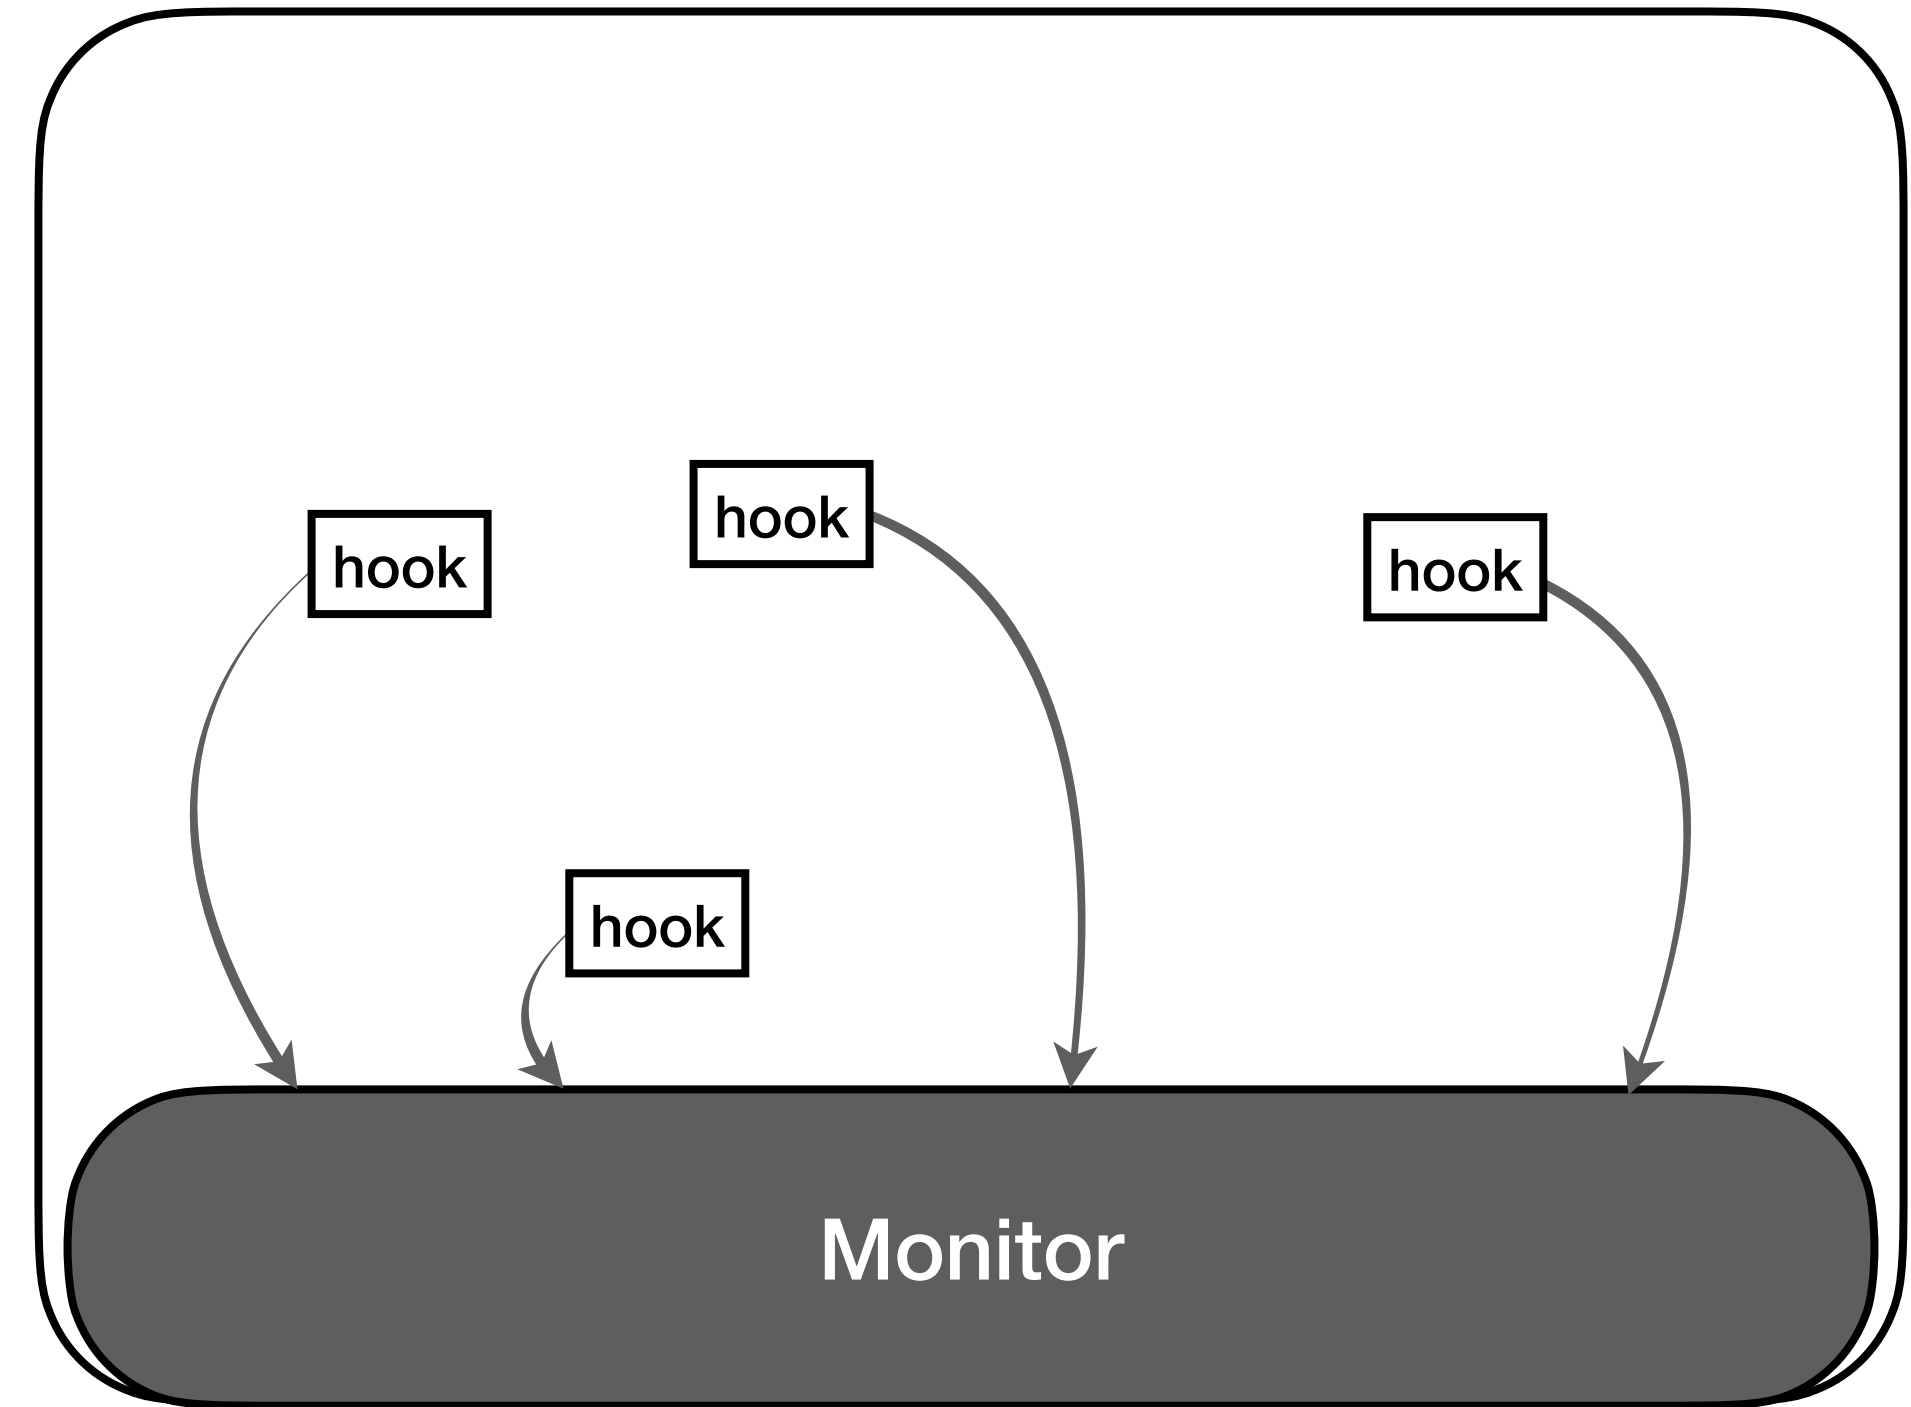


Distance: 5
#Files: 3

Exposure Reduction Metric

System	System Interfaces										
	system	fork	exec*	popen	open*	read*	write*	ioctl	dup	mmap	mprotect
libexpat	0 / 16	0 / 16	0 / 16	0 / 16	1 / 16	1 / 16	0 / 16	0 / 16	0 / 16	0 / 16	0 / 16
libxml	0 / 40	0 / 40	0 / 40	0 / 40	2 / 40	1 / 40	1 / 40	0 / 40	1 / 40	0 / 40	0 / 40
jansson	0 / 37	0 / 37	0 / 37	0 / 37	3 / 37	2 / 37	1 / 37	0 / 37	0 / 37	0 / 37	0 / 37
json-c	0 / 56	0 / 56	0 / 56	0 / 56	2 / 56	2 / 56	1 / 56	0 / 56	0 / 56	0 / 56	0 / 56
facil.io	0 / 102	1 / 102	0 / 102	0 / 102	4 / 102	5 / 102	4 / 102	0 / 102	0 / 102	1 / 102	0 / 102
libxml2	0 / 141	0 / 141	0 / 141	0 / 141	3 / 141	3 / 141	4 / 141	0 / 141	1 / 141	0 / 141	0 / 141
Nginx	0 / 260	2 / 260	1 / 260	0 / 260	10 / 260	4 / 260	4 / 260	4 / 260	2 / 260	2 / 260	0 / 260
CPython	1 / 600	1 / 600	1 / 600	0 / 600	4 / 600	6 / 600	4 / 600	2 / 600	1 / 600	2 / 600	1 / 600
MuPDF	1 / 613	0 / 613	0 / 613	0 / 613	8 / 613	8 / 613	5 / 613	0 / 613	0 / 613	0 / 613	0 / 613
PHP	0 / 1056	2 / 1056	1 / 1056	4 / 1056	16 / 1056	8 / 1056	7 / 1056	0 / 1056	2 / 1056	5 / 1056	1 / 1056
Gimp	0 / 4086	1 / 4086	1 / 4086	0 / 4086	4 / 4086	4 / 4086	1 / 4086	0 / 4086	0 / 4086	0 / 4086	0 / 4086



Why is this useful?



Objective: can we build tools to help refactor?

Several research challenges obstruct

- **Representation and Abstractions:** What is a privilege?
- **Flexible, Scalable, Systematic Policy Synthesis:**
 - What are the subjects? objects? operations?
 - How to get as non-expert?
 - How to protect all resources not just known critical?
- **Security:** How to measure the security of the system?

Today

Program Capability Graph

Object-Encapsulation Model: automated program reasoning through object-oriented grouping

An end-to-end compiler for automated privilege and compartmentalization analysis

Fully automated program partitioning and privilege reduction measurement with ****Inductive Approach****

Security measurement using deductive sensitive and suspicious type labels

Key results: 50% private, 75% for linux, syscall exposure very low, automated library analysis, privilege locality arises in both good and bad cases

FINI!



RICE

**TRAIL
OF
BITS**

Appendix


Nginx External RO/RW Global Variable Rank

#	Encapsulated-Object	Ext. R/O ↑
1	os/unix/nginx_process_cycle.c	47
2	event/nginx_event_timer.h	43
3	event/nginx_event.c	41
4	http/nginx_http_core_module.c	40
5	http/nginx_http_upstream.c	39
6	http/nginx_http_request.c	38
7	core/nginx.c	30
8	event/modules/nginx_epoll_module.c	27
9	http/modules/nginx_http_proxy_module.c	27
10	http/nginx_http_variables.c	25

#	Encapsulated-Object	Ext. R/W ↑
1	os/unix/nginx_process_cycle.c	36
2	core/nginx_times.c	30
3	core/nginx.c	26
4	os/unix/nginx_process.c	20
5	event/nginx_event.c	17
6	event/modules/nginx_epoll_module.c	16
7	core/nginx_regex.c	9
8	core/nginx_cycle.c	7
9	event/nginx_event_accept.c	6
10	http/modules/nginx_http_ssi_filter_module.c	6



External Access Ratio Metric (Cont.)

Private Objs. Public Objs.





External Access Distribution per Object

Code Objects



Data Objects



PDF

CDF

Read / Write Relation for LLVM Instructions and Operations

Instruction	Read-Only / Writable
LoadInst	RO
StoreInst	RW
GetElementPtrInst	Depends on forward uses
CallInst	RO
CallBrInst	RO
ICmpInst	RO
AtomicRMWInst	RW
AtomicCmpXchgInst	RW
ReturnInst	RO
SelectInst	RO
PtrToIntInst	RO
PtrToIntOpr	RO
GEPOperator	Depends on forward uses
BitCastOperator	RO
PHINode	RO
Constant	N/A

Evaluated Programs and Libraries & Generation Time

Parsing Libraries	Software System
libroxml	Nginx
jansson	CPython
facil.io	PHP
libxml2	MuPDF
json-c	gimp

Target Program	Generation Runtime
libroxml	0.854s
jansson	0.824s
facil.io	8.290s
libxml2	2m58.852s
json-c	0.884s
Nginx	6m3.503s
CPython	119m24.766s
PHP	180m44.243s
MuPDF	60m43.136s
gimp	1m26.717s