AlphaEXP: An Expert System for Identifying Security-Sensitive Kernel Objects

Ruipeng Wang, Kaixiang Chen, Chao Zhang, Zulie Pan
Qianyu Li, Siliang Qin, Shenglin Xu, Min Zhang, Yang Li

Motivation

Memory corruption vulnerabilities are one of the major threats to software.

[1]

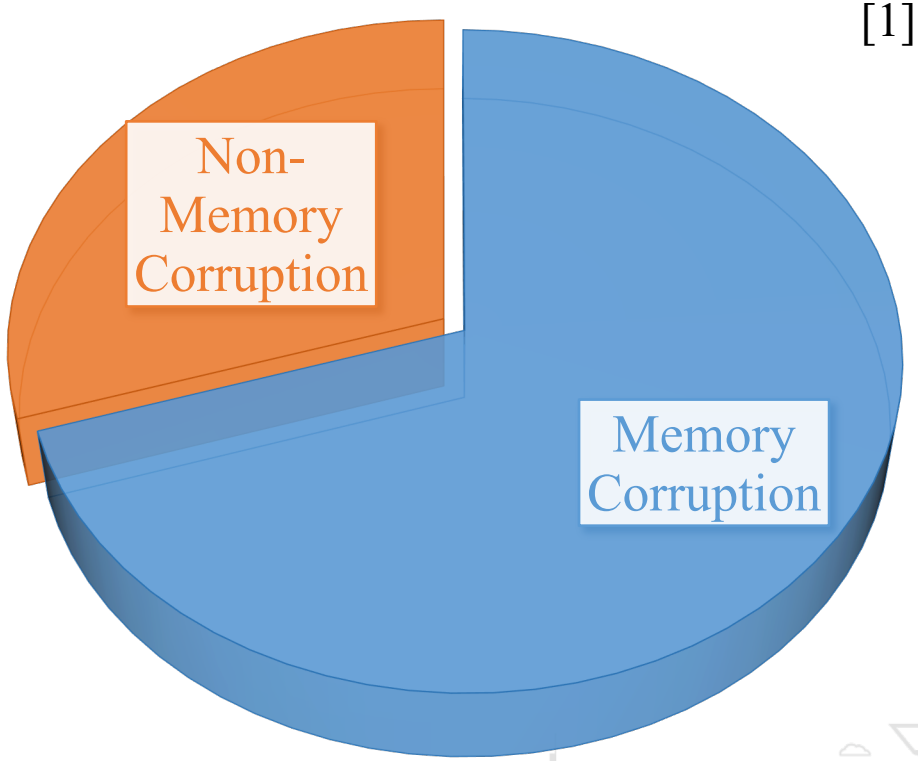

APPLICATION SECURITY

MITRE Updates CWE Top 25 Most Dangerous Software Weaknesses

Use-after-free and OS command injection vulnerabilities reach the top five most dangerous software weaknesses in the 2023 CWE Top 25 list.

SECURITY BUGS CISA

Buffer overflow-type memory bugs remain the most dangerous vulnerabilities out there

Source: Google Project Zero, 0day "In the Wild" spreadsheet. Last updated: 2023-04-20

Motivation

There are three types of solutions proposed and deployed in practice：

1. Vulnerability patching

- It cannot mitigate unknown 0-day vulnerabilities

2. Software and system hardening

- Such solutions would introduce performance costs to the system

3. Object-specific protections

- Object-specific protection has a good balance between security and performance.

Motivation

There are three types of solutions proposed and deployed in practice：

1. Vulnerability patching

   - It cannot mitigate unknown 0-day vulnerabilities

How to identify sensitive objects that need to be protected?

2. Software and system hardening

   - Such solutions would introduce performance costs to the system

3. Object-specific protections

   - Object-specific protection has a good balance between security and performance.

Motivation

How to identify sensitive objects that need to be protected?

◆ **Analyzing publicly exposed exploits to find out objects that are abused**

this solution heavily relies on the human experience, and cannot find sensitive data that have not been abused yet

◆ **Classifying objects based on developers' intentions and the program's semantics**

Its results (i.e., sensitive objects) may deviate from the adversary's

◆ **Analyzing the target code following specific attack knowledge**

SLAKE （CCS '19）ELOISE（CCS '20）
they are not generic solutions for identifying sensitive objects, and cannot distinguish the sensitivity of the objects.

Motivation

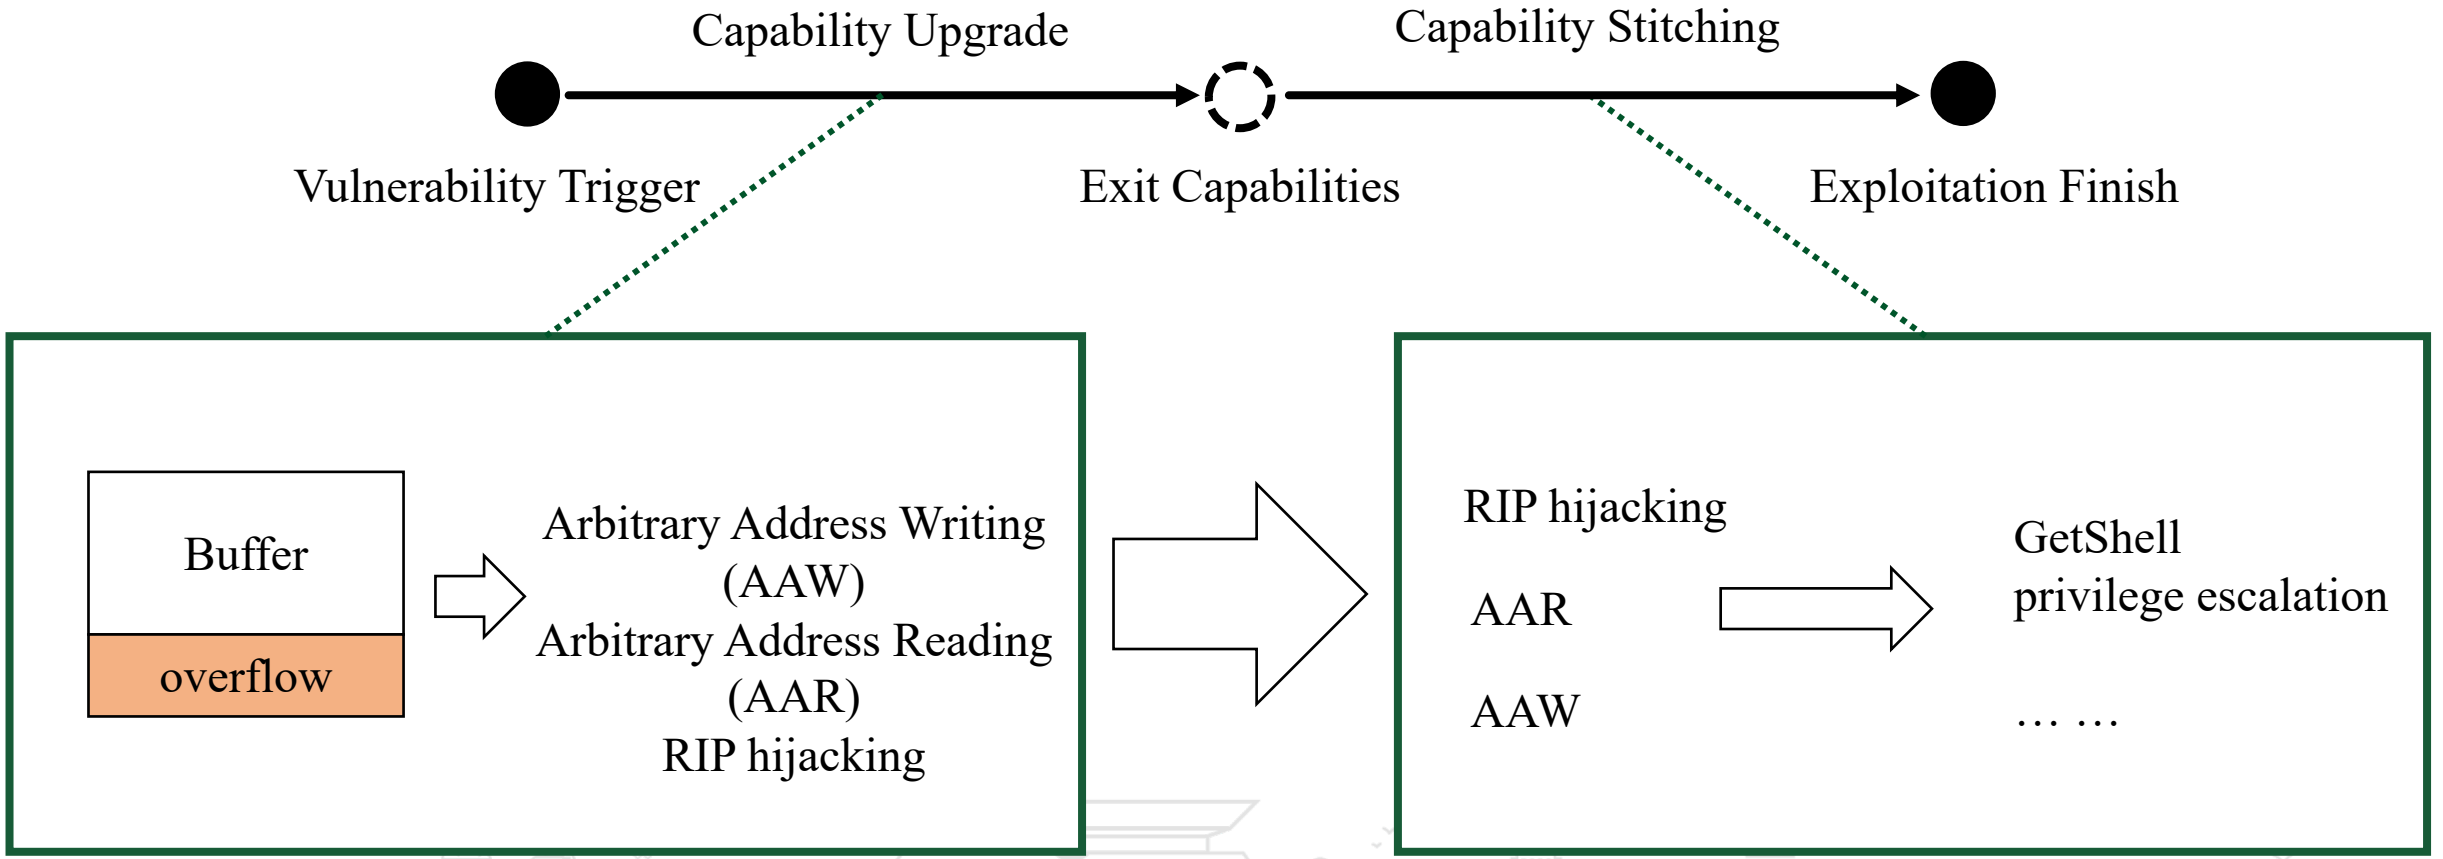
Capability Upgrade
Capability Stitching
Vulnerability Trigger
Exit Capabilities
Exploitation Finish
Buffer
overflow
Arbitrary Address Writing (AAW)
Arbitrary Address Reading (AAR)
RIP hijacking
RIP hijacking
AAR
AAW
GetShell
privilege escalation
… …

Motivation

(a) Vulnerable Function

```
1  static bool tipc_crypto_key_rcv(struct tipc_crypto *rx, struct tipc_msg *hdr)
2  {
3     ...
4     struct tipc_aead_key *skey = NULL;
5     ...
6     u16 size = msg_data_sz(hdr);
7     u8 *data = msg_data(hdr);
8     ...
9     skey = kmalloc(size, GFP_ATOMIC);
10    ...
11    skey->keylen = ntohl(*((__be32 *)(data + TIPC_AEAD_ALG_NAME)));
12    memcpy(skey->alg_name, data, TIPC_AEAD_ALG_NAME);
13    memcpy(skey->key, data + TIPC_AEAD_ALG_NAME + sizeof(__be32), skey-
14 >keylen);
15    ...
```

(b) Vulnerability

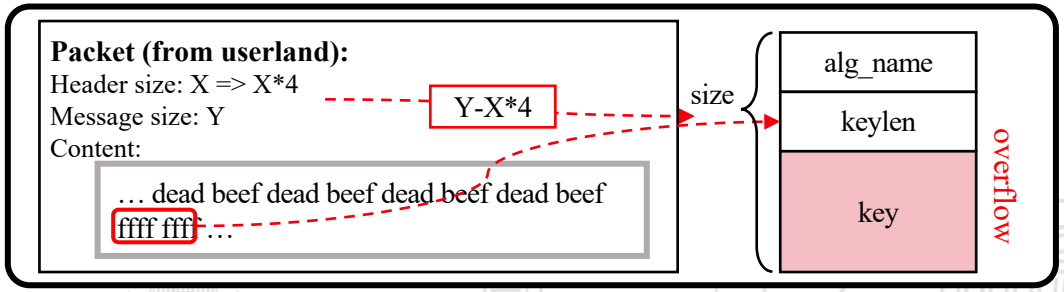

(c) Structure Definition

```
struct tipc_aead_key {
  char alg_name[TIPC_AEAD_ALG_NAME];
  unsigned int keylen;
  char key[ ];
};

struct msg_msg {
  struct list_head m_list;
  long m_type;
  size_t m_ts;
  struct msg_msgseg *next;
  void * security;
};

struct tty_struct {
  int magic;
  struct kref kref;
  struct device *dev;
  struct tty_driver *driver;
  const struct tty_operations *ops;
  int index;
  ...
};
```

(d) Exploitation

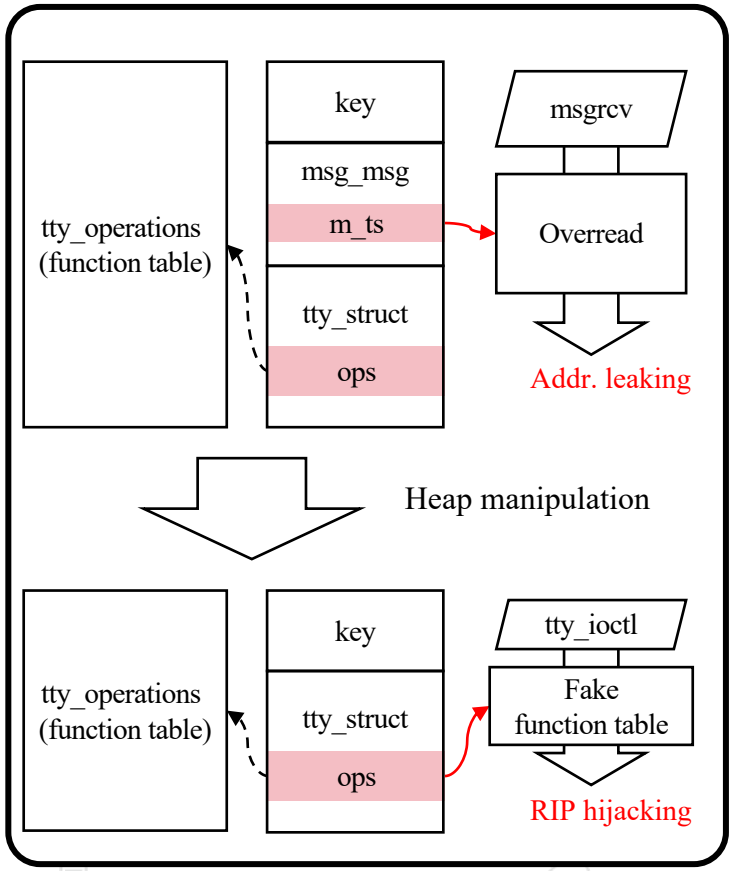

Motivation

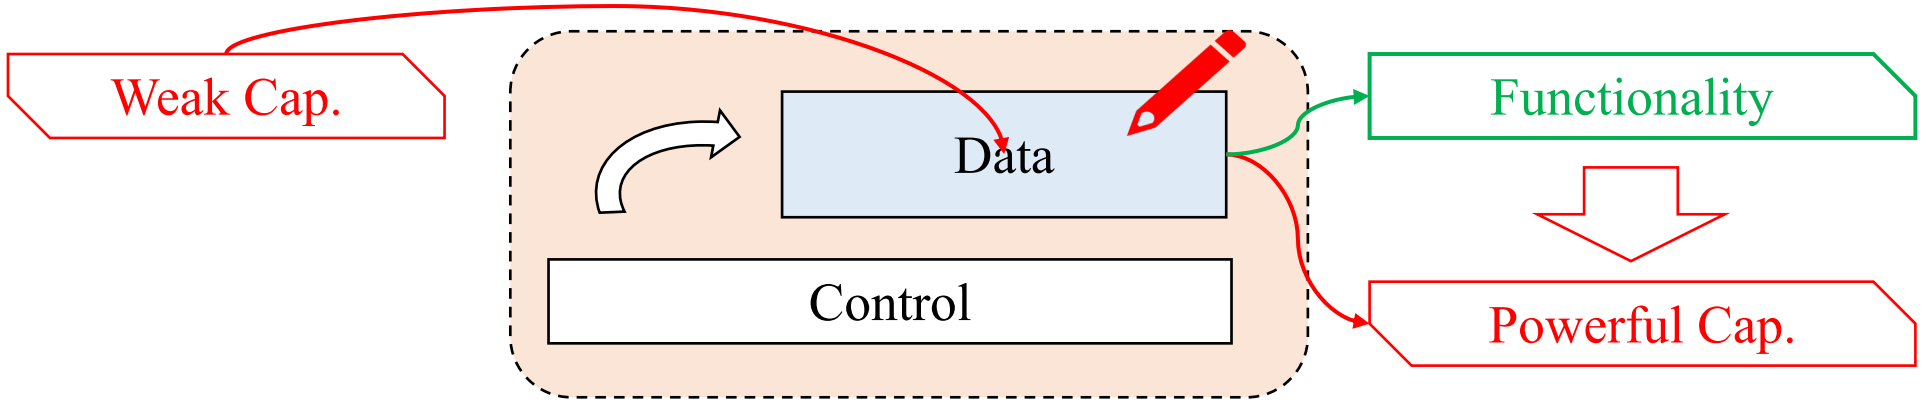


```
struct tty_struct {
  int magic;
  struct kref kref;
  struct device *dev;
  struct tty_driver *driver;
  const struct tty_operations *ops;
  int index;
  ...
};
```

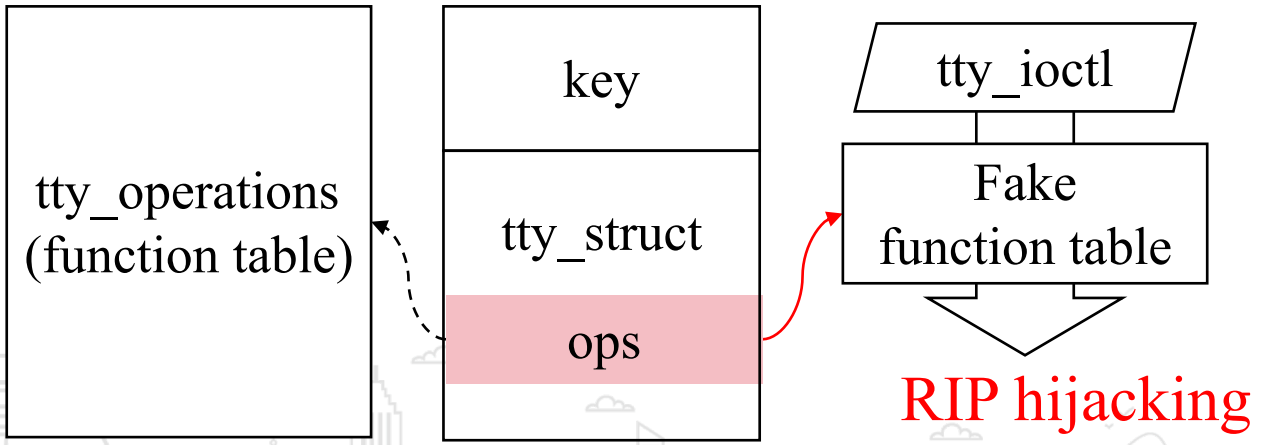

Motivation

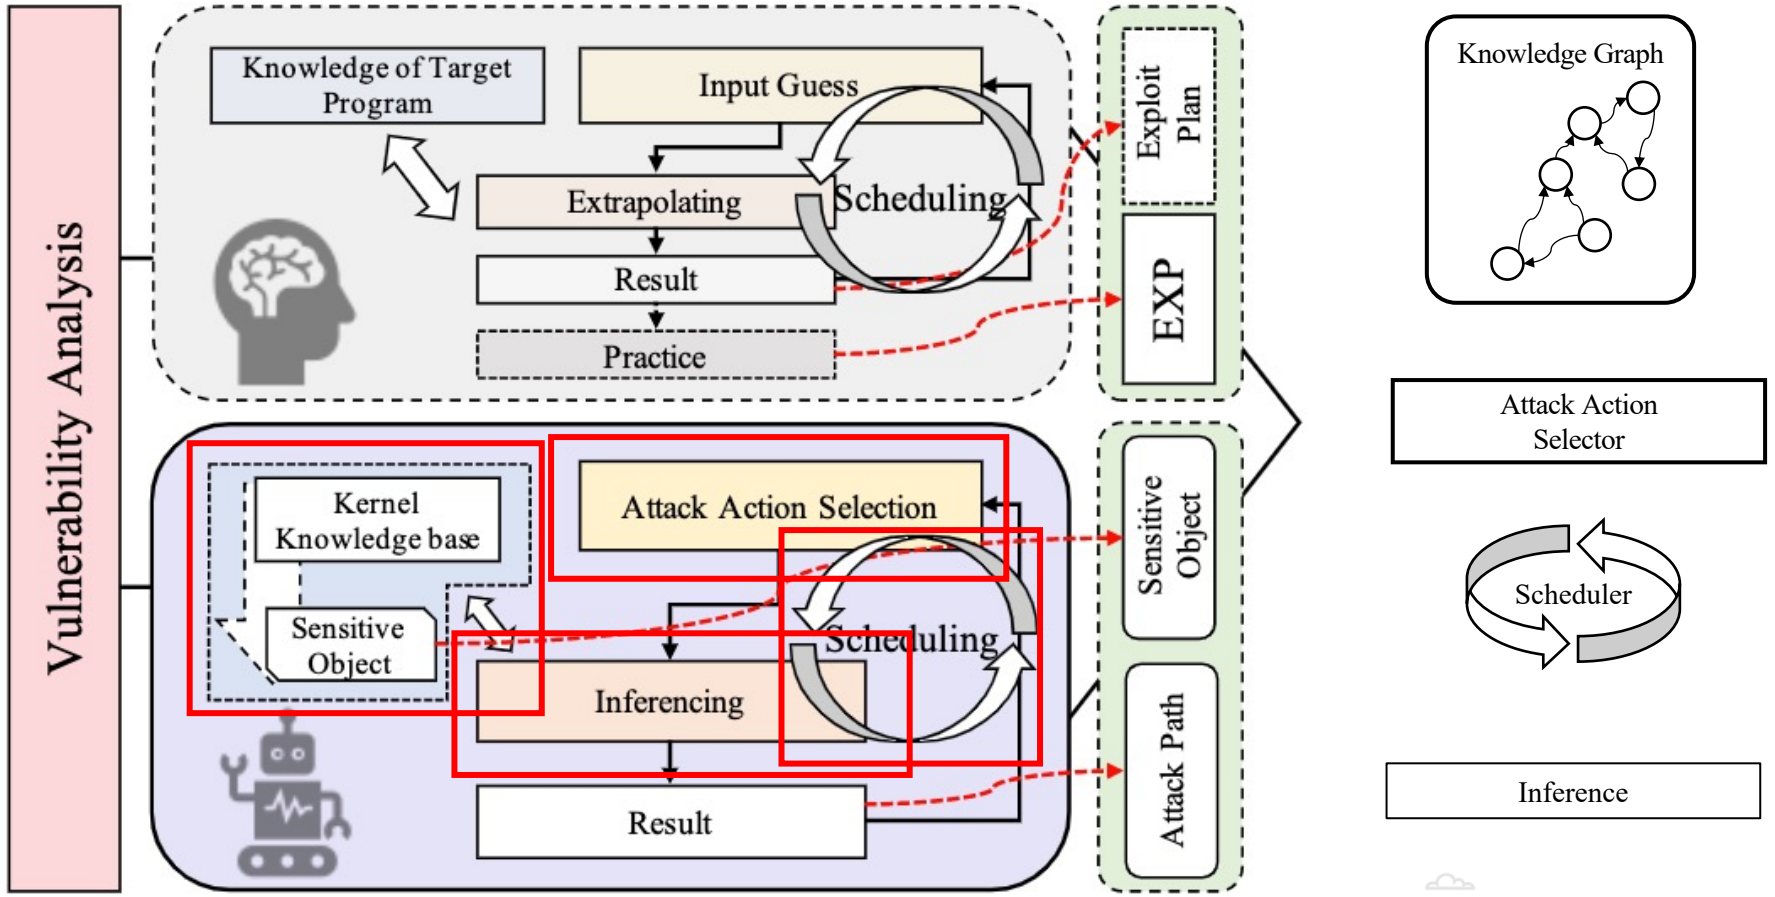
Vulnerability Analysis
Knowledge of Target Program
Input Guess
Extrapolating
Scheduling
Result
Practice
Exploit Plan
EXP
Kernel Knowledge base
Sensitive Object
Attack Action Selection
Scheduling
Inferencing
Result
Sensitive Object
Attack Path
Knowledge Graph
Attack Action Selector
Scheduler
Inference

Our Solution: AlphaEXP

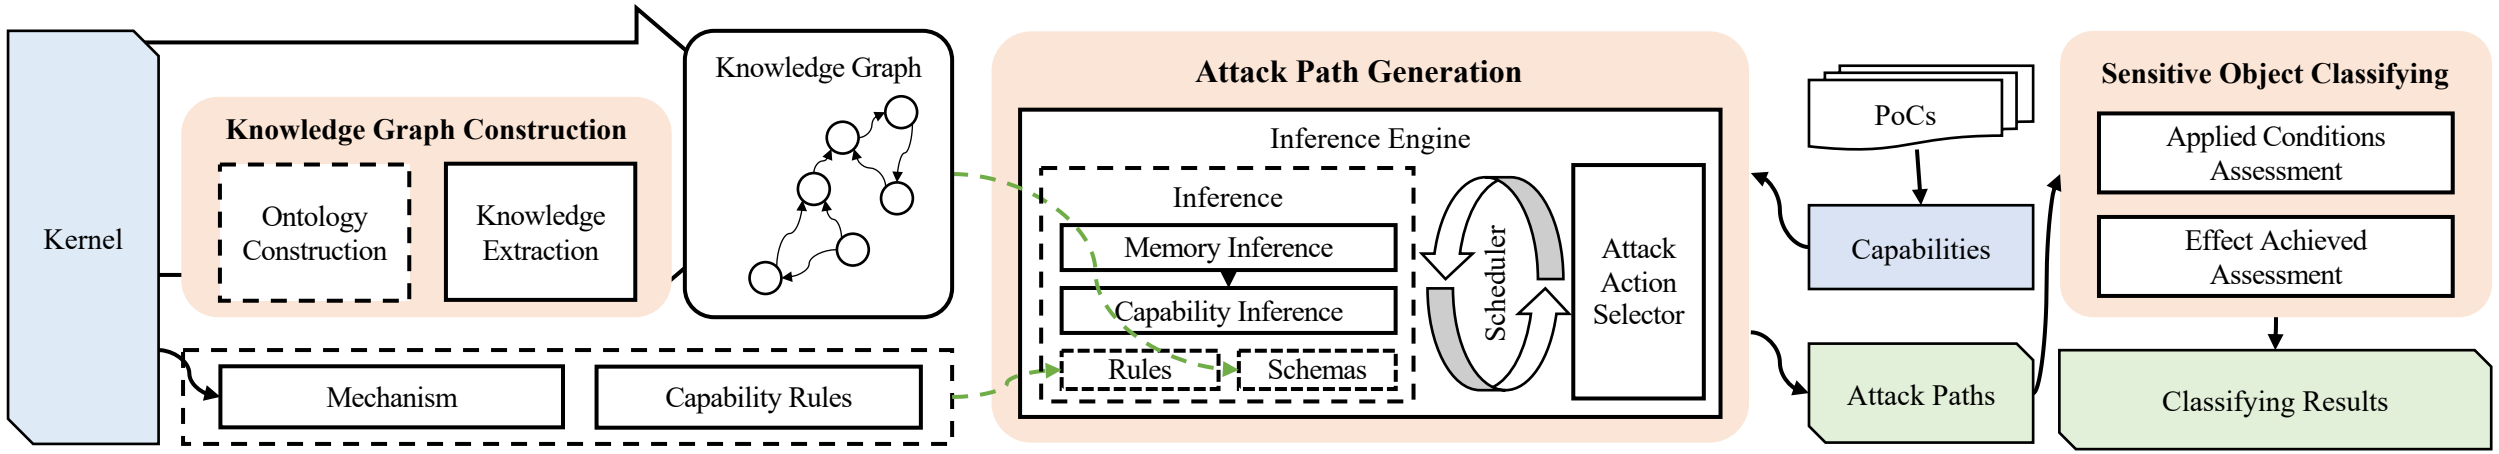

Our Solution: AlphaEXP

Knowledge Graph Construction

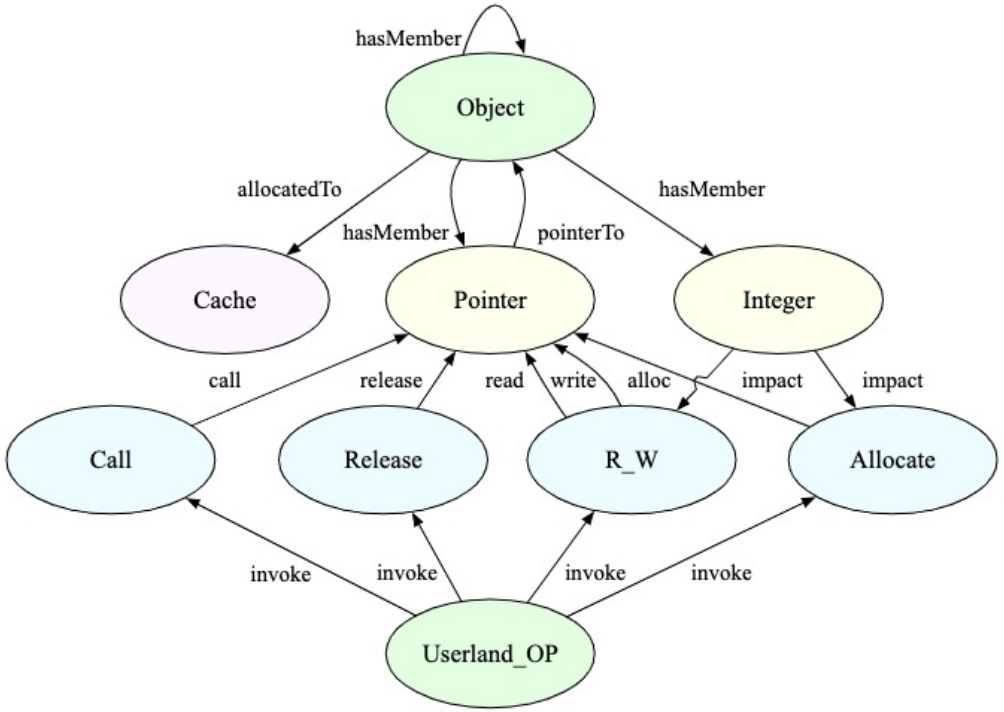

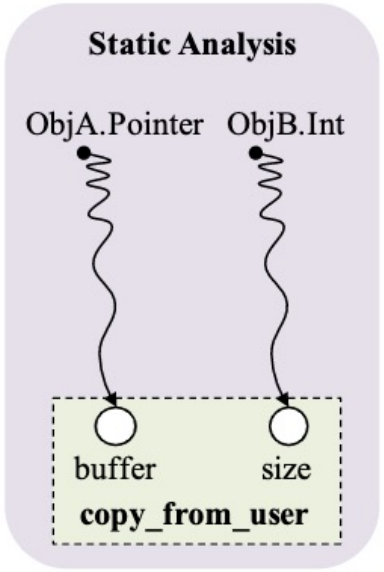

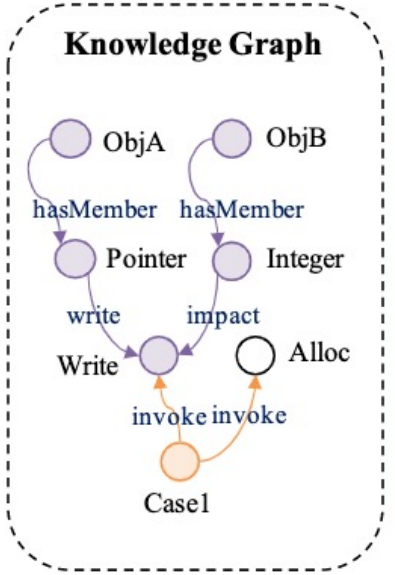

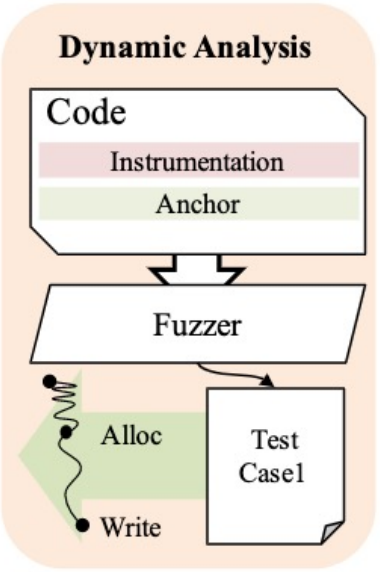

Our Solution: AlphaEXP

Attack Path Generation

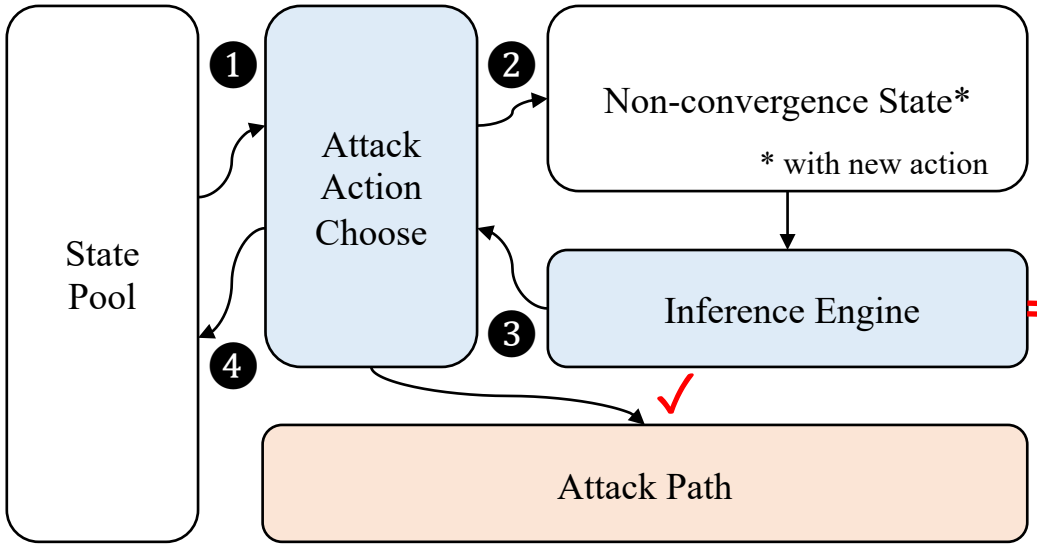

1 Choose an state.
2 Choose an attack action.
3 Inference attack action effect.
4 Adding the converged state to state pool.
✓ If the attack action is as expected, add it to the attack path.

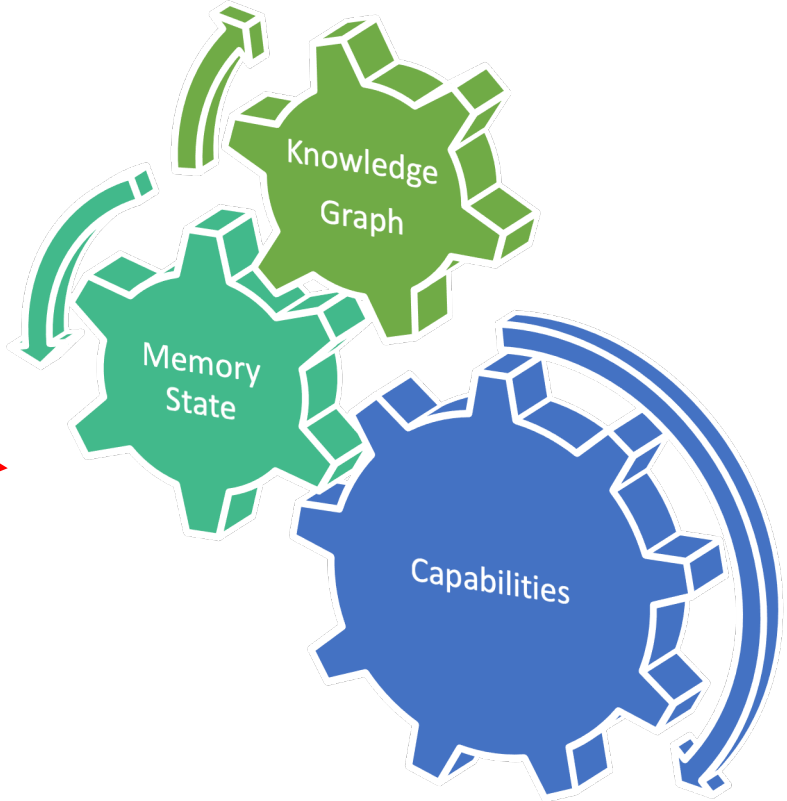

Inference Rules based on Datalog

Our Solution: AlphaEXP

Sensitive Object Classifying

Perspective	Factor	Description
Applied Conditions	kmem-cache	Can be applied in the exploitation of different vulnerability object memory kmem-cache.
	entry capability	Modification of sensitive object requires unintended writing capability over 0x80 size
	vulnerability type	Sensitive object can both be applied in the exploitation of overflows and UAF
Effect Achieved	writing capability	Sensitive object can be used to upgrade writing capability in exploitation
	executing capability	Sensitive object can be used to upgrade executing capability in exploitation
	reading capability	Sensitive object can be used to upgrade reading capability in exploitation

Evaluation

- RQ1: How effective is AlphaEXP in sensitive objects identifying and classifying?
- RQ2: Is AlphaEXP better at identifying sensitive objects compared to current SOTA techniques?
- RQ3: What is the cost of building a knowledge graph?
- RQ4: How effective is attack path generation?

Evaluation

- RQ1: How effective is AlphaEXP in sensitive objects identifying and classifying?
- RQ2: Is AlphaEXP better at identifying sensitive objects compared to current SOTA techniques?

50 objects that could be abused to get writing capability, 81 objects with reading capability, and 112 objects with execution capability

	Sensitive Objects
Write	keyctl_update_key✦, msg_msg✦, add_key✦, ip_options_get_from_user✦, scsi_request, hiddev_ioctl_usage, proc_ioctl, kexec_segment, do_ipv6_setsockopt snd_info_buffer, xt_table_info, gss_pipe_downcall, snd_ctl_elem_info, vt_do_kdgkb_ioctl, drm_ioctl, proc_bulk, create_entry, move_addr_to_kernel elf_prpsinfo, simple_transaction_argresp, simple_attr_write, proc_do_submiturb, drm_syncobj_array_find, fb_sys_write, fb_write, drm_mode_dirtyfb_ioctl usblp_write, drm_crtc, drm_property_blob, drm_i915_gem_object, tty_struct, drm_syncobj_array_wait_timeout,kernfs_fop_write, kexec_segment snd_ctl_elem_id, map_lookup_elem, map_update_elem, do_semtimedop, compat_agpioc_reserve_wrap, sk_buff, sendmsg, sock_filter, setxattr __get_filter, ethtool_set_eeprom, rawv6_seticmpfilter, ipv6_txoptions, agpioc_reserve_wrap, memfd_create✲, drm_syncobj_timeline_signal_ioctl✲ raw_seticmpfilter✱, cpumask✱
Read	ipv6_opt_hdr★, sock_fprog_kern★, policy_load_memory★, ldt_struct★, ip_options★, seq_file★, xfrm_policy★, xfrm_algo_aead★, xfrm_algo★ cfg80211_pkt_pattern★, user_key_payload★, xfrm_replay_state_esn★, ip_sf_socklist★, proc_dir_entry★, ext4_dir_entry_2★, station_info★ cache_reader★, tc_cookie★, cfg80211_bss_ies★, sg_header★, inotify_event_info★, audit_rule_data★, fb_info★, cfg80211_sched_scan_request★ fb_cmap_user★, cache_request★, fname★, ieee80211_mgd_auth_data★, mon_reader_bin★, mon_reader_text★, cfg80211_scan_request★ tcp_fastopen_context★, request_key_auth★, xfrm_algo_auth★, cfg80211_wowlan_tcp★, msg_msg★, tcp_sock✰, user_element, neighbour pneigh_entry, net_device, netdev_phys_item_id, netlink_ext_ack, cfg80211_nan_match_params, wiphy, wiphy_iftype_ext_capab, wireless_dev hidraw_report, hid_device, sg_request, fb_cmap, usb_device, urb, usblp, drm_crtc, drm_plane, cfg80211_connect_resp_params, kobj_uevent_env beacon_data, probe_resp, cfg80211_roam_info, cfg80211_wowlan_wakeup, cfg80211_ssid, cfg80211_mgmt_tx_params, ieee80211_mgd_assoc_data cfg80211_ft_event_params, fat_ioctl_filldir_callback, key_params, drm_property_blob, rpc_pipe_msg, geneve_opt, tcp_fastopen_cookie __kfifo, seq_buf, rchan_buf, drm_master, cfg80211_pmsr_ftm_result✲, cfg80211_update_owe_info✲, cfg80211_fils_resp_params✲, sg_scsi_ioctl✱
Exec	seq_operations✦, perf_event_context✦, linux_binprm✦, vmap_area✦, kioctx_table✦, snd_seq_timer✦, sock✦, tty_ldisc✦, tty_struct✦, seq_file✦ sk_security_struct✦, assoc_array_edit✦, cgroup_namespace✦, ext4_allocation_context✦, tty_file_private✦, file✦, subprocess_info✦, ccid✦, timerfd_ctx✦ ip_options✦, ip_mc_list✦, ip_sf_socklist✦, request_key_auth✦, pid_namespace✦, k_itimer✦, avc_node✦, kioctx✦, key✦ ✰, ip_mc_socklist✦ ✰,packet_sock✰ fsnotify_group✰, blk_plug_cb, blk_stat_callback, snd_timer, nfs_io_completion, hci_dev, drm_i915_gem_object, pipe_buffer, vga_device snd_pcm_runtime, snd_pcm, snd_hwdep, snd_kcontrol, link_master, snd_kctl_ioctl, snd_timer_instance, scsi_cmnd, clk_fractional_divider, fbcon_ops , snd_pcm_hw_rule, snd_seq_device, snd_info_entry, snd_card, snd_jack, net_device, acpi_cpufreq_data, hid_device, sony_sc, ahci_host_priv, udp_sock snd_seq_client_port, hda_jack_callback, kprobe, hashtab, shm_file_data, iommu_group, loop_device, input_polled_dev, sched_domain_topology_level crypto_ahash, ahash_request_priv, crypto_tfm, skcipher_instance, akcipher_instance, journal_s, input_dev_poller, serio, ml_device, alps_data, psmouse crypto_skcipher, aead_instance, crypto_acomp, nf_conntrack_expect, ubuf_info, proc_inode, proc_dir_entry, fib6_walker, aio_kiocb, dio, simple_attr, rtnl_link inet_connection_sock, kthread_create_info, async_entry, ring_buffer, filter_pred, nfs_renamedata, nfs_server, nfs_pgio_header, nfs_commit_data, hda_codec, tracer rpc_task, rpc_rqst, flow_block_cb✲, flow_indr_block_cb✲, tcf_filter_chain_list_item✲, io_wq✲, context_barrier_task✲, execute_cb✲
✦: Identified by SLAKE as well, ★: Identified by ELOISE as well, ✰: Identified by KOOBE as well, ✲: Not present in v4.15, ✱: False Positives	

Evaluation

- RQ1: How effective is AlphaEXP in sensitive objects identifying and classifying?

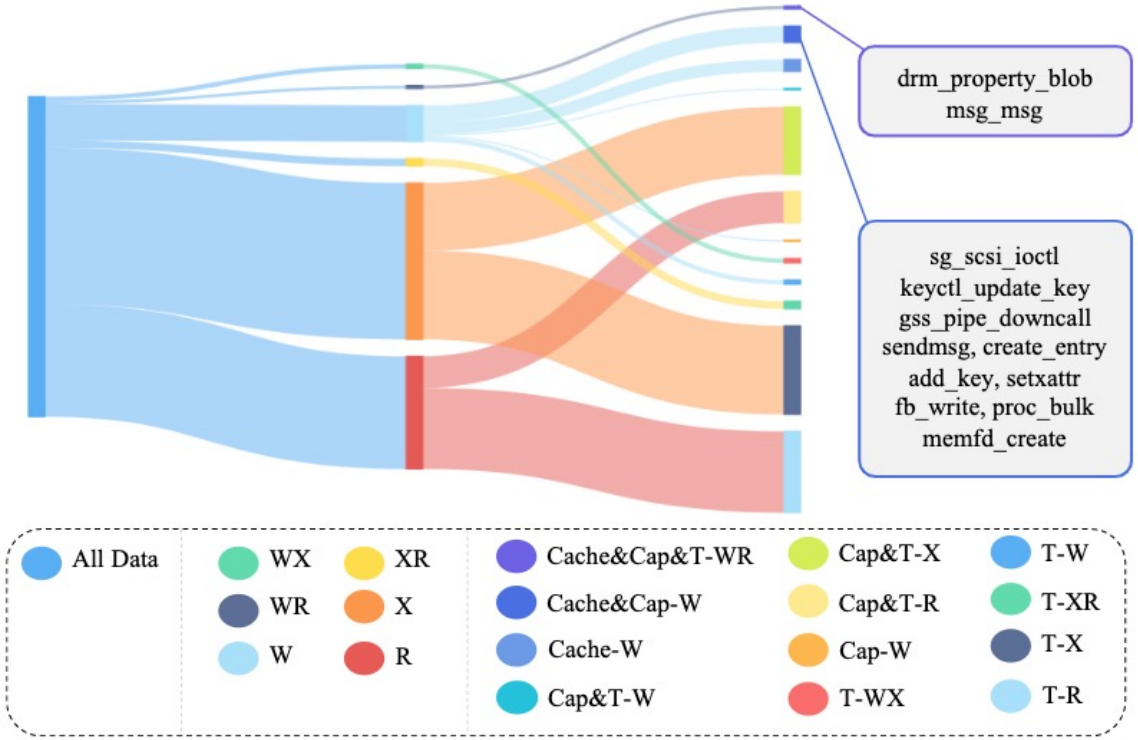

Cache&Cap&T-WR	drm_property_blob, msg_msg
Cache&Cap-W	sg_scsi_ioctl, keyctl_update_key, sendmsg, gss_pipe_downcall, create_entry, add_key, setxattr, fd_write, proc_bulk, memfd_create
Cache-W	xt_table_info, simple_attr_write, proc_do_submiturb, usblp_write, ip_options_get_from_user, ipv6_txoptions, do_ipv6_setsockopt, sk_buff
Cap&T-W	do_semtimedop
Cap&T-X	seq_operations, assoc_array_edit, cgroup_namespace, ext4_allocation_context, ip_options_rcu, ip_sf_socklist, pid_namespace, avc_node, tty_ldisc tty_file_private, file, ccid, blk_plug_cb, snd_timer_instance, link_master, snd_info_entry, hda_jack_callback, hashtab, shm_file_data, crypto_ahash ahash_request_priv, crypto_tfm, skcipher_instance, akcipher_instance, crypto_skcipher, aead_instance, crypto_acomp, ubuf_info, flow_block_cb, rpc_task kthread_create_info, tracer, nfs_io_completion, io_wq, simple_attr, input_polled_dev, input_dev_poller, acpi_cpufreq_data, clk_fractional_divider fbcon_ops, pipe_buffer, ip_mc_socklist
Cap&T-R	user_element, request_key_auth, user_key_payload, seq_buf, pneigh_entry, netdev_phys_item_id, tc_cookie, cfg80211_nan_match_params wiphy_iftype_ext_capab, cfg80211_connect_resp_params, cfg80211_fils_resp_params, cfg80211_roam_info, cfg80211_ssid, cfg80211_update_owe_info key_params, cfg80211_pkt_pattern, cache_request, tcp_fastopen_cookie, beacon_data, fat_ioctl_filldir_callback, hidraw_report
Cap-W	snd_info_buffer
T-WX	tty_struct, ip_options, drm_i915_gem_object
T-W	rawv6_seticmpfilter, kernfs_fop_write, fb_sys_write
T-XR	seq_file,ip_sf_socklist, net_device, hid_device
T-X	perf_event_context, linux_binprm, vmap_area, kioctx_table, kioctx, ip_mc_list, k_itimer, sk_security_struct, snd_seq_timer, timerfd_ctx, subprocess_info key, sock, blk_stat_callback, snd_timer, snd_pcm_runtime, snd_pcm, snd_hwdep, snd_kcontrol, snd_kctl_ioctl, snd_pcm_hw_rule, snd_seq_device, snd_card snd_jack, snd_seq_client_port, hda_codec, kprobe, nf_conntrack_expect, rtnl_link, flow_indr_block_cb, tcf_filter_chain_list_item, fib6_walker inet_connection_sock, packet_sock, rpc_rqst, hci_dev, udp_sock, sched_domain_topology_level, async_entry, ring_buffer, filter_pred, nfs_renamedata nfs_server, nfs_pgio_header, nfs_commit_data, proc_inode, proc_dir_entry, aio_kiocb, dio, journal_s, serio, ml_device, alps_data, psmouse, iommu_group loop_device, sony_sc, ahci_host_priv, scsi_cmnd, nvmem_device, vga_device, context_barrier_task, execute_cb
T-R	ipv6_opt_hdr, sock_fprog_kern, policy_load_memory, ldt_struct, ip_options, xfrm_replay_state_esn, cache_reader, cfg80211_bss_ies, sg_header inotify_event_info, fb_cmap_user, fname, ieee80211_mgd_auth_data, tcp_fastopen_context, xfrm_algo_auth, cfg80211_wowlan_tcp, xfrm_algo xfrm_algo_aead, cfg80211_scan_request, mon_reader_bin, cfg80211_sched_scan_request, mon_reader_text, station_info, ext4_dir_entry_2, xfrm_policy fb_info, audit_rule_data, n_tty_data, proc_dir_entry, kobj_uevent_env, sk_buff, neighbour, netlink_ext_ack, wiphy, wireless_dev, cfg80211_wowlan_wakeup cfg80211_ft_event_params, cfg80211_pmsr_ftm_result, rpc_pipe_msg, geneve_opt, tcp_sock, probe_resp, cfg80211_mgmt_tx_params, ieee80211_mgd_assoc_data rchan_buf, sg_request, fb_cmap, usb_device, urb, usblp, drm_crtc, drm_plane, drm_master, console_font

Evaluation

- RQ3: What is the cost of building a knowledge graph?

The static knowledge extraction process takes **19** minutes

The dynamic knowledge extraction process takes **72** hours

100,723 entities and **180,204** relationships

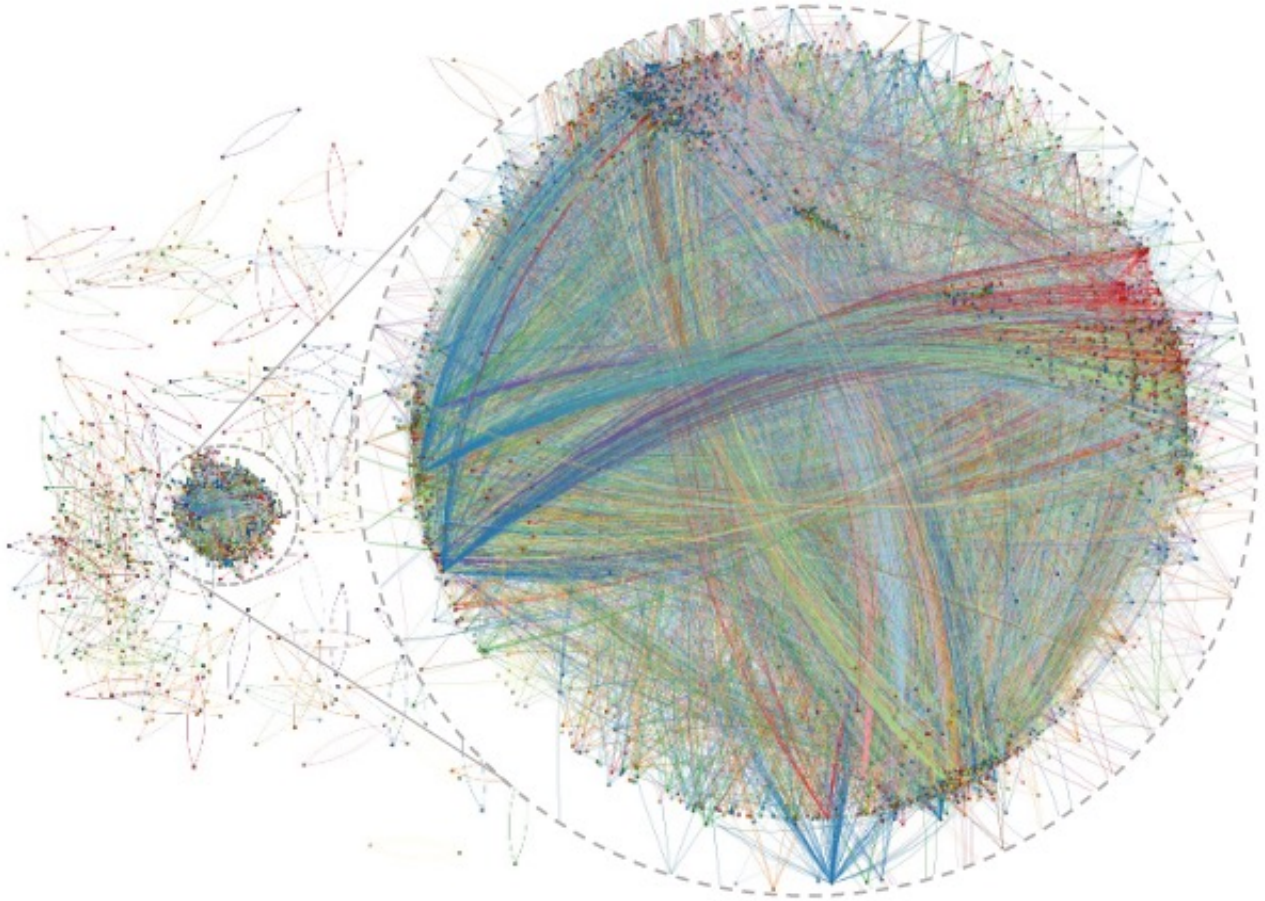

Evaluation

- RQ4: How effective is attack path generation?

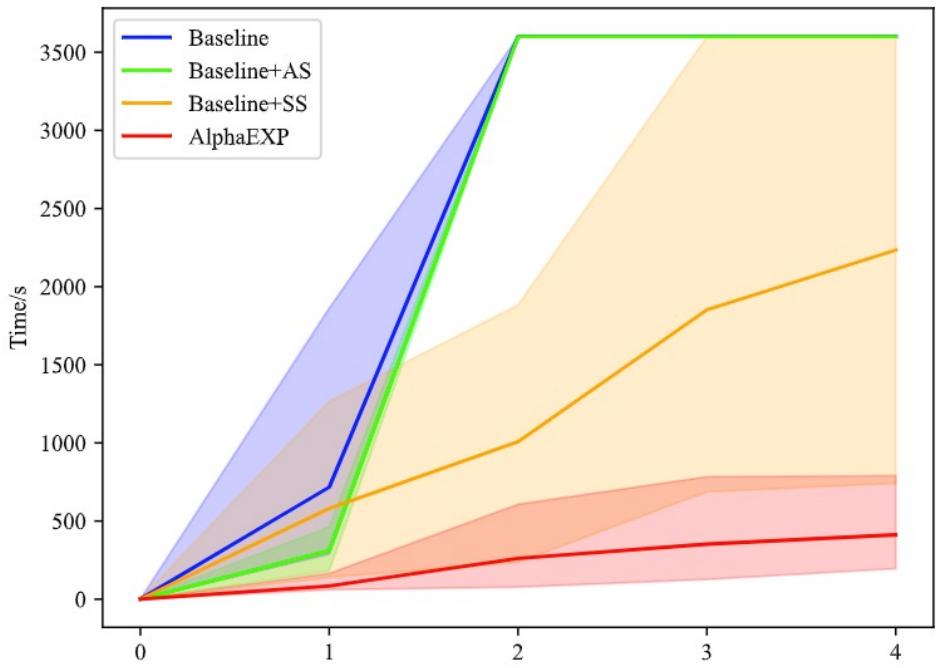

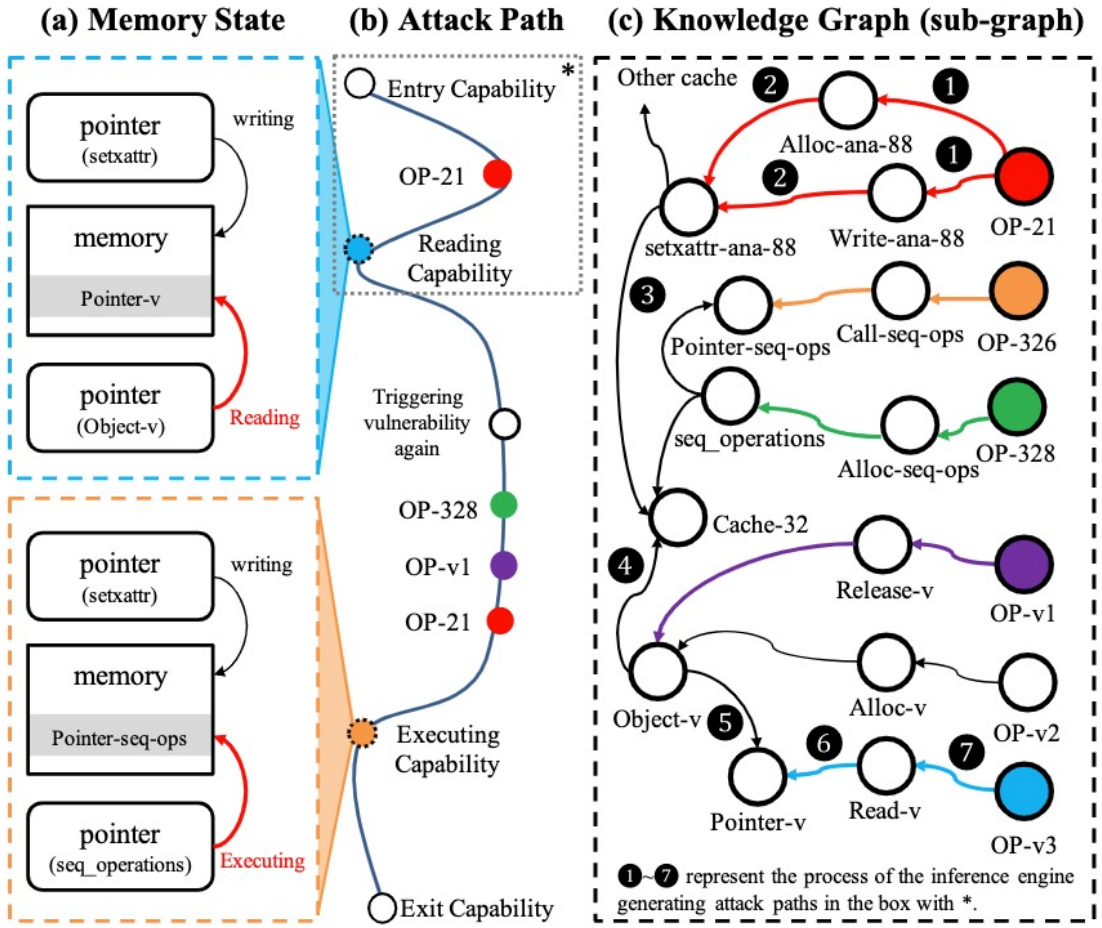

Conclusion

- AlphaEXP can identify sensitive kernel objects and classify their sensitivity, able to help defenders build cost-effective defenses.
- AlphaEXP constructs a knowledge graph of the kernel.
- AlphaEXP reports several hundreds of sensitive kernel objects and classifies them into 12 sensitivity levels.
- AlphaEXP provides a new idea for automated exploit generation (AEG).

Thanks for listening!

Q&A

Contact: Ruipeng Wang, wangruipeng@nudt.edu.cn