3 июня 2025 Москва, LOFT HALL#2

БЕКОН '25

Конференция по БЕзопасности
КОНтейнеров и контейнерных сред

Расширение политик фильтрации трафика в Cilium

Антон Баранов

ГК Астра

Обо мне

БЕКОН

- Ведущий разработчик в команде контейнеризации и оркестрации ГК «Астра»
- 10 лет коммерческой разработки на C++
- Чем занимаюсь сейчас: Kubernetes, eBPF, Cilium, сетевые политики и безопасность контейнеров.
- Люблю глубоко разбираться в технологиях и делиться знаниями.

Антон Баранов

План доклада

БЕКОН

- Предпосылки и цели
- Cilium, eBPF и наши патчи
- Политики Cilium в Kubernetes
- Мониторинг и метрики
- Выводы

БЕКОН

Предпосылки и цели

Предпосылки и цели

БЕКОН

Классификационная метка

Level

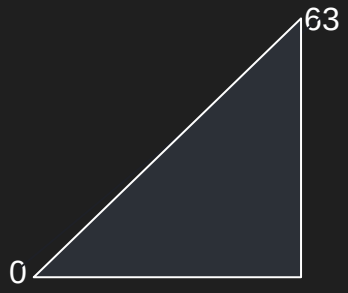

Categories

☐☐☐☐☐☐☐☐
☐☑☐☐☐☐☐☐
☐☐☐☐☐☐☐☐
☐☐☐☐☑☐☐☐
☐☐☐☐☐☐☐☐
☐☑☐☐☐☐☐☐
☐☐☐☐☐☑☐☐
☐☐☐☐☐☐☐☐

Всего 64

Предпосылки и цели

БЕКОН

Пример

Level: 1

Level: 3

Предпосылки и цели

БЕКОН

Пример

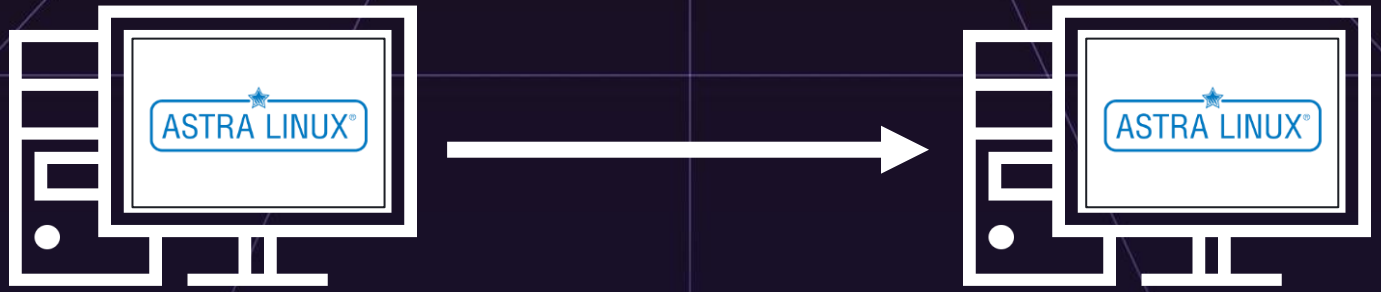

Level: 1

Level: 3

Предпосылки и цели

БЕКОН

Пример

Level: 1

Level: 3

Предпосылки и цели

БЕКОН

В сетевые пакеты IPv4 вставлены метки

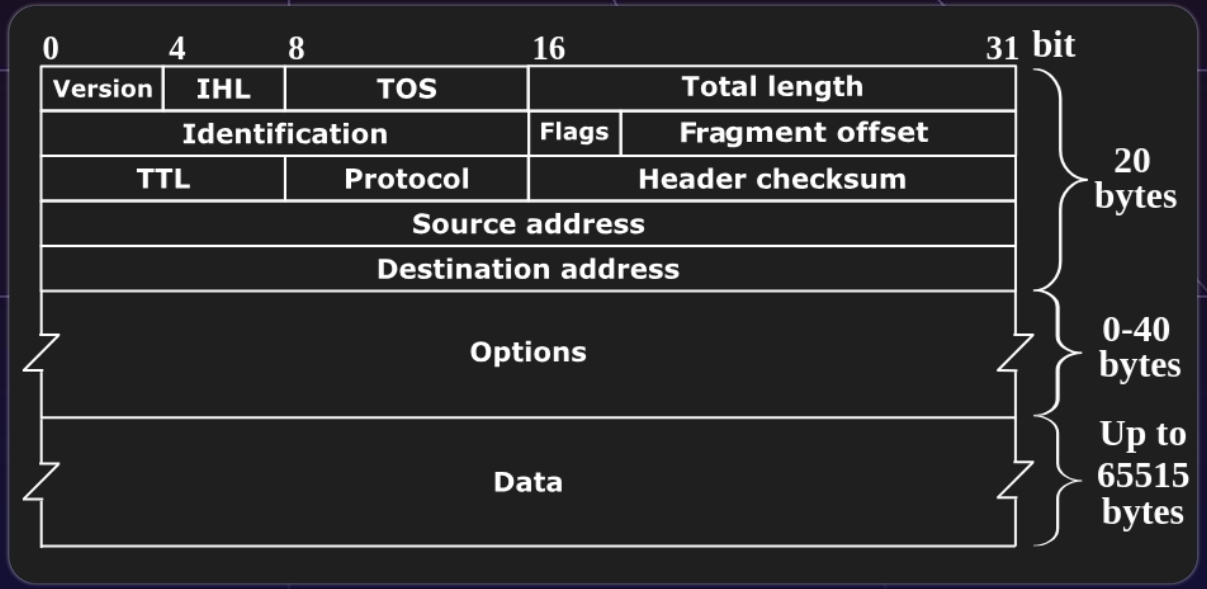

Предпосылки и цели

БЕКОН

В сетевые пакеты IPv4 вставлены метки

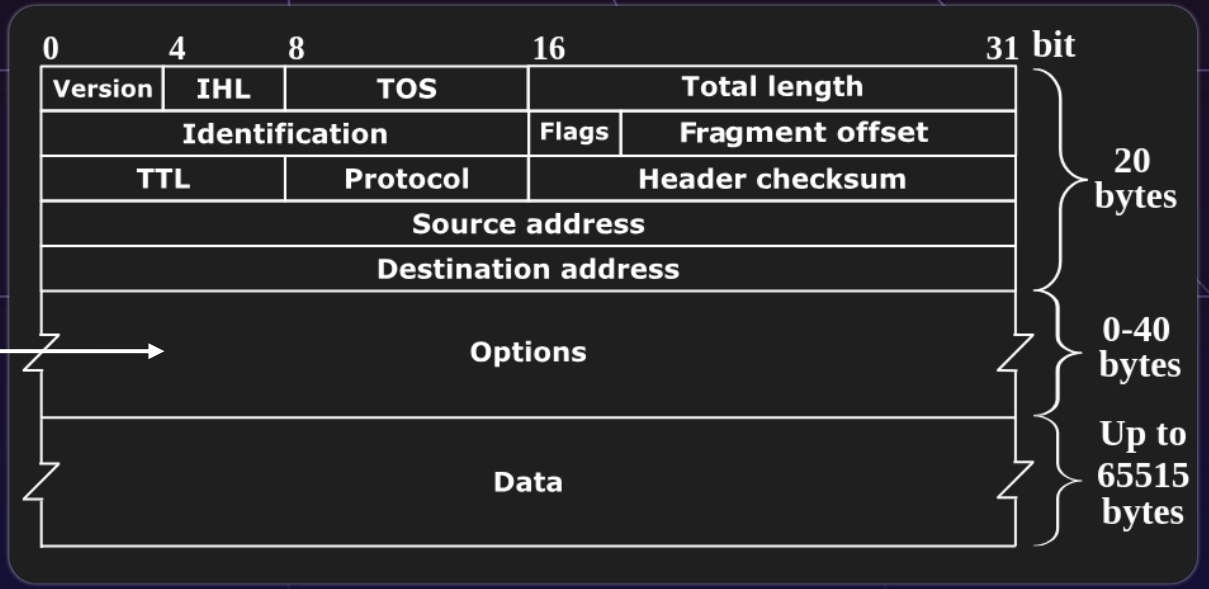

Предпосылки и цели

БЕКОН

В сетевые пакеты IPv4 вставлены метки

Нужны правила,
которые позволят фильтровать пакеты с метками

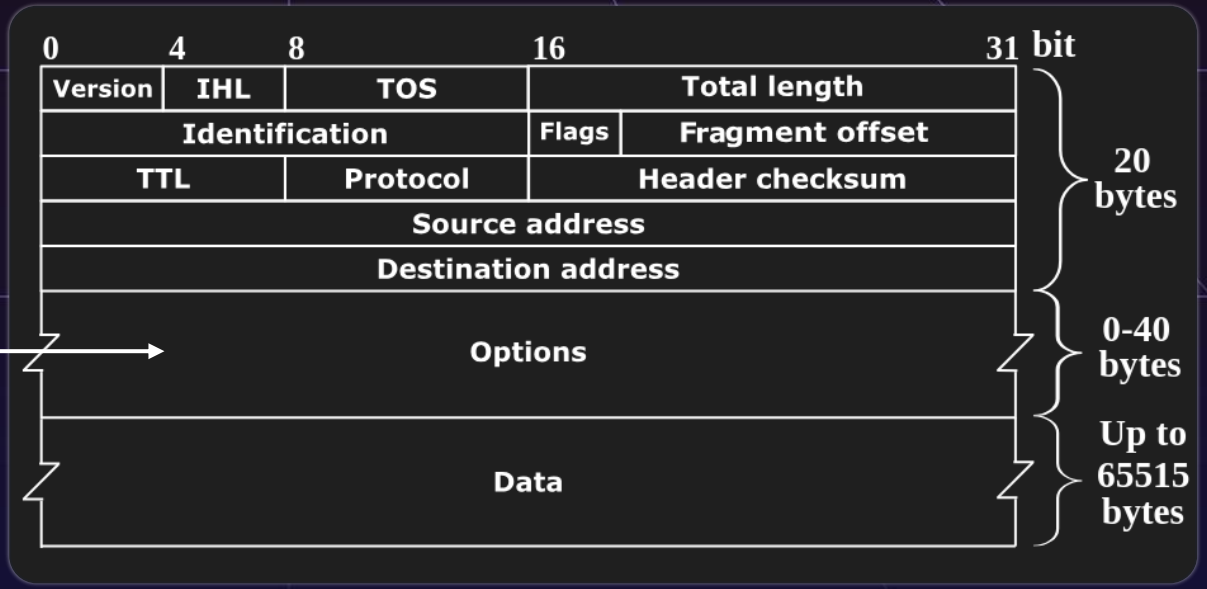

Классификационная метка

БЕКОН

ГОСТ Р 58256- 2018

TYPE (8 бит)	LENGTH (8 бит)	CLASSIFICATION LEVEL (8 бит)	PROTECTION AUTHORITY FLAGS (11+ байт)
10000010 (Dec 130) Security	XXXXXXXX	10101011 Unclassified	AAAAAAA[1|0] AAAAAAA0

Классификационная метка

БЕКОН

ГОСТ Р 58256- 2018

TYPE (8 бит)	LENGTH (8 бит)	CLASSIFICATION LEVEL (8 бит)	PROTECTION AUTHORITY FLAGS (11+ байт)
10000010 (Dec 130) Security	XXXXXXXX	10101011 Unclassified	AAAAAAA[1|0] AAAAAAA0

Общий вид заполнения поля PROTECTION AUTHORITY FLAGS

L - биты значения уровня

C - битовая маска категорий

```
LLLLLLL1 CCCCCCL1 CCCCCCC1 CCCCCCC1 CCCCCCC1 CCCCCCC1
CCCCCCC1 CCCCCCC1 CCCCCCC1 CCCCCCC1 00000CC0
```

Классификационная метка

БЕКОН

ГОСТ Р 58256- 2018

TYPE (8 бит)	LENGTH (8 бит)	CLASSIFICATION LEVEL (8 бит)	PROTECTION AUTHORITY FLAGS (11+ байт)
10000010 (Dec 130) Security	XXXXXXXX	10101011 Unclassified	AAAAAAA[1|0] AAAAAAA0

Общий вид заполнения поля PROTECTION AUTHORITY FLAGS

L - биты значения уровня

C - битовая маска категорий

```
LLLLLLL1 CCCCCCL1 CCCCCCC1 CCCCCCC1 CCCCCCC1 CCCCCCC1
CCCCCCC1 CCCCCCC1 CCCCCCC1 CCCCCCC1 00000CC0
```

Классификационная метка

БЕКОН

ГОСТ Р 58256- 2018

TYPE (8 бит)	LENGTH (8 бит)	CLASSIFICATION LEVEL (8 бит)	PROTECTION AUTHORITY FLAGS (11+ байт)
10000010 (Dec 130) Security	XXXXXXXX	10101011 Unclassified	AAAAAAA[1|0] AAAAAAA0

Общий вид заполнения поля PROTECTION AUTHORITY FLAGS

L - биты значения уровня

C - битовая маска категорий

LLLLLLL1 CCCCCCL1 CCCCCCC1 CCCCCCC1 CCCCCCC1 CCCCCCC1
CCCCCCC1 CCCCCCC1 CCCCCCC1 CCCCCCC1 00000CC0

Классификационная метка

БЕКОН

ГОСТ Р 58256- 2018

TYPE (8 бит)	LENGTH (8 бит)	CLASSIFICATION LEVEL (8 бит)	PROTECTION AUTHORITY FLAGS (11+ байт)
10000010 (Dec 130) Security	XXXXXXXX	10101011 Unclassified	AAAAAAA[1|0] AAAAAAA0

Общий вид заполнения поля PROTECTION AUTHORITY FLAGS

L - биты значения уровня

C - битовая маска категорий

```
LLLLLLL1 CCCCCCL1 CCCCCCC1 CCCCCCC1 CCCCCCC1 CCCCCCC1
CCCCCCC1 CCCCCCC1 CCCCCCC1 CCCCCCC1 00000CC0
```

БЕКОН

Cilium 

Cilium

БЕКОН

- Cilium - CNI plugin Kubernetes
- Cilium использует eBPF

+ Производительность за счет eBPF
+ Мониторинг сетевых взаимодействий
+ Открытый код и комьюнити

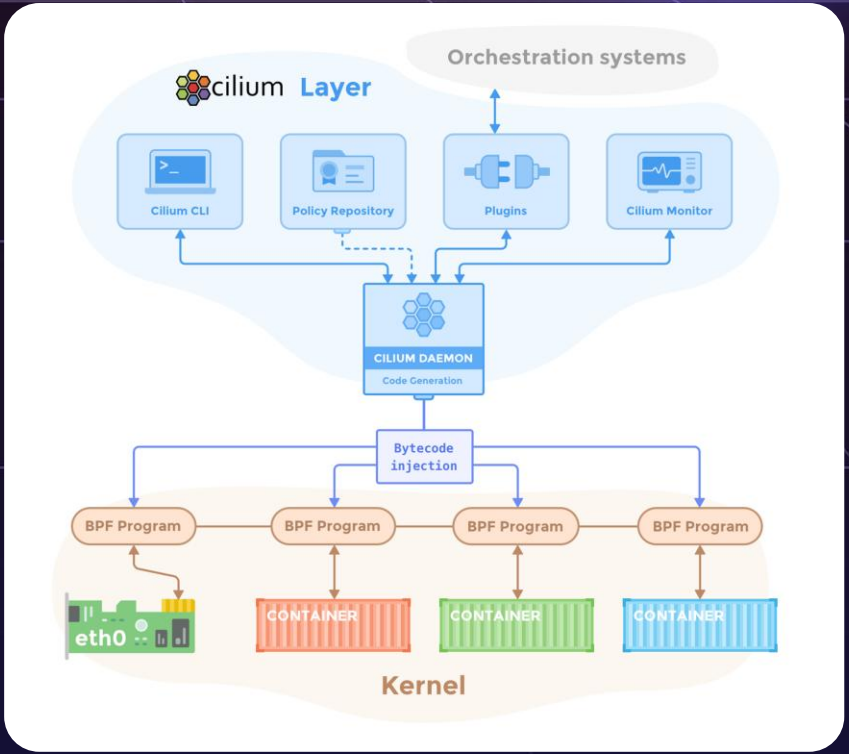

БЕКОН

eBPF

БЕКОН

User Space

Kernel Space

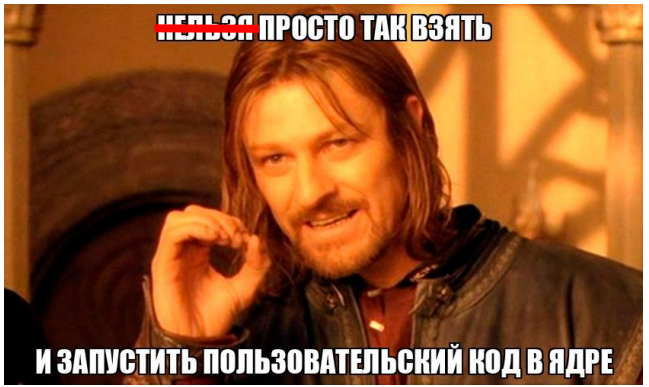

eBPF

БЕКОН

User Space

eBPF code

Kernel Space

eBPF

БЕКОН

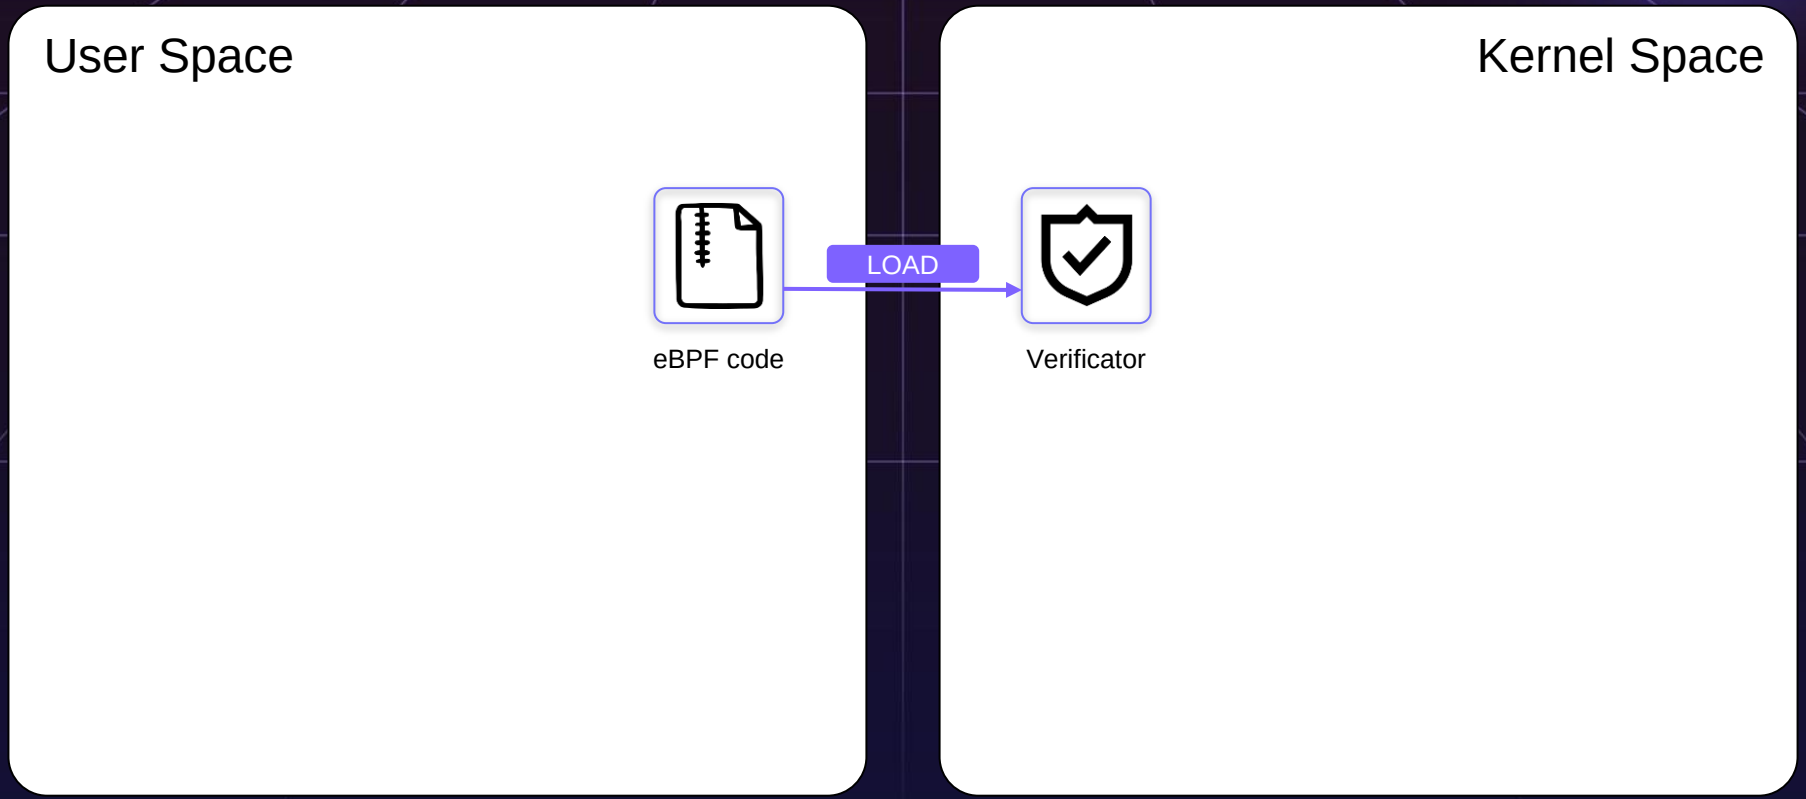

eBPF

БЕКОН

User Space

Kernel Space

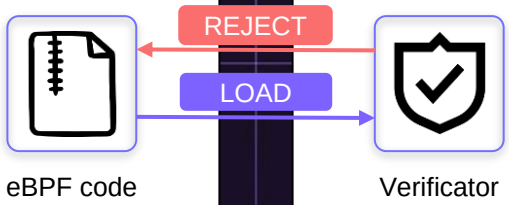

eBPF

БЕКОН

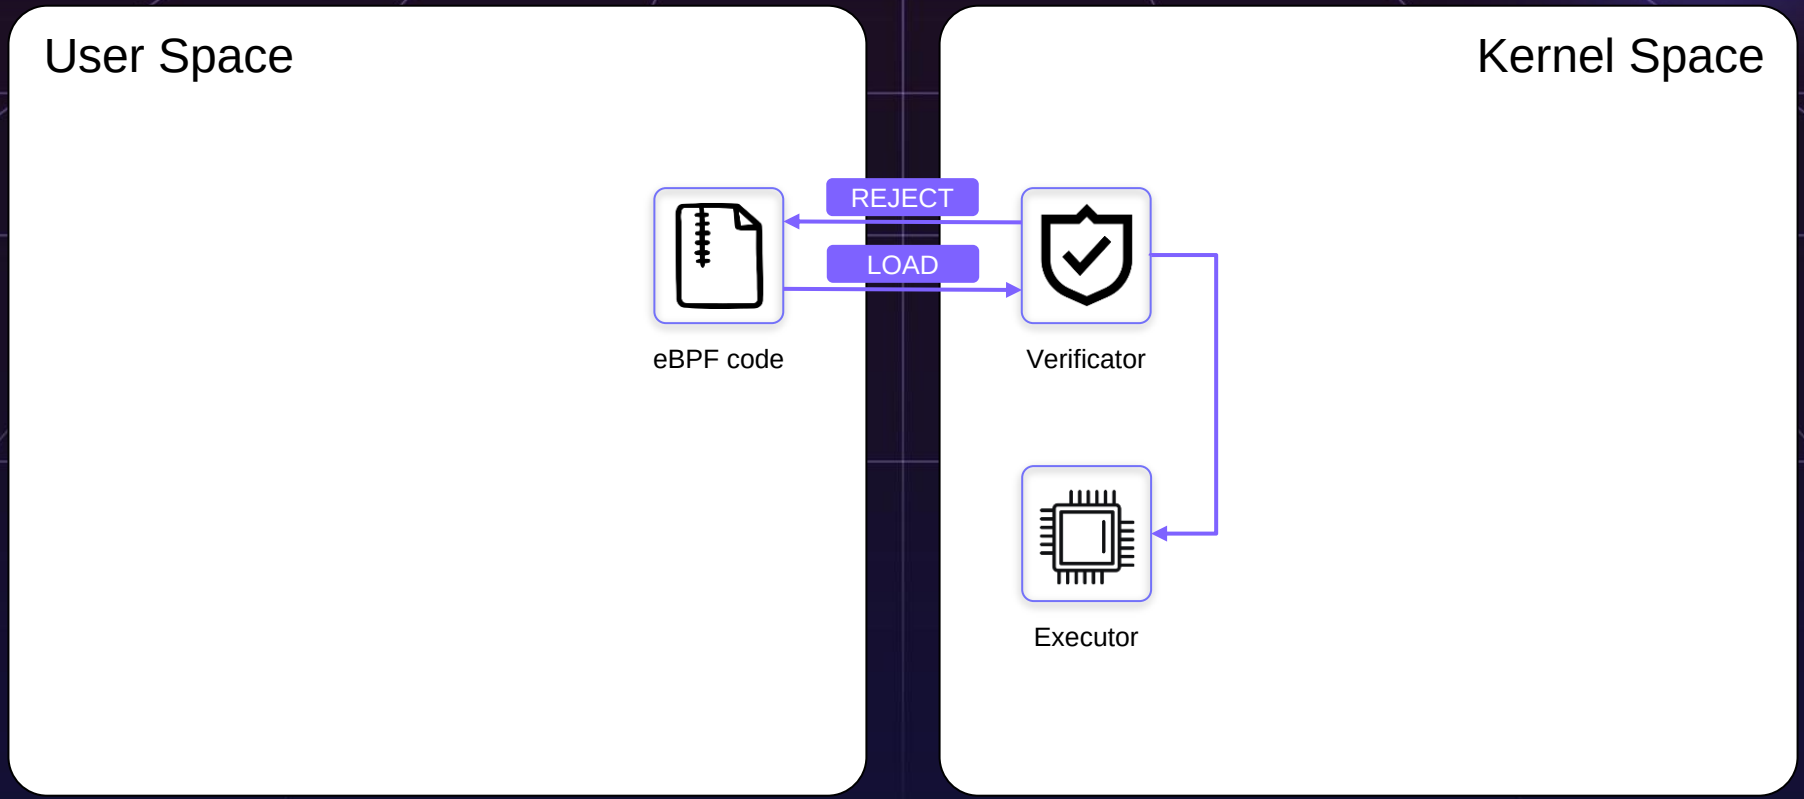
User Space
Kernel Space
REJECT
LOAD
eBPF code
Verificator
Executor

eBPF

БЕКОН

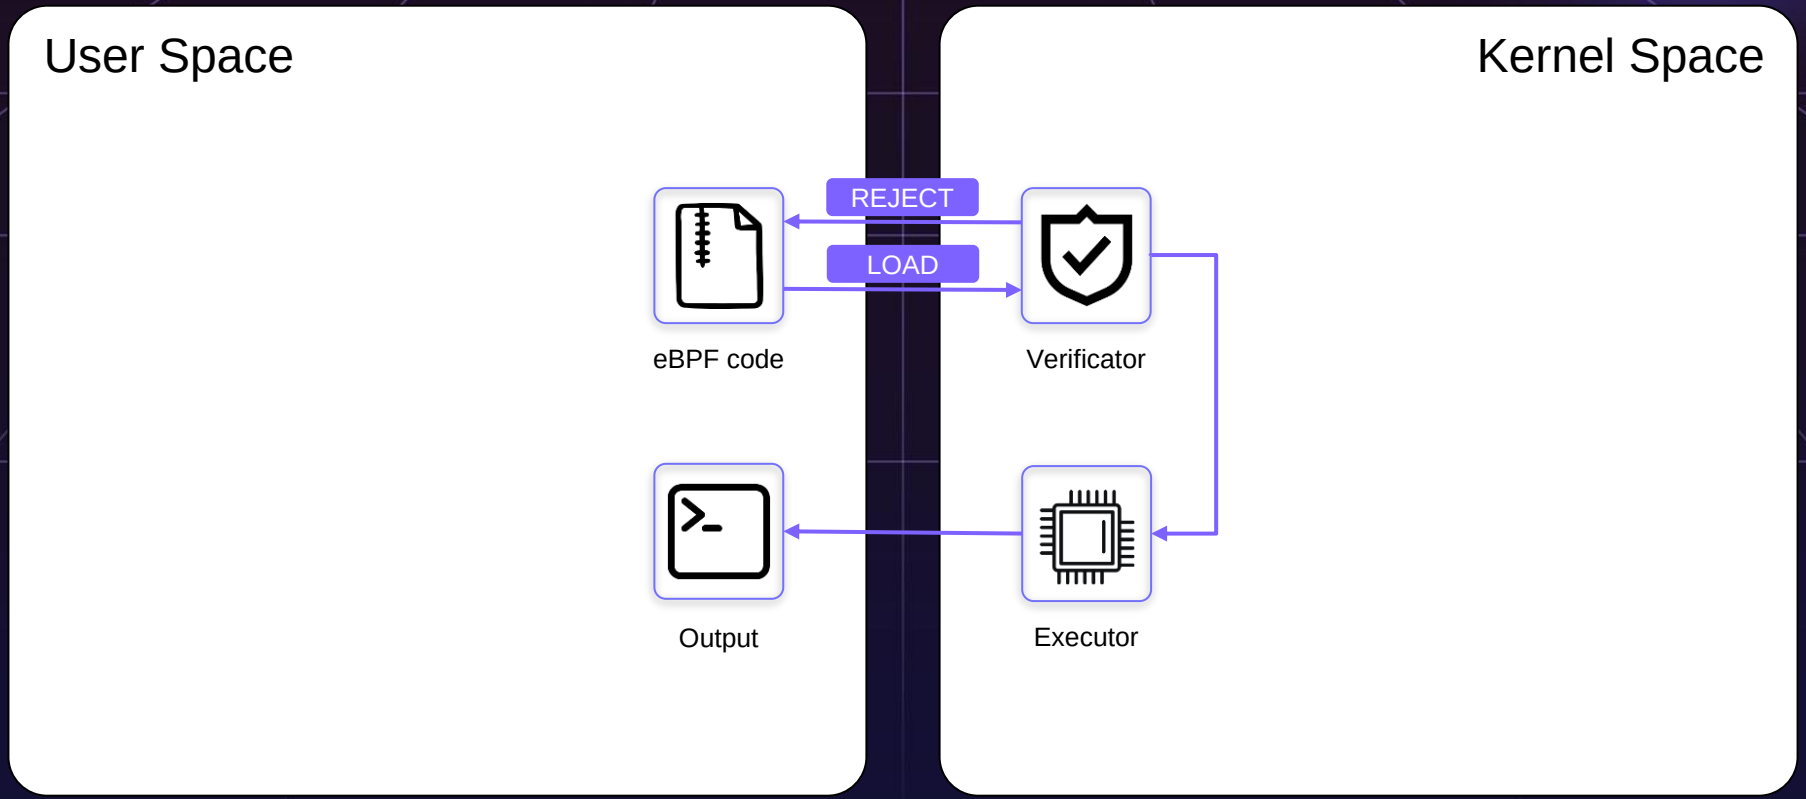
User Space
Kernel Space
REJECT
LOAD
eBPF code
Verificator
Output
Executor

eBPF

БЕКОН

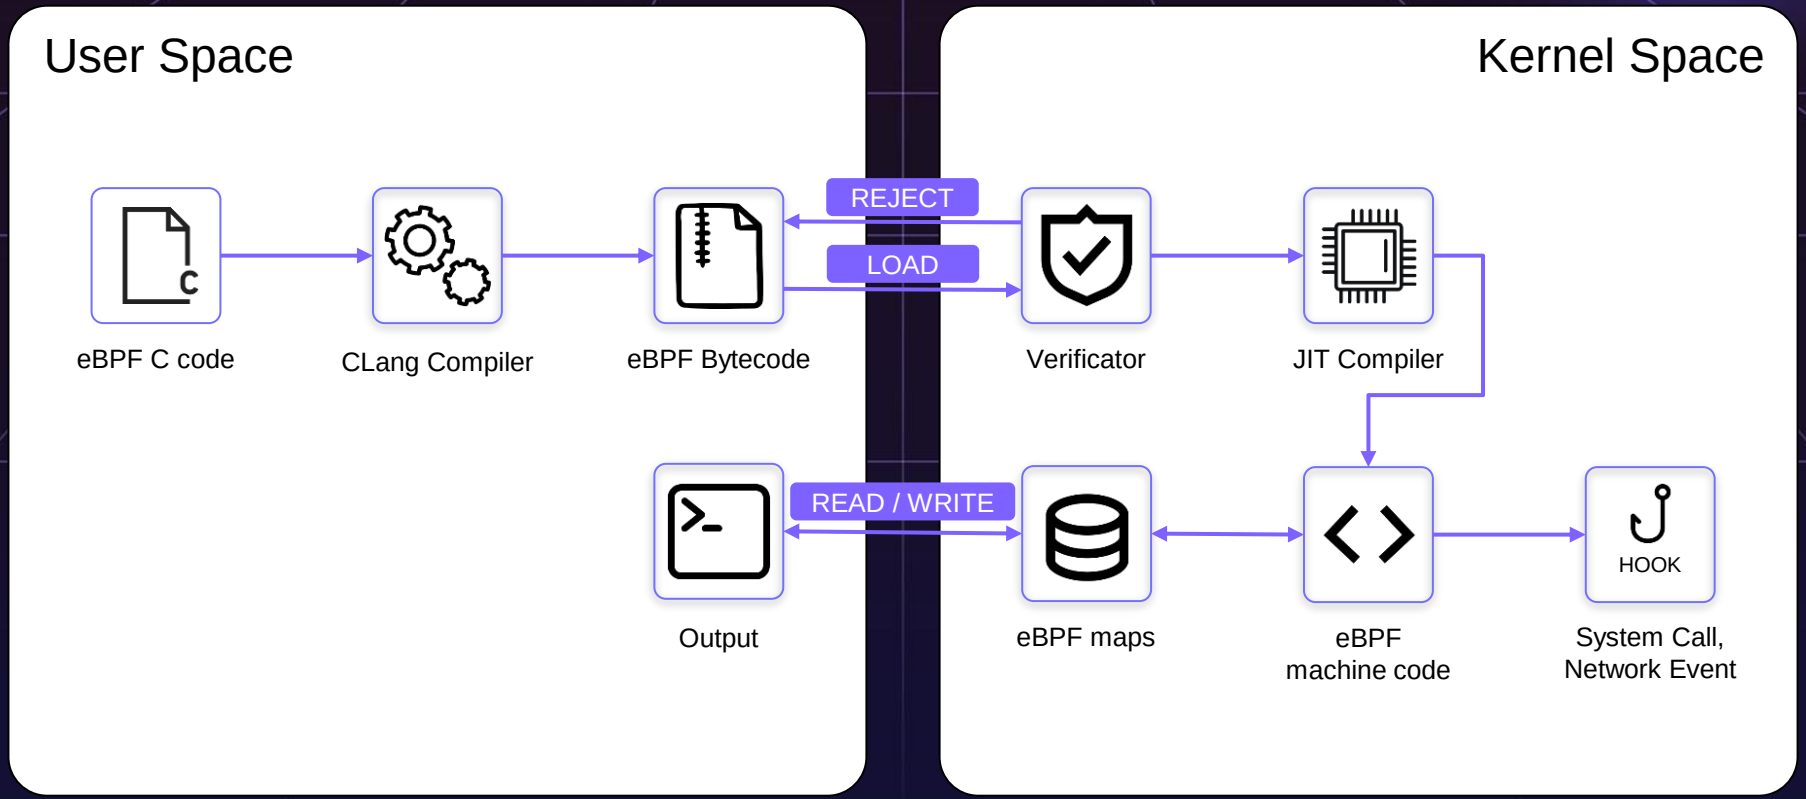
User Space
Kernel Space
eBPF C code
CLang Compiler
eBPF Bytecode
REJECT
LOAD
Verificator
JIT Compiler
Output
READ / WRITE
eBPF maps
eBPF
machine code
HOOK
System Call,
Network Event

БЕКОН

Cilium Policy

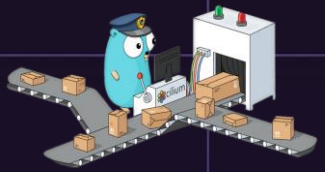

Cilium Policy

БЕКОН

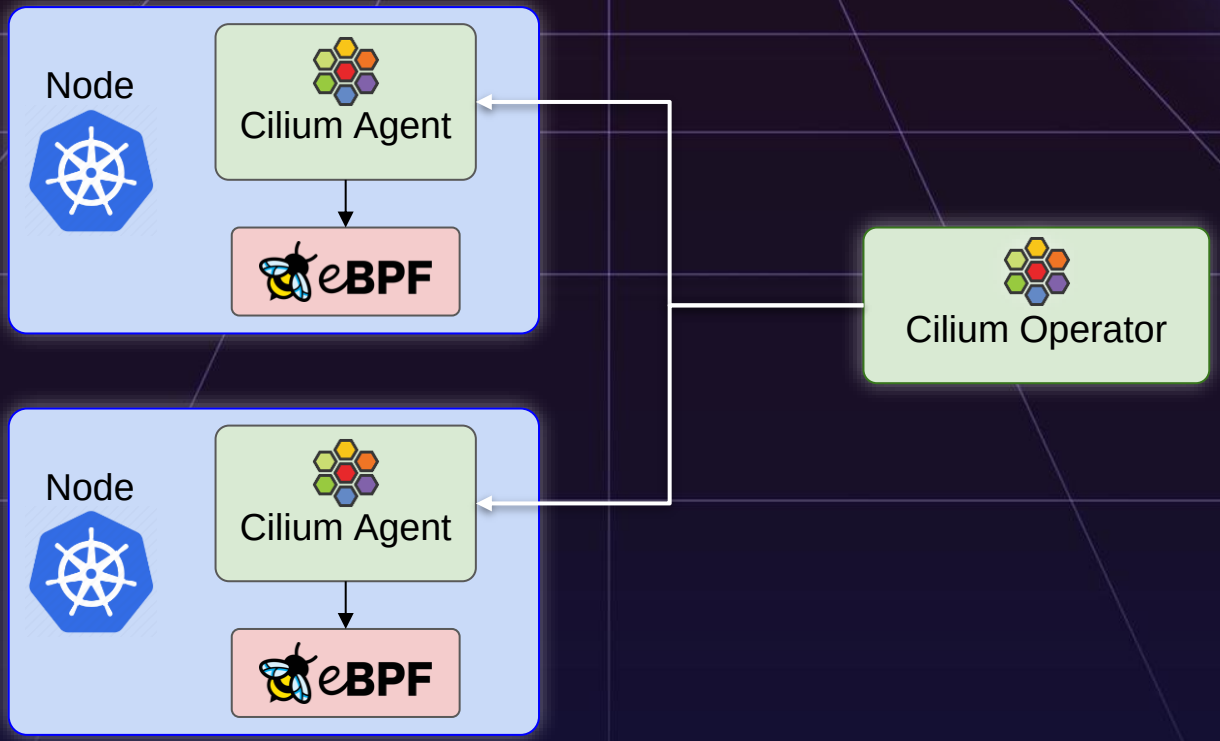

Cilium Policy

БЕКОН

Policy map

```
#ifdef POLICY_MAP
/* Per-endpoint policy enforcement map */
struct {
    __uint(type, BPF_MAP_TYPE_LPM_TRIE);
    __type(key, struct policy_key);
    __type(value, struct policy_entry);
    __uint(pinning, LIBBPF_PIN_BY_NAME);
    __uint(max_entries, POLICY_MAP_SIZE);
    __uint(map_flags, BPF_F_NO_PREALLOC);
} POLICY_MAP __section_maps_btf;
#endif
```

Cilium Policy

БЕКОН

Policy map

```
#ifdef POLICY_MAP
/* Per-endpoint policy enforcement map */
struct {
    __uint(type, BPF_MAP_TYPE_LPM_TRIE);
    __type(key, struct policy_key);
    __type(value, struct policy_entry);
    __uint(pinning, LIBBPF_PIN_BY_NAME);
    __uint(max_entries, POLICY_MAP_SIZE);
    __uint(map_flags, BPF_F_NO_PREALLOC);
} POLICY_MAP __section_maps_btf;
#endif
```

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        egress:1,
                pad:7;
    __u8        protocol;
    __u16       dport;
};
struct policy_entry {
    __be16      proxy_port;
    __u8        deny:1,
                wildcard_protocol:1,
                wildcard_dport:1,
                pad:5;
    __u8        auth_type;
    __u16       pad1;
    __u16       pad2;
    __u64       packets;
    __u64       bytes;
};
```

Cilium Policy

БЕКОН

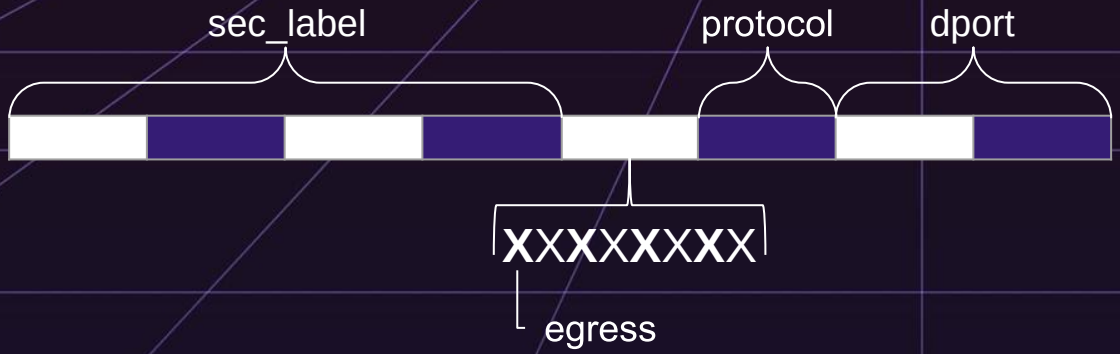

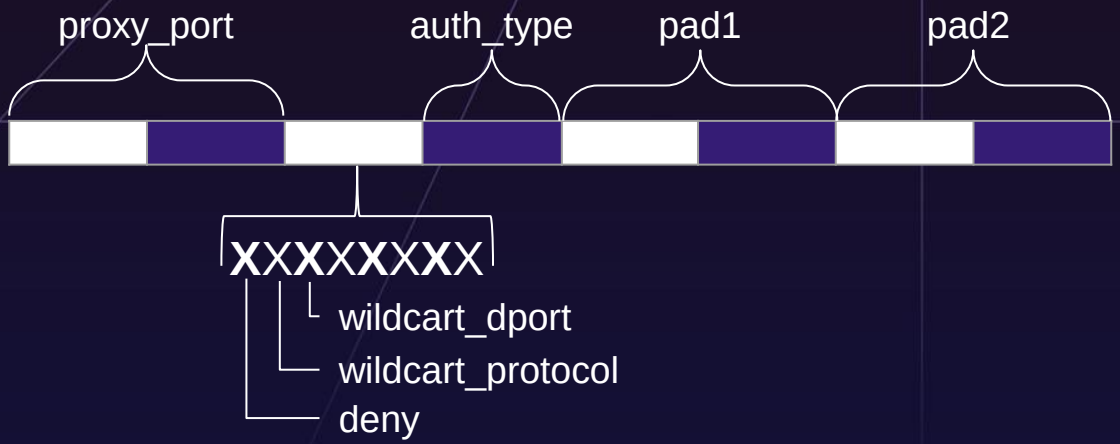


```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32           sec_label;
    __u8            egress:1,
                    pad:7;
    __u8            protocol;
    __u16           dport;
};
struct policy_entry {
    __be16          proxy_port;
    __u8            deny:1,
                    wildcard_protocol:1,
                    wildcard_dport:1,
                    pad:5;
    __u8            auth_type;
    __u16           pad1;
    __u16           pad2;
    __u64           packets;
    __u64           bytes;
};
```

Cilium Policy

БЕКОН

Changes

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol; /* can be wildcarded if 'dport' is fully wildcarded */
    __u16       dport; /* can be wildcarded with CIDR-like prefix */
};
struct policy_entry {
    __be16      proxy_port;
    __u8        deny:1,
            wildcard_protocol:1, /* protocol is fully wildcarded */
            wildcard_dport:1, /* dport is fully wildcarded */
            pad:5;
    __u8        auth_type;
    __u16       pad1;
    __u16       pad2;
    __u64       packets;
    __u64       bytes;
    __u64       categories;
};
```

БЕКОН

Поиск политик

Поиск политик

БЕКОН

Longest Prefix Match (LPM)

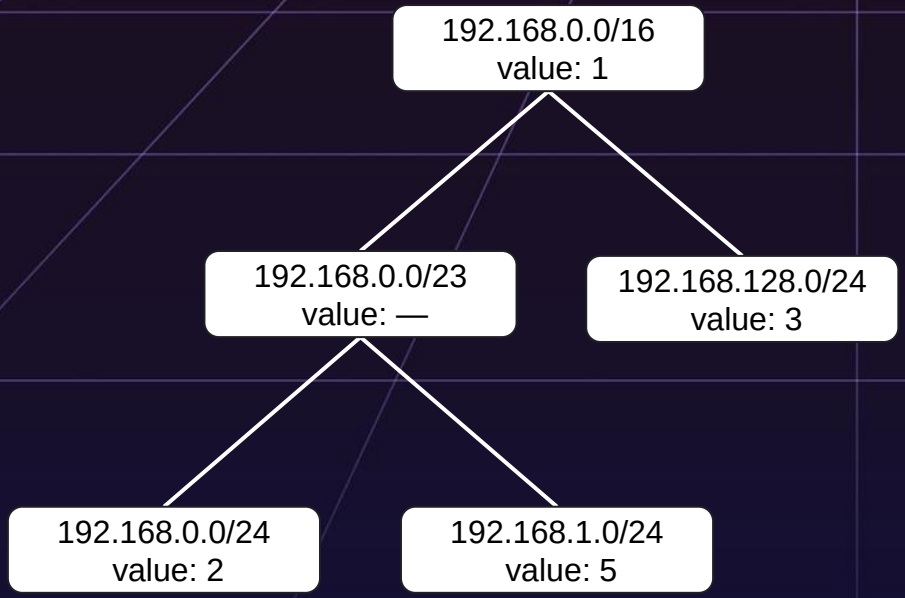

Поиск политик

БЕКОН

Longest Prefix Match (LPM)

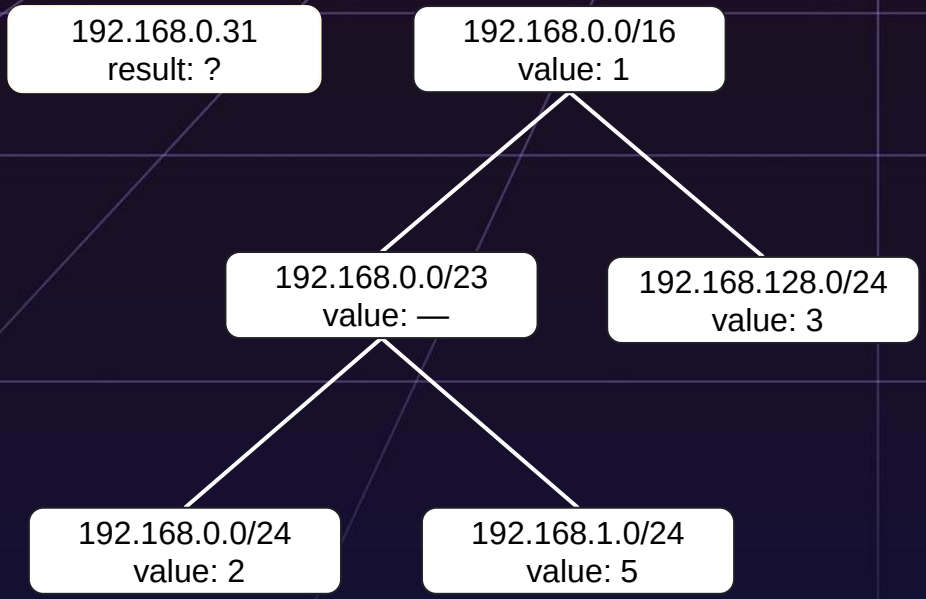

Поиск политик

БЕКОН

Longest Prefix Match (LPM)

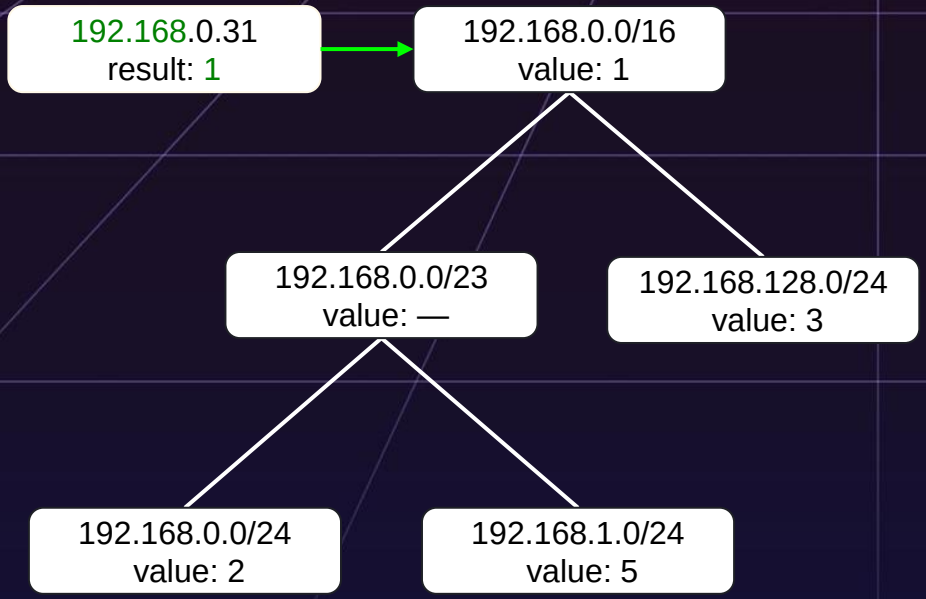

Поиск политик

БЕКОН

Longest Prefix Match (LPM)

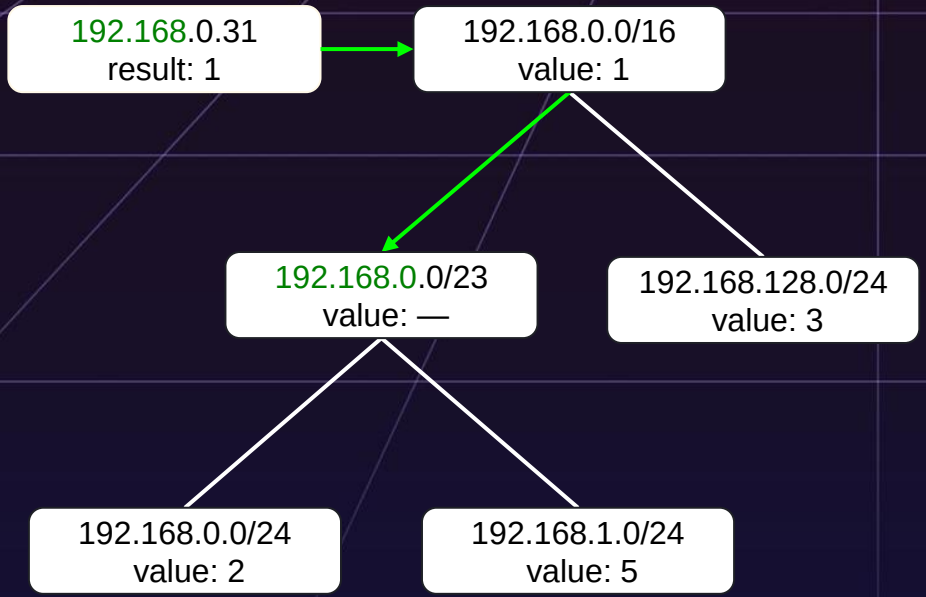

Поиск политик

БЕКОН

Longest Prefix Match (LPM)

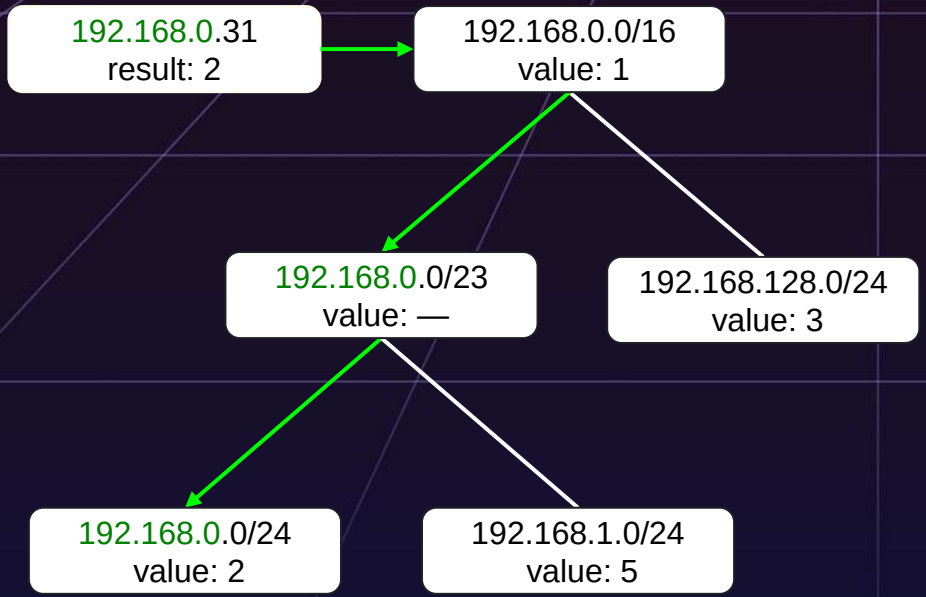

Поиск политик

БЕКОН

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};

struct bpf_lpm_trie_key {
    __u32   prefixlen;  /* up to 32 for AF_INET, 128 for AF_INET6 */
    __u8    data[0];    /* Arbitrary size */
};
```

Поиск политик

БЕКОН

1. Заполняем ВСЕ поля ключа

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};

static __always_inline int
__policy_can_access(const void *map,
            struct __ctx_buff *ctx, __u32 local_id,
            __u32 remote_id,
            __u16 ethertype __maybe_unused,
            __u16 dport,
            __u8 proto, int off __maybe_unused,
            int dir,
            bool is_untracked_fragment,
            __u8 *match_type, __s8 *ext_err,
            __u16 *proxy_port,
            __u8 mac_level, __u64 categories)
```

Поиск политик

БЕКОН

1. Заполняем ВСЕ поля ключа
2. Ищем по полному ключу

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};

policy = map_lookup_elem(map, &key);
if (likely(policy)) {
    if ((policy->categories & categories) == policy->categories) {
        astra_printk("AMAC Policy FOUND (1st lookup)");
        *match_type = POLICY_MATCH_L3_ONLY;
        goto check_policy;
    }
}
```

Поиск политик

БЕКОН

1. Заполняем ВСЕ поля ключа
2. Ищем по полному ключу
3. Зануляем sec_label

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};

key.sec_label = 0;
policy = map_lookup_elem(map, &key);
if (likely(policy)) {
    if ((policy->categories & categories) == policy->categories) {
        astra_printk("AMAC Policy FOUND (2nd lookup)");
        *match_type = POLICY_MATCH_L3_ONLY;
        goto check_policy;
    }
}
```

Поиск политик

БЕКОН

1. Заполняем ВСЕ поля ключа
2. Ищем по полному ключу
3. Зануляем sec_label
4. Возвращаем sec_label и зануляем astra_mac_level

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
key.astra_mac_level = 0;
key.dport = dport;
key.sec_label = remote_id;
policy = map_lookup_elem(map, &key);
if (likely(policy && !policy->wildcard_dport)) {
    printk("1. full L3/l4 match");
    cilium_dbg3(ctx, DBG_L4_CREATE, remote_id, local_id,
            dport << 16 | proto);
    *match_type = POLICY_MATCH_L3_L4;       /* 1. id/proto/port */
    goto check_policy;
}
```

Поиск политик

БЕКОН

1. Заполняем ВСЕ поля ключа
2. Ищем по полному ключу
3. Зануляем sec_label
4. Возвращаем sec_label и зануляем astra_mac_level
5. Зануляем sec_label

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
                pad:7;
    __u8        protocol;
    __u16       dport;
};


/* L4-only lookup. */
key.sec_label = 0;
l4policy = map_lookup_elem(map, &key);
```

Поиск политик

БЕКОН

1. Заполняем ВСЕ поля ключа
2. Ищем по полному ключу
3. Зануляем sec_label
4. Возвращаем sec_label и зануляем astra_mac_level
5. Зануляем sec_label
6. Решаем какая политика больше подходит

```
if (likely(l4policy && !l4policy->wildcard_dport)) {
    astra_printk("2. ANY/proto/port");
    *match_type = POLICY_MATCH_L4_ONLY;      /* 2. ANY/proto/port */
    goto check_l4_policy;
}

if (likely(policy && !policy->wildcard_protocol)) {
    astra_printk("3. id/proto/ANY");
    *match_type = POLICY_MATCH_L3_PROTO;         /* 3. id/proto/ANY */
    goto check_policy;
}

if (likely(l4policy && !l4policy->wildcard_protocol)) {
    astra_printk("4. ANY/proto/ANY");
    *match_type = POLICY_MATCH_PROTO_ONLY;       /* 4. ANY/proto/ANY */
    goto check_l4_policy;
}

if (likely(policy)) {
    astra_printk("5. id/ANY/ANY");
    *match_type = POLICY_MATCH_L3_ONLY;      /* 5. id/ANY/ANY */
    goto check_policy;
}

if (likely(l4policy)) {
    astra_printk("6. ANY/ANY/ANY");
    *match_type = POLICY_MATCH_ALL;          /* 6. ANY/ANY/ANY */
    goto check_l4_policy;
}
```

БЕКОН

Cilium в Kubernetes

Cilium в Kubernetes

БЕКОН

Добавление политик

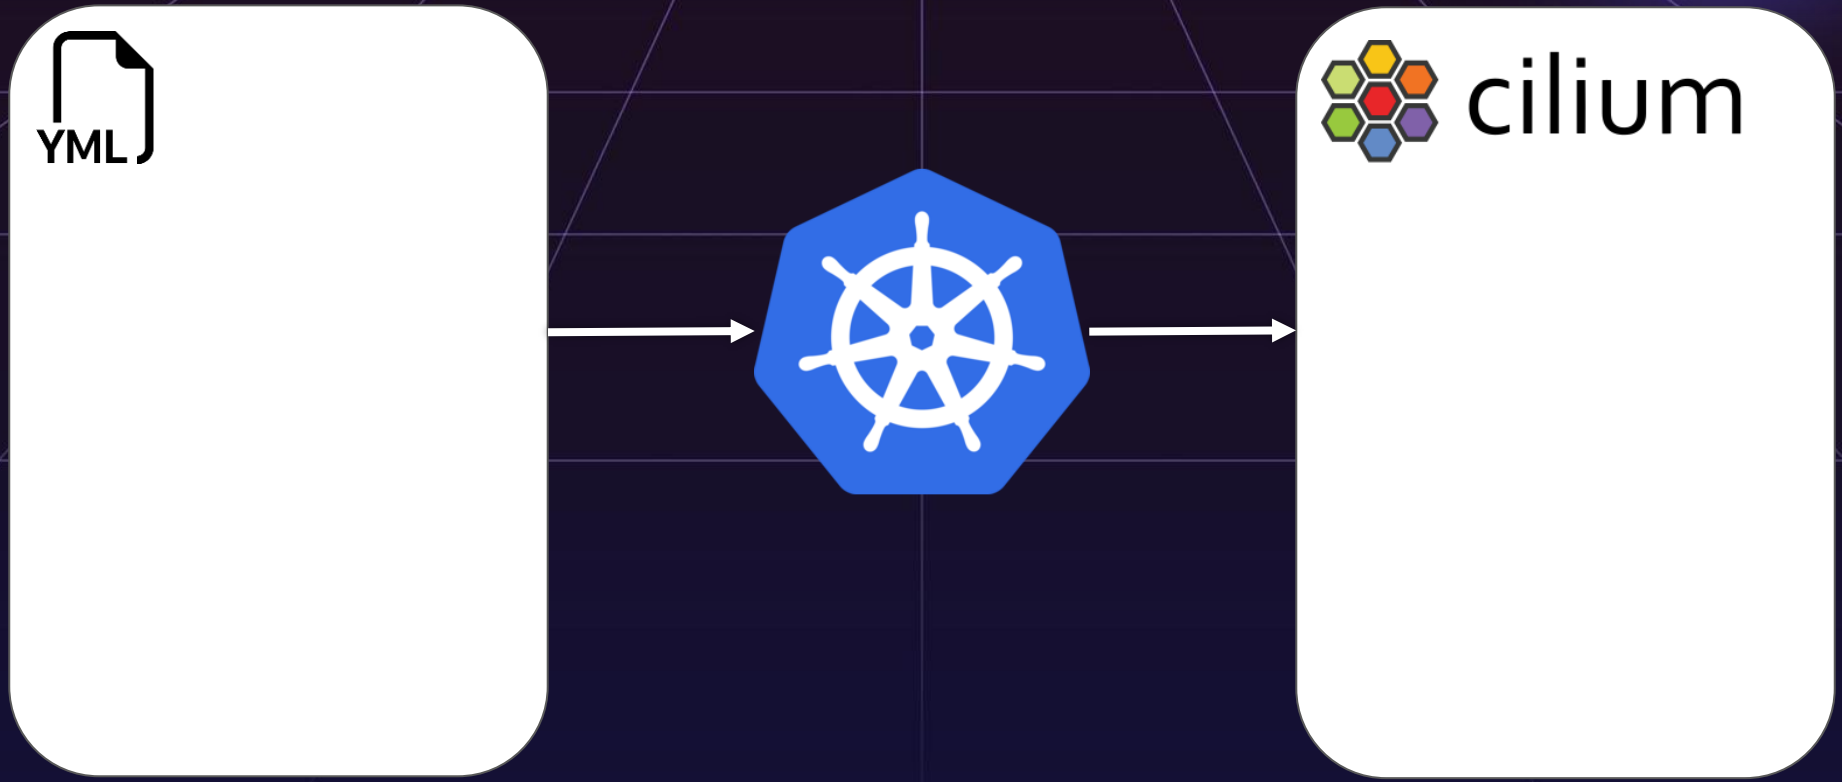

Cilium в Kubernetes

БЕКОН

Добавление политик

YML

```yaml
apiVersion: cilium.io/v2
kind: CiliumNetworkPolicy
metadata:
 name: "awesome-policy"
spec:
 ingress:
   - fromEntities:
       - all
 ingressDeny:
   - amacs:
       - level: 3
         categories: "0x1"
 egress:
   - toEntities:
       - all
```

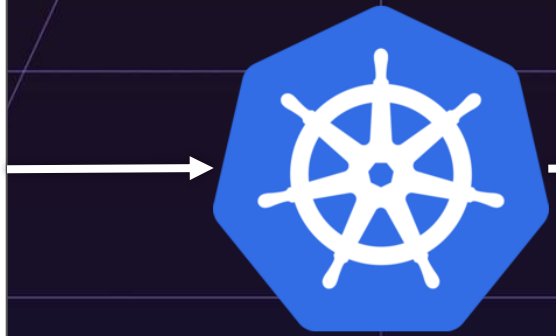

cilium

Cilium в Kubernetes

БЕКОН

Добавление политик

YML

```yaml
apiVersion: cilium.io/v2
kind: CiliumNetworkPolicy
metadata:
  name: "awesome-policy"
spec:
  ingress:
    - fromEntities:
        - all
  ingressDeny:
    - amacs:
        - level: 3
          categories: "0x1"
  egress:
    - toEntities:
        - all
```

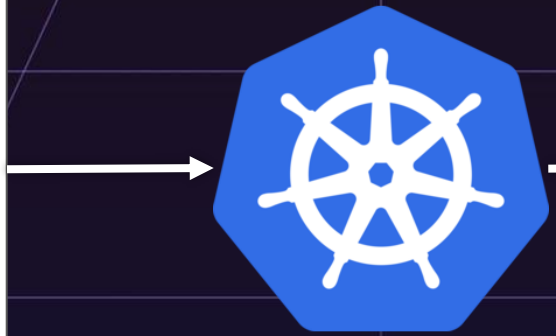

cilium

Cilium в Kubernetes

БЕКОН

Добавление политик

YML

```yaml
apiVersion: cilium.io/v2
kind: CiliumNetworkPolicy
metadata:
  name: "awesome-policy"
spec:
  ingress:
    - fromEntities:
        - all
  ingressDeny:
    - amacs:
        - level: 3
          categories: "0x1"
  egress:
    - toEntities:
        - all
```

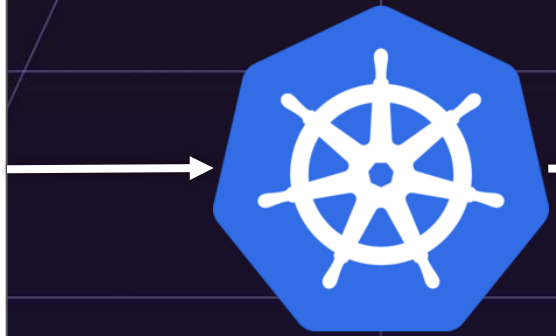

cilium

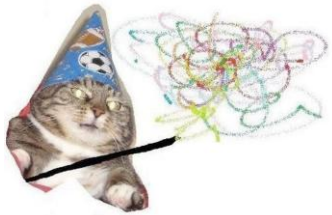

Cilium в Kubernetes

БЕКОН

На практике

```
± % k get po -A -o wide                                                    !140
NAMESPACE            NAME                                         READY   STATUS    RESTARTS       AGE    IP           NODE
kube-system          cilium-f6h5n                                 1/1     Running   0              113m   172.18.0.3   kind-control-plane
kube-system          cilium-kh47l                                 1/1     Running   0              113m   172.18.0.2   kind-worker
kube-system          cilium-operator-86649d6bc8-dfsbc             1/1     Running   0              113m   172.18.0.2   kind-worker
kube-system          coredns-668d6bf9bc-hmshb                     1/1     Running   0              12h    10.0.0.57    kind-control-plane
kube-system          coredns-668d6bf9bc-rdg5f                     1/1     Running   0              12h    10.0.0.34    kind-control-plane
kube-system          etcd-kind-control-plane                      1/1     Running   1 (137m ago)   12h    172.18.0.3   kind-control-plane
kube-system          kube-apiserver-kind-control-plane            1/1     Running   1 (137m ago)   12h    172.18.0.3   kind-control-plane
kube-system          kube-controller-manager-kind-control-plane   1/1     Running   1 (137m ago)   12h    172.18.0.3   kind-control-plane
kube-system          kube-proxy-4ppk2                             1/1     Running   1 (137m ago)   12h    172.18.0.2   kind-worker
kube-system          kube-proxy-j4cjq                             1/1     Running   1 (137m ago)   12h    172.18.0.3   kind-control-plane
kube-system          kube-scheduler-kind-control-plane            1/1     Running   1 (137m ago)   12h    172.18.0.3   kind-control-plane
local-path-storage   local-path-provisioner-7dc846544d-swgzq      1/1     Running   0              12h    10.0.0.101   kind-control-plane
```

Cilium в Kubernetes

БЕКОН

На практике

```
CLI for interacting with the local Cilium Agent

Usage:
  cilium-dbg [command]

Available Commands:
  bgp          Access to BGP control plane
  bpf          Direct access to local BPF maps
  build-config Resolve all of the configuration sources that apply to this node
  cgroups      Cgroup metadata
  cleanup      Remove system state installed by Cilium at runtime
  completion   Output shell completion code
  config       Cilium configuration options
  debuginfo    Request available debugging information from agent
  encrypt      Manage transparent encryption
  endpoint     Manage endpoints
  fqdn         Manage fqdn proxy
  help         Help about any command
  identity     Manage security identities
  ip           Manage IP addresses and associated information
  kvstore      Direct access to the kvstore
  lrp          Manage local redirect policies
  map          Access userspace cached content of BPF maps
  metrics      Access metric status
  monitor      Display BPF program events
  node         Manage cluster nodes
  nodeid       List node IDs and associated information
  policy       Manage security policies
  prefilter    Manage XDP CIDR filters
  preflight    Cilium upgrade helper
  recorder     Introspect or mangle pcap recorder
  service      Manage services & loadbalancers
  statedb      Inspect StateDB
  status       Display status of daemon
  version      Print version information

Flags:
      --config string   Config file (default is $HOME/.cilium.yaml)
  -D, --debug           Enable debug messages
  -h, --help            help for cilium-dbg
  -H, --host string     URI to server-side API
```

Cilium в Kubernetes

БЕКОН

На практике

```
CLI for interacting with the local Cilium Agent

Usage:
  cilium-dbg [command]

Available Commands:
  bgp           Access to BGP control plane
  bpf           Direct access to local BPF maps
  build-config  Resolve all of the configuration sources that apply to this node
  cgroups       Cgroup metadata
  cleanup       Remove system state installed by Cilium at runtime
  completion    Output shell completion code
  config        Cilium configuration options
  debuginfo     Request available debugging information from agent
  encrypt       Manage transparent encryption
  endpoint      Manage endpoints
  fqdn          Manage fqdn proxy
  help          Help about any command
  identity      Manage security identities
  ip            Manage IP addresses and associated information
  kvstore       Direct access to the kvstore
  lrp           Manage local redirect policies
  map           Access userspace cached content of BPF maps
  metrics       Access metric status
  monitor       Display BPF program events
  node          Manage cluster nodes
  nodeid        List node IDs and associated information
  policy        Manage security policies
  prefilter     Manage XDP CIDR filters
  preflight     Cilium upgrade helper
  recorder      Introspect or mangle pcap recorder
  service       Manage services & loadbalancers
  statedb       Inspect StateDB
  status        Display status of daemon
  version       Print version information

Flags:
      --config string   Config file (default is $HOME/.cilium.yaml)
  -D, --debug           Enable debug messages
  -h, --help            help for cilium-dbg
  -H, --host string     URI to server-side API
```

Cilium в Kubernetes

БЕКОН

На практике

```
CLI for interacting with the local Cilium Agent

Usage:
  cilium-dbg [command]

Available Commands:
  bgp           Access to BGP control plane
  bpf           Direct access to local BPF maps
  build-config  Resolve all of the configuration sources that apply to this node
  cgroups       Cgroup metadata
  cleanup       Remove system state installed by Cilium at runtime
  completion    Output shell completion code
  config        Cilium configuration options
  debuginfo     Request available debugging information from agent
  encrypt       Manage transparent encryption
  endpoint      Manage endpoints
  fqdn          Manage fqdn proxy
  help          Help about any command
  identity      Manage security identities
  ip            Manage IP addresses and associated information
  kvstore       Direct access to the kvstore
  lrp           Manage local redirect policies
  map           Access userspace cached content of BPF maps
  metrics       Access metric status
  monitor       Display BPF program events
  node          Manage cluster nodes
  nodeid        List node IDs and associated information
  policy        Manage security policies
  prefilter     Manage XDP CIDR filters
  preflight     Cilium upgrade helper
  recorder      Introspect or mangle pcap recorder
  service       Manage services & loadbalancers
  statedb       Inspect StateDB
  status        Display status of daemon
  version       Print version information

Flags:
      --config string   Config file (default is $HOME/.cilium.yaml)
  -D, --debug           Enable debug messages
  -h, --help            help for cilium-dbg
  -H, --host string     URI to server-side API
```

Cilium в Kubernetes

БЕКОН

На практике

```
CLI for interacting with the local Cilium Agent

Usage:
  cilium-dbg [command]

Available Commands:
  bgp           Access to BGP control plane
  bpf           Direct access to local BPF maps
  build-config  Resolve all of the configuration sources that apply to this node
  cgroups       Cgroup metadata
  cleanup       Remove system state installed by Cilium at runtime
  completion    Output shell completion code
  config        Cilium configuration options
  debuginfo     Request available debugging information from agent
  encrypt       Manage transparent encryption
  endpoint      Manage endpoints
  fqdn          Manage fqdn proxy
  help          Help about any command
  identity      Manage security identities
  ip            Manage IP addresses and associated information
  kvstore       Direct access to the kvstore
  lrp           Manage local redirect policies
  map           Access userspace cached content of BPF maps
  metrics       Access metric status
  monitor       Display BPF program events
  node          Manage cluster nodes
  nodeid        List node IDs and associated information
  policy        Manage security policies
  prefilter     Manage XDP CIDR filters
  preflight     Cilium upgrade helper
  recorder      Introspect or mangle pcap recorder
  service       Manage services & loadbalancers
  statedb       Inspect StateDB
  status        Display status of daemon
  version       Print version information

Flags:
      --config string   Config file (default is $HOME/.cilium.yaml)
  -D, --debug           Enable debug messages
  -h, --help            help for cilium-dbg
  -H, --host string     URI to server-side API
```

Cilium в Kubernetes

БЕКОН

На практике

```
root@kind-worker:~# cilium map list
Name                          Num entries   Num errors   Cache enabled
cilium_runtime_config         0             0            false
cilium_policy_03829           0             0            false
cilium_tunnel_map             2             0            true
cilium_lb6_services_v2        0             0            true
cilium_lb6_reverse_nat        0             0            true
cilium_lb4_reverse_nat        4             0            true
cilium_lb4_source_range       0             0            true
cilium_lxc                    7             0            true
cilium_ipcache                21            0            true
cilium_lb6_source_range       0             0            true
cilium_policy_00084           0             0            false
cilium_lb4_services_v2        10            0            true
cilium_lb4_backends_v3        6             0            true
cilium_lb6_backends_v3        0             0            true
cilium_auth_map               0             0            false
cilium_node_map               0             0            false
cilium_metrics                0             0            false
cilium_l2_responder_v4        0             0            false
```

Cilium в Kubernetes

БЕКОН

На практике

```
root@kind-worker:~# cilium bpf policy list
Endpoint ID: 84
Path: /sys/fs/bpf/tc/globals/cilium_policy_00084

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0


Endpoint ID: 3829
Path: /sys/fs/bpf/tc/globals/cilium_policy_03829

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES    PACKETS   PREFIX
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    176762   1976      0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    209149   2254      0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0        0         0
```

Cilium в Kubernetes

БЕКОН

На практике

```
root@kind-worker:~# cilium bpf policy list
Endpoint ID: 84
Path: /sys/fs/bpf/tc/globals/cilium_policy_00084

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0


Endpoint ID: 3829
Path: /sys/fs/bpf/tc/globals/cilium_policy_03829

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES    PACKETS   PREFIX
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    176762   1976      0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    209149   2254      0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0        0         0



root@kind-worker:~# cilium endpoint list
ENDPOINT   POLICY (ingress)   POLICY (egress)   IDENTITY   LABELS (source:key[=value])   IPv6         IPv4          STATUS
           ENFORCEMENT        ENFORCEMENT
84         Disabled           Disabled          1          reserved:host                                            ready
3829       Disabled           Disabled          4          reserved:health               fd00::1dd    10.0.1.132    ready
```

Cilium в Kubernetes

БЕКОН

На практике

```
root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_03829
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  fa 08 00 00 00 00 00 00
df 22 03 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  3e 0a 00 00 00 00 00 00
51 b6 03 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 3 elements
```

Cilium в Kubernetes

БЕКОН

На практике

```yaml
apiVersion: v1
kind: Pod
metadata:
 name: alpine
 labels:
   app: alpine-based
spec:
 containers:
   - name: alpine
     tty: true
     stdin: true
     image: "alpine:latest"
```

```yaml
apiVersion: "cilium.io/v2"
kind: CiliumNetworkPolicy
metadata:
 name: "block-80-tcp"
spec:
 endpointSelector:
   matchLabels:
     app: alpine-based
 egressDeny:
   - toPorts:
       - ports:
           - port: "80"
             protocol: TCP
 egress:
   - toEntities:
       - all
```

```
% kubectl get po --show-labels -o wide                                                                    !148
NAME     READY   STATUS    RESTARTS   AGE    IP           NODE          NOMINATED NODE   READINESS GATES   LABELS
alpine   1/1     Running   0          4m1s   10.0.1.186   kind-worker   <none>           <none>            app=alpine-based
```

Cilium в Kubernetes

БЕКОН

На практике

```
root@kind-worker:~# cilium map list
Name                          Num entries   Num errors   Cache enabled
cilium_runtime_config         0             0            false
cilium_policy_03829           0             0            false
cilium_tunnel_map             2             0            true
cilium_lb6_services_v2        0             0            true
cilium_lb6_reverse_nat        0             0            true
cilium_lb4_reverse_nat        4             0            true
cilium_lb4_source_range       0             0            true
cilium_lxc                    9             0            true
cilium_ipcache                23            0            true
cilium_lb6_source_range       0             0            true
cilium_policy_00084           0             0            false
cilium_policy_00489           0             0            false
cilium_lb4_services_v2        10            0            true
cilium_lb4_backends_v3        6             0            true
cilium_lb6_backends_v3        0             0            true
cilium_metrics                0             0            false
cilium_l2_responder_v4        0             0            false
cilium_auth_map               0             0            false
cilium_node_map               0             0            false
```

NEW

Cilium в Kubernetes

БЕКОН

На практике

```
root@kind-worker:~# cilium endpoint list
ENDPOINT   POLICY (ingress)   POLICY (egress)   IDENTITY   LABELS (source:key[=value])                                                IPv6        IPv4         STATUS
           ENFORCEMENT        ENFORCEMENT
84         Disabled           Disabled          1          reserved:host                                                                                   ready
489        Disabled           Disabled          22989      k8s:app=alpine-based                                                       fd00::138   10.0.1.186   ready
                                                           k8s:io.cilium.k8s.namespace.labels.kubernetes.io/metadata.name=default
                                                           k8s:io.cilium.k8s.policy.cluster=kind-kind
                                                           k8s:io.cilium.k8s.policy.serviceaccount=default
                                                           k8s:io.kubernetes.pod.namespace=default
3829       Disabled           Disabled          4          reserved:health                                                            fd00::1dd   10.0.1.132   ready
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

port

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

protocol

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

egress

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

pad1, pad2

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

astra_mac_level

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

sec_label

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

Cilium в Kubernetes

БЕКОН

На практике

```
Endpoint ID: 489
Path: /sys/fs/bpf/tc/globals/cilium_policy_00489

POLICY   DIRECTION   LABELS (source:key[=value])   PORT/PROTO   AMAC:LEVEL/CATEGORIES   PROXY PORT   AUTH TYPE   BYTES   PACKETS   PREFIX
Deny     Egress      reserved:unknown              80/TCP       0/0x0                   NONE         disabled    0       0         24
Allow    Ingress     reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Ingress     reserved:host                 ANY          0/0x0                   NONE         disabled    0       0         0
Allow    Egress      reserved:unknown              ANY          0/0x0                   NONE         disabled    0       0         0

root@kind-worker:~# bpftool map dump pinned /sys/fs/bpf/tc/globals/cilium_policy_00489
key:
48 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
60 00 00 00 00 00 00 00  00 00 00 00 01 06 00 50
value:
00 00 01 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 00 00 00 00  00 00 00 00 01 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
key:
48 00 00 00 01 00 00 00  00 00 00 00 00 00 00 00
value:
00 00 06 00 00 00 00 00  00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
Found 4 elements
```

lpm_key

```
struct policy_key {
    struct bpf_lpm_trie_key lpm_key;
    __u32       sec_label;
    __u8        astra_mac_level;
    __u8    pad1;
    __u16   pad2;
    __u8        egress:1,
            pad:7;
    __u8        protocol;
    __u16       dport;
};
```

БЕКОН

Мониторинг и метрики

Мониторинг и метрики

БЕКОН

Hubble

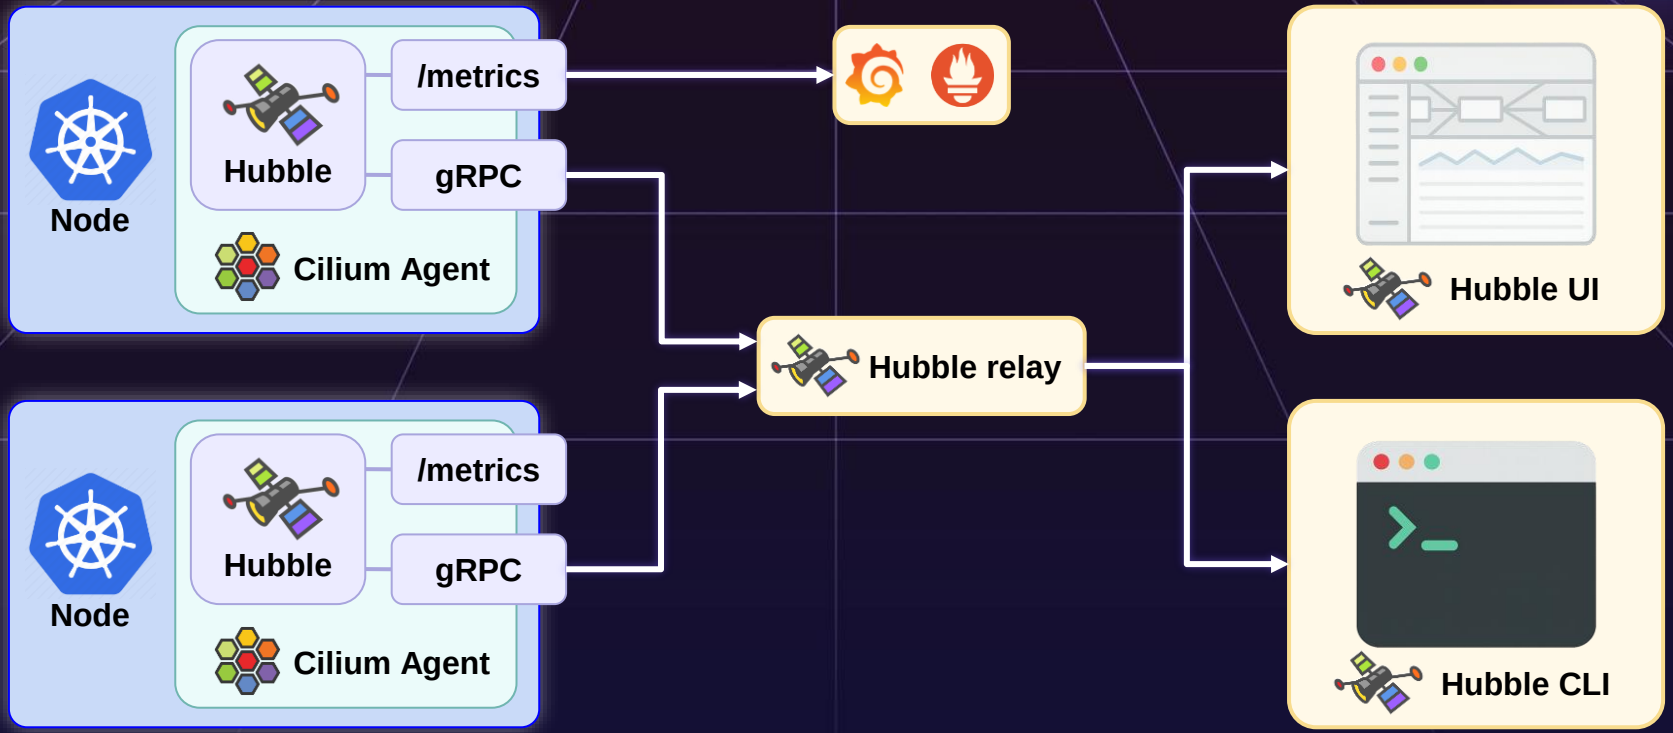

Мониторинг и метрики

БЕКОН

Hubble Flow

```
message Flow {
  …
  Ethernet ethernet = 4;
  IP IP = 5;
  Layer4 l4 = 6;
  …
  repeated Policy egress_allowed_by = 21001;
  repeated Policy ingress_allowed_by = 21002;
  repeated Policy egress_dropped_by = 21003;
  repeated Policy ingress_dropped_by = 21004;
}
```

```
message IP {
   …
   AstraMac astra_mac = 5;
}
```

```
message AstraMac {
  uint32 level = 1;
  string categories = 2;
}
```

Мониторинг и метрики

БЕКОН

На практике

```
default/ppg1:44959 (ID:454) <> default/ppg0:27015 (ID:16978) policy-verdict:L3-Only INGRESS DENIED (UDP)
default/ppg1:44959 (ID:454) <> default/ppg0:27015 (ID:16978) Policy denied by denylist DROPPED (UDP)
```


```
"flow": {
  …
  "IP": {
    "source": "10.0.1.179",
    "destination": "10.0.1.72",
    "ipVersion": "IPv4",
    "astra_mac": {
      "level": 64,
      "categories": "0x3"
    }
  },
  …
  "ingress_dropped_by": [
    {
      "name": "rule-with-1-level-categories-set",
      "namespace": "default",
      "labels": [
        "k8s:io.cilium.k8s.policy.derived-from=CiliumNetworkPolicy",
        "k8s:io.cilium.k8s.policy.name=rule-with-1-level-categories-set",
        "k8s:io.cilium.k8s.policy.namespace=default",
        "k8s:io.cilium.k8s.policy.uid=b58ba4f8-588e-4674-be03-ddb07a4cbee8"
      ],
      "revision": "3"
    }
  ]
}
```


Мониторинг и метрики

БЕКОН

На практике

```
default/ppg1:44959 (ID:454) <> default/ppg0:27015 (ID:16978) policy-verdict:L3-Only INGRESS DENIED (UDP)
default/ppg1:44959 (ID:454) <> default/ppg0:27015 (ID:16978) Policy denied by denylist DROPPED (UDP)
```


```
"flow": {
  …
  "IP": {
    "source": "10.0.1.179",
    "destination": "10.0.1.72",
    "ipVersion": "IPv4",
    "astra_mac": {
      "level": 64,
      "categories": "0x3"
    }
  },
  …
  "ingress_dropped_by": [
    {
      "name": "rule-with-1-level-categories-set",
      "namespace": "default",
      "labels": [
        "k8s:io.cilium.k8s.policy.derived-from=CiliumNetworkPolicy",
        "k8s:io.cilium.k8s.policy.name=rule-with-1-level-categories-set",
        "k8s:io.cilium.k8s.policy.namespace=default",
        "k8s:io.cilium.k8s.policy.uid=b58ba4f8-588e-4674-be03-ddb07a4cbee8"
      ],
      "revision": "3"
    }
  ]
}
```


Мониторинг и метрики

БЕКОН

Метрики

pkg/hubble/metrics/drop/handler.go

```
...  ...  @@ -26,7 +26,7 @@ func (d *dropHandler) Init(registry *prometheus.Registry, options api.Options) e
26   26           d.context = c
27   27
28   28           contextLabels := d.context.GetLabelNames()
29        -       labels := append(contextLabels, "reason", "protocol")
     29   +       labels := append(contextLabels, "reason", "protocol", "astra_mac", "node")
30   30
31   31           d.drops = prometheus.NewCounterVec(prometheus.CounterOpts{
32   32                   Namespace: api.DefaultPrometheusNamespace,
...  ...  @@ -62,6 +62,13 @@ func (d *dropHandler) ProcessFlow(ctx context.Context, flow *flowpb.Flow) error
62   62
63   63           labels := append(contextLabels, flow.GetDropReasonDesc().String(), v1.FlowProtocol(flow))
64   64
     65   +       if flow.IP != nil && flow.IP.AstraMac != nil {
     66   +               labels = append(labels, "Exist")
     67   +       } else {
     68   +               labels = append(labels, "None")
     69   +       }
     70   +       labels = append(labels, flow.NodeName)
     71   +
65   72           d.drops.WithLabelValues(labels...).Inc()
66   73           return nil
67   74   }
```

Мониторинг и метрики

БЕКОН

Grafana

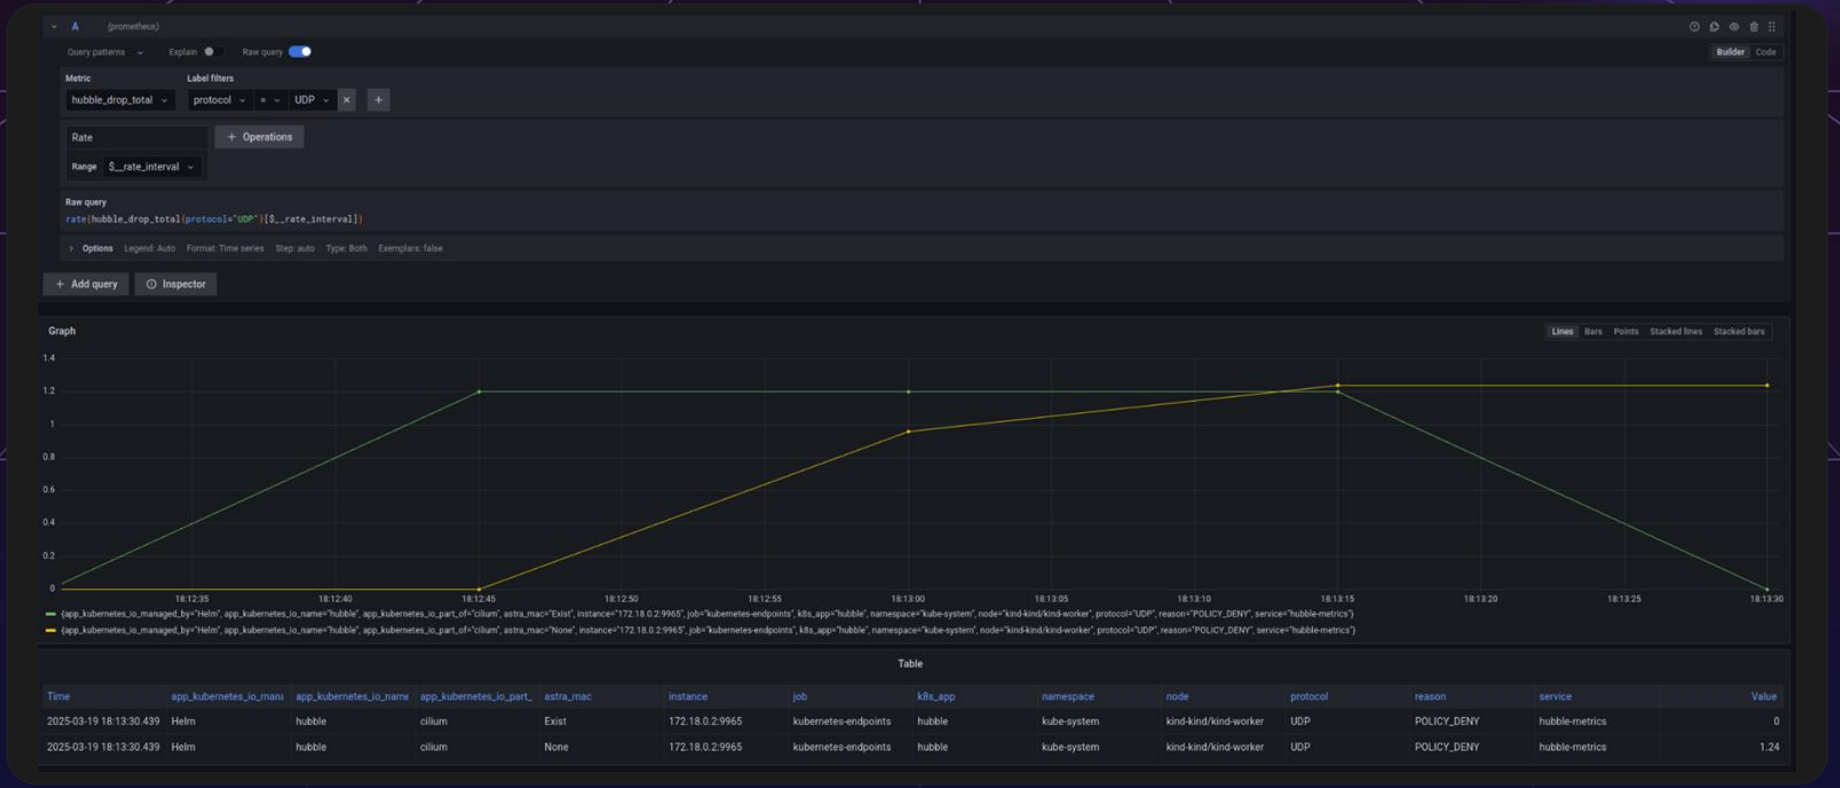

Мониторинг и метрики

БЕКОН

Grafana

БЕКОН

Итоги

Итоги

БЕКОН

- ✓ Реализовали eBPF-логику
- ✓ Обновили структуру политик
- ✓ Добавили Go-обвязку
- ✓ Прокачали Hubble UI/CLI
- ✓ Запустили интеграционные и connectivity-тесты

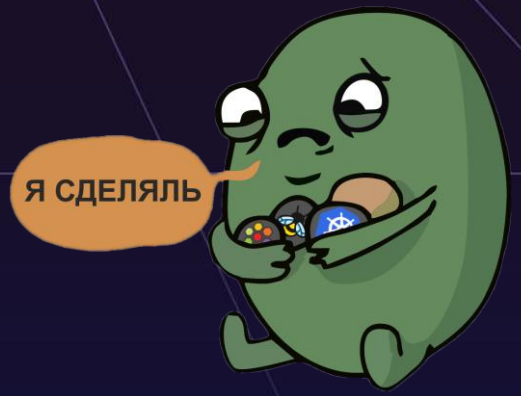

3 июня 2025 Москва, LOFT HALL#2

Конференция по БЕзопасности
КОНтейнеров и контейнерных сред

БЕКОН

LUNTRY

@Toshka_31
abaranov@astralinux.ru
astralinux.ru

astralinux.ru