

# Demystifying eBPF

## eBPF Firewall from scratch

Filip Nikolic

Kubernetes Engineer at PostFinance

# CNCF Projects

- Falco
  - Security
- Pixie
  - Observability
- Cilium
  - Networking







**User Space**





User Space



Kernel Space





User Space



Kernel Space



NIC

Disk

...



User Space



System Call Interface

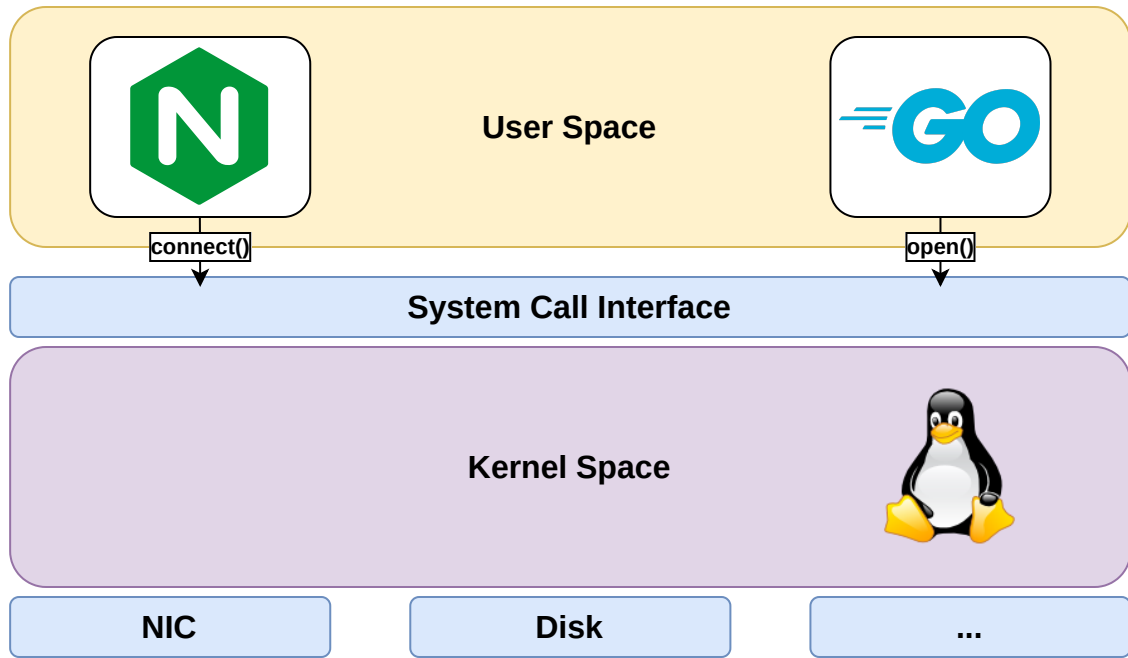
Kernel Space



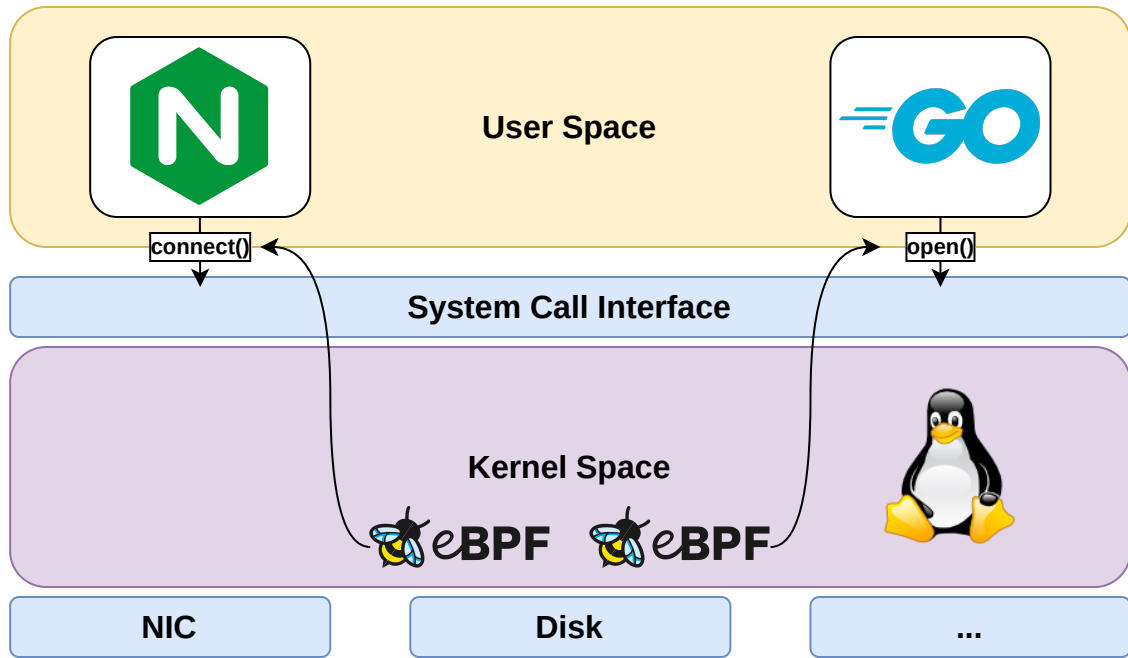
NIC

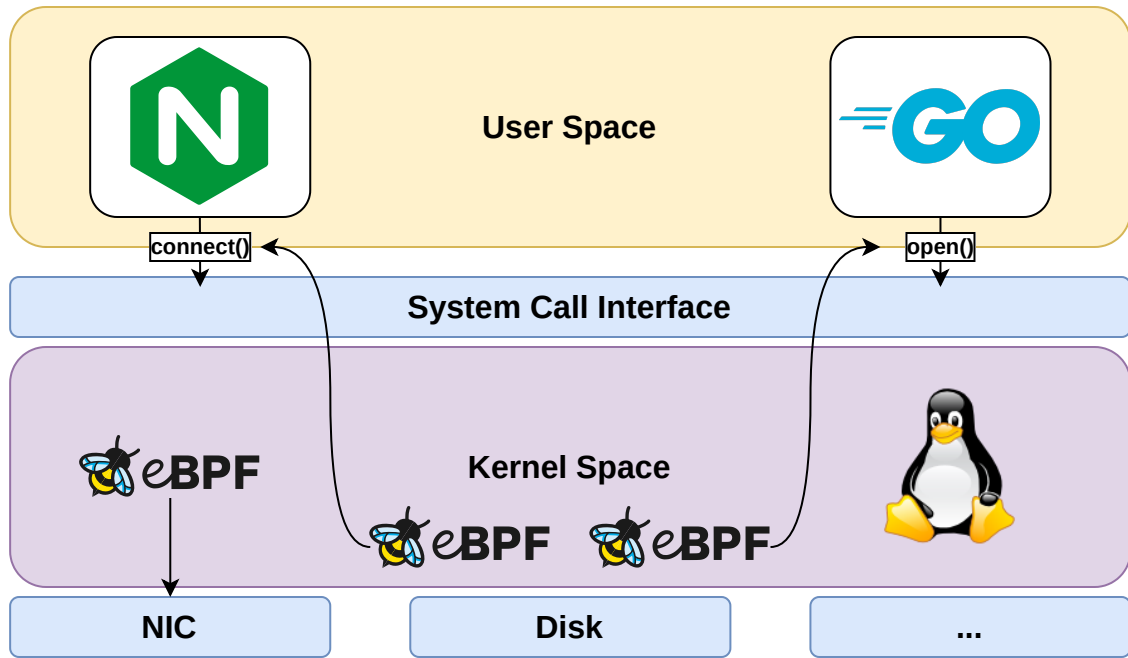
Disk


...











# Pass everything

```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // pass all packets
5     return XDP_PASS;
6 }
```

# Pass everything

```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // pass all packets
5     return XDP_PASS;
6 }
```

# Pass everything

```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // pass all packets
5     return XDP_PASS;
6 }
```

# Pass everything

```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // pass all packets
5     return XDP_PASS;
6 }
```


# Pass everything

```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // pass all packets
5     return XDP_PASS;
6 }
```



# Pass everything

```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // pass all packets
5     return XDP_PASS;
6 }
```



# Drop everything

```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // drop all packets
5     return XDP_DROP;
6 }
```

# Drop everything

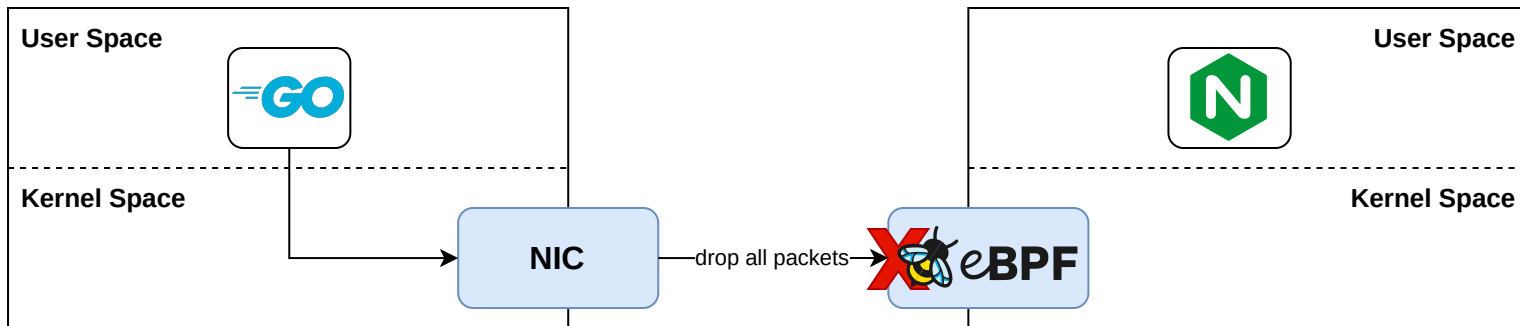
```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // drop all packets
5     return XDP_DROP;
6 }
```

# Drop everything

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      // drop all packets
5      return XDP_DROP;
6  }
```

# Drop everything

```
1 SEC("xdp_kcd_zurich")
2 int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3 {
4     // drop all packets
5     return XDP_DROP;
6 }
```



# Pseudocode (what we want)

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      protocol = find_out_protocol;
5
6      // drop all packets of type unwanted protocol
7      if (protocol == UNWANTED_PROTOCOL)
8          return XDP_DROP;
9
10     // pass all other packets
11     return XDP_PASS;
12 }
```

# Pseudocode (what we want)

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      protocol = find_out_protocol;
5
6      // drop all packets of type unwanted protocol
7      if (protocol == UNWANTED_PROTOCOL)
8          return XDP_DROP;
9
10     // pass all other packets
11     return XDP_PASS;
12 }
```

# Pseudocode (what we want)

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      protocol = find_out_protocol;
5
6      // drop all packets of type unwanted protocol
7      if (protocol == UNWANTED_PROTOCOL)
8          return XDP_DROP;
9
10     // pass all other packets
11     return XDP_PASS;
12 }
```



# Pseudocode (what we want)


```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      protocol = find_out_protocol;
5
6      // drop all packets of type unwanted protocol
7      if (protocol == UNWANTED_PROTOCOL)
8          return XDP_DROP;
9
10     // pass all other packets
11     return XDP_PASS;
12 }
```

# Pseudocode (what we want)

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      protocol = find_out_protocol;
5
6      // drop all packets of type unwanted protocol
7      if (protocol == UNWANTED_PROTOCOL)
8          return XDP_DROP;
9
10     // pass all other packets
11     return XDP_PASS;
12 }
```


# Pseudocode (what we want)

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      protocol = find_out_protocol;
5
6      // drop all packets of type unwanted protocol
7      if (protocol == UNWANTED_PROTOCOL)
8          return XDP_DROP;
9
10     // pass all other packets
11     return XDP_PASS;
12 }
```



# Pseudocode (what we want)

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      protocol = find_out_protocol;
5
6      // drop all packets of type unwanted protocol
7      if (protocol == UNWANTED_PROTOCOL)
8          return XDP_DROP;
9
10     // pass all other packets
11     return XDP_PASS;
12 }
```



# Drop IPv6

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      // define variables that represent the packet
5      void *packet_start = (void *)(long)ctx->data;
6      void *packet_end = (void *)(long)ctx->data_end;
7      struct ethhdr *eth = packet_start;
8
9      // satisfy eBPF verifier
10     if (packet_start + sizeof(*eth) > packet_end)
11         return XDP_DROP;
12
13     // find out the next protocol
14     __u16 protocol = eth->h_proto;
15
16     // drop all IPv6 packets
17     if (protocol == bpf_htons(ETH_P_IPV6))
18         return XDP_DROP;
19
20     // pass all other packets
21     return XDP_PASS;
22 }
```

# Drop IPv6

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      // define variables that represent the packet
5      void *packet_start = (void *) (long) ctx->data;
6      void *packet_end = (void *) (long) ctx->data_end;
7      struct ethhdr *eth = packet_start;
8
9      // satisfy eBPF verifier
10     if (packet_start + sizeof(*eth) > packet_end)
11         return XDP_DROP;
12
13     // find out the next protocol
14     __u16 protocol = eth->h_proto;
15
16     // drop all IPv6 packets
17     if (protocol == bpf_htons(ETH_P_IPV6))
18         return XDP_DROP;
19
20     // pass all other packets
21     return XDP_PASS;
22 }
```

# Drop IPv6

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      // define variables that represent the packet
5      void *packet_start = (void *)(long)ctx->data;
6      void *packet_end = (void *)(long)ctx->data_end;
7      struct ethhdr *eth = packet_start;
8
9      // satisfy eBPF verifier
10     if (packet_start + sizeof(*eth) > packet_end)
11         return XDP_DROP;
12
13     // find out the next protocol
14     __u16 protocol = eth->h_proto;
15
16     // drop all IPv6 packets
17     if (protocol == bpf_htons(ETH_P_IPV6))
18         return XDP_DROP;
19
20     // pass all other packets
21     return XDP_PASS;
22 }
```

# Drop IPv6

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      // define variables that represent the packet
5      void *packet_start = (void *) (long) ctx->data;
6      void *packet_end = (void *) (long) ctx->data_end;
7      struct ethhdr *eth = packet_start;
8
9      // satisfy eBPF verifier
10     if (packet_start + sizeof(*eth) > packet_end)
11         return XDP_DROP;
12
13     // find out the next protocol
14     __u16 protocol = eth->h_proto;
15
16     // drop all IPv6 packets
17     if (protocol == bpf_htons(ETH_P_IPV6))
18         return XDP_DROP;
19
20     // pass all other packets
21     return XDP_PASS;
22 }
```



# Drop IPv6

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      // define variables that represent the packet
5      void *packet_start = (void *) (long) ctx->data;
6      void *packet_end = (void *) (long) ctx->data_end;
7      struct ethhdr *eth = packet_start;
8
9      // satisfy eBPF verifier
10     if (packet_start + sizeof(*eth) > packet_end)
11         return XDP_DROP;
12
13     // find out the next protocol
14     __u16 protocol = eth->h_proto;
15
16     // drop all IPv6 packets
17     if (protocol == bpf_htons(ETH_P_IPV6))
18         return XDP_DROP;
19
20     // pass all other packets
21     return XDP_PASS;
22 }
```

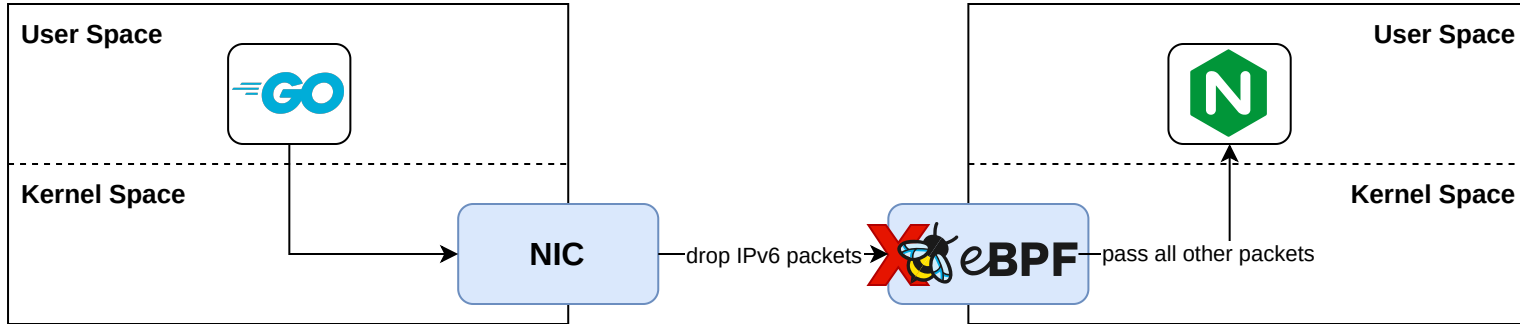
# Drop IPv6

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      // define variables that represent the packet
5      void *packet_start = (void *) (long) ctx->data;
6      void *packet_end = (void *) (long) ctx->data_end;
7      struct ethhdr *eth = packet_start;
8
9      // satisfy eBPF verifier
10     if (packet_start + sizeof(*eth) > packet_end)
11         return XDP_DROP;
12
13     // find out the next protocol
14     __u16 protocol = eth->h_proto;
15
16     // drop all IPv6 packets
17     if (protocol == bpf_htons(ETH_P_IPV6))
18         return XDP_DROP;
19
20     // pass all other packets
21     return XDP_PASS;
22 }
```

# Drop IPv6

```
1  SEC("xdp_kcd_zurich")
2  int xdp_kcd_zurich_firewall(struct xdp_md *ctx)
3  {
4      // define variables that represent the packet
5      void *packet_start = (void *)(long)ctx->data;
6      void *packet_end = (void *)(long)ctx->data_end;
7      struct ethhdr *eth = packet_start;
8
9      // satisfy eBPF verifier
10     if (packet_start + sizeof(*eth) > packet_end)
11         return XDP_DROP;
12
13     // find out the next protocol
14     __u16 protocol = eth->h_proto;
15
16     // drop all IPv6 packets
17     if (protocol == bpf_htons(ETH_P_IPV6))
18         return XDP_DROP;
19
20     // pass all other packets
21     return XDP_PASS;
22 }
```

# Outcome



# Summary

- Event-driven
- Versatile
- Fast
- Secure
- Increasingly more popular

Questions ?