

The eBPF Systems Revolution

Joe Stringer
11-JUNE-2025



ISOVALENT
now part of **cisco**

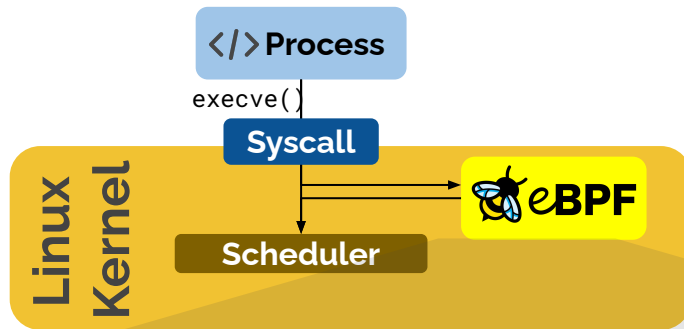
Agenda

- How eBPF changes the pace of innovation
- Real-world Applications
 - Katran Layer 4 Load Balancer (L4LB)
 - Cilium Networking for Kubernetes
- A look at eBPF beyond networking



Makes the OS kernel programmable in a secure and efficient way.

“What JavaScript is to the browser, eBPF is to the Linux Kernel”



```
SEC("fexit/__x64_sys_execve")
int BPF_PROG(fexit_execve, struct filename *name,
             const char *const __argv, const char *const
             __envp)
{
    struct comm_event event = {
        .pid = bpf_get_current_pid_tgid() >>32,
        .type = TYPE_RETURN,
    };

    bpf_get_current_comm(&event.comm, sizeof(event.comm));
    bpf_perf_event_output(ctx, events, BPF_F_CURRENT_CPU,
                        event, sizeof(event));

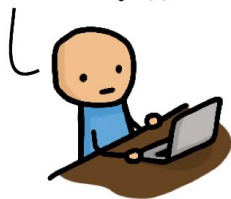
    return 0;
}
```

Why  eBPF ?

Before eBPF

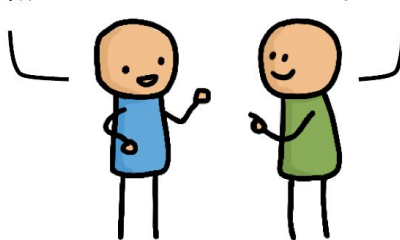
Application Developer:

I want this new feature to observe my app



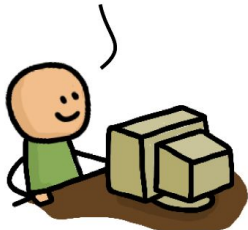
Hey kernel developer! Please add this new feature to the Linux kernel

OK! Just give me a year to convince the entire community that this is good for everyone.

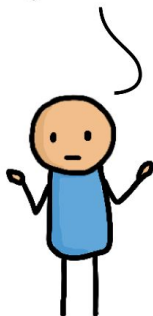


1 year later...

I'm done. The upstream kernel now supports this.



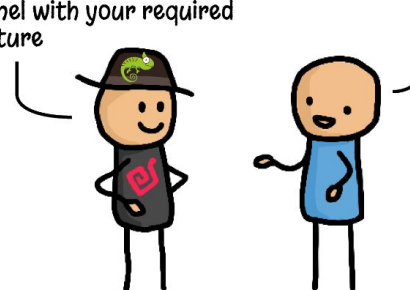
But I need this in my Linux distro



5 year later...

Good news. Our Linux distribution now ships a kernel with your required feature

OK but my requirements have changed since...



After eBPF

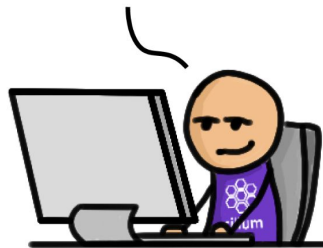
Application Developer:

I want this new feature
to observe my app



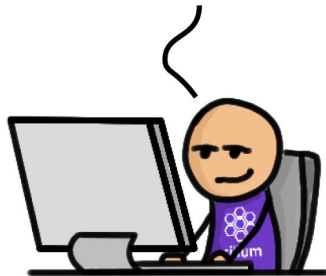
eBPF Developer:

OK! The kernel can't do this so let
me quickly solve this with eBPF.

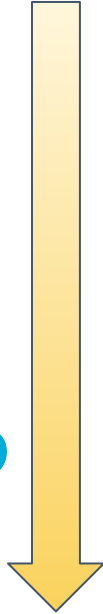
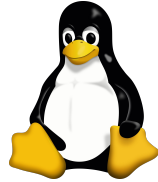
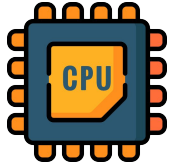


A couple of days later...

Here is a release of our eBPF project that has this feature
now. BTW, you don't have to reboot your machine.



Long Innovation
Cycle



Rapid Innovation
Cycle

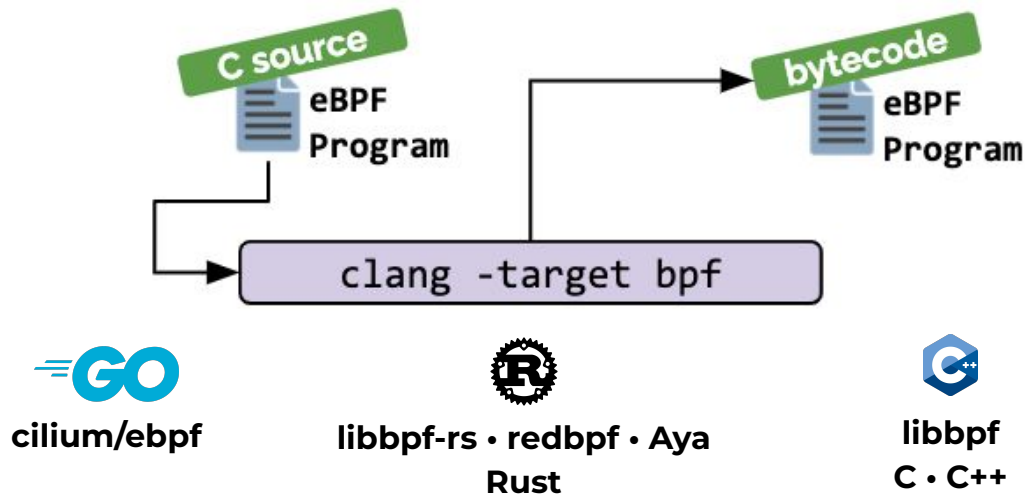
Long innovation cycle results in need to predict use cases or stick to providing building blocks

Programmability allows to continuously adapt to changing requirements and innovate quickly

How does eBPF work?

Language • Runtime

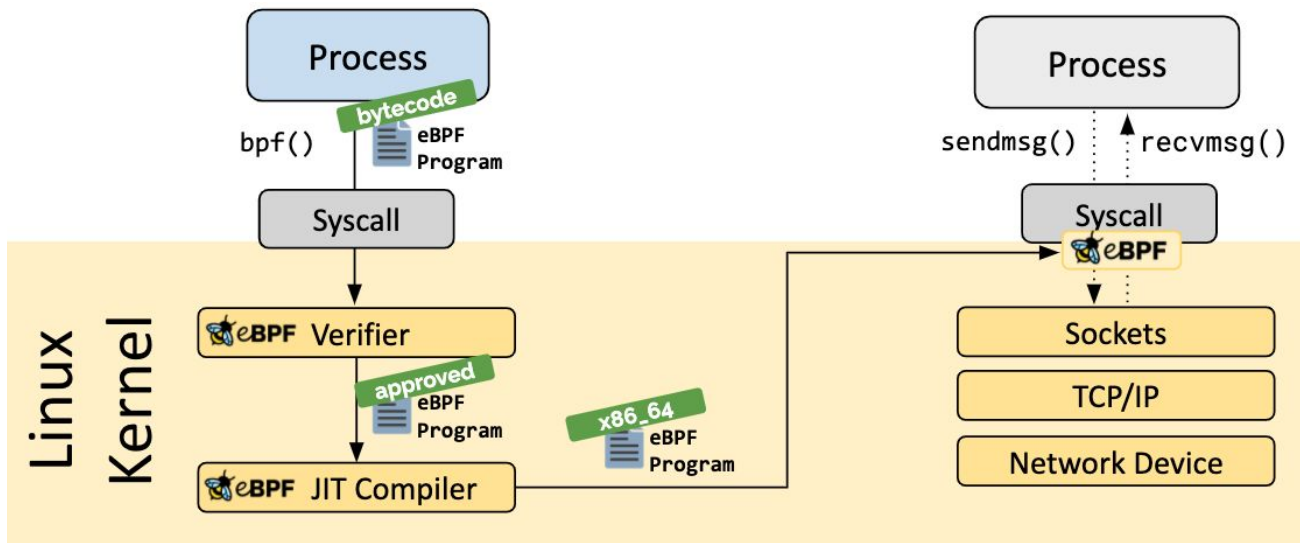
eBPF Language



RFC 9669 BPF Instruction Set Architecture (ISA)

Multiple SDKs and Compilers exist to get to eBPF bytecode

eBPF Runtime



The runtime accepts bytecode, verifies it, just-in-time compiles it, and runs it at the requested hook point.

eBPF Platform

Secure

Runtime
verification
Program Signing



















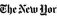
















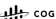
































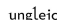

Efficient

JIT compilation
Embedded in OS
Per-CPU data
structures

Portable

Generic Bytecode
Data Type
Discovery
Stable API to OS

Where is eBPF used today?

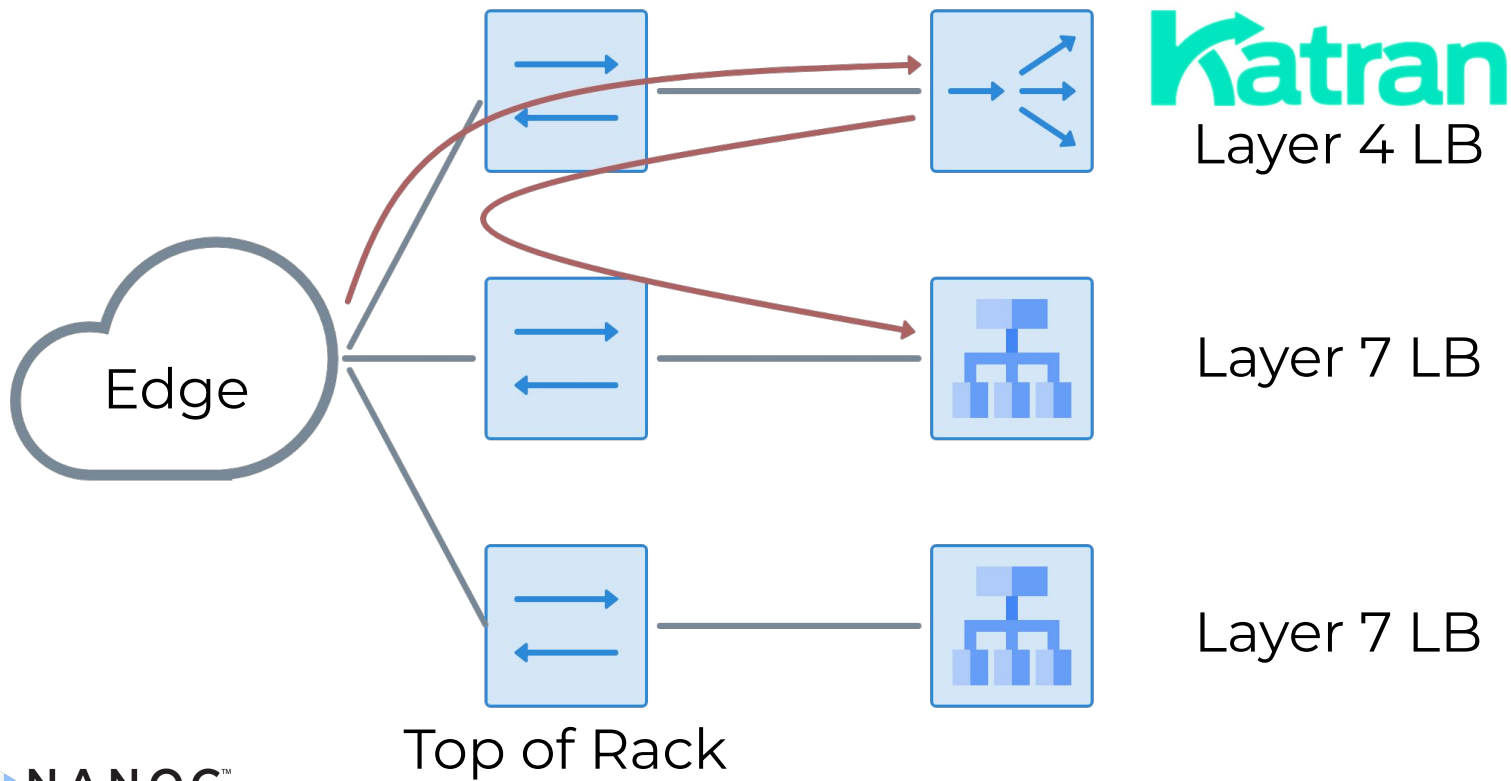
 <p>Adobe</p> <p>What Makes a Good Multi-Tenant Kubernetes Solution</p> <p>VIDEO 1 - VIDEO 2</p>	 <p>Alibaba Cloud</p> <p>Building High-Performance Cloud Native Pod Networks</p> <p>READ BLOG</p>	 <p>aws</p> <p>AWS picks Cilium for Networking & Security on EKS Anywhere</p> <p>READ BLOG</p>	 <p>Bell</p> <p>Bell uses Cilium and eBPF for telco networking</p> <p>VIDEO 1 - VIDEO 2</p>	 <p>nexiot</p> <p>Nexiot using Cilium as the CNI plugin on EKS for its IoT SaaS</p> <p>READ USER STORY</p>	 <p>PostFinance</p> <p>PostFinance is using Cilium as their CNI for all mission critical, on premise k8s clusters</p> <p>CASE STUDY - VIDEO</p>	 <p>S&P Global</p> <p>S&P Global uses Cilium as its CNI</p> <p>WATCH VIDEO</p>	 <p>SEZNAM.CZ</p> <p>Seznam.cz uses Cilium for Layer 4 Load Balancing with XDP</p> <p>BLOG - VIDEO</p>	 <p>Switch</p> <p>Uswitch uses Cilium ClusterMesh for multi cluster networking</p> <p>BLOG - VIDEO</p>	 <p>Utmost</p> <p>Utmost is using Cilium in all tiers of its Kubernetes ecosystem to implement zero trust networking</p> <p>USER STORY - BLOG</p>	 <p>VSHN</p> <p>VSHN uses Cilium for multi-tenant networking on APPIUO Cloud</p> <p>READ CASE STUDY</p>	 <p>VZP</p> <p>Building a Global Multi-Cluster Networking Infrastructure with Cilium</p> <p>BLOG - VIDEO</p>
 <p>Capital One</p> <p>Building a Secure and Maintainable PaaS</p> <p>WATCH VIDEO</p>	 <p>GEGEN</p> <p>Cloud Native Networking with eBPF</p> <p>WATCH VIDEO</p>	 <p>DATADOG</p> <p>Datadog is using Cilium in AWS (self-hosted k8s)</p> <p>VIDEO - USER STORY</p>	 <p>DigitalOcean</p> <p>Managed Kubernetes: 1.5 Years of Cilium Usage at DigitalOcean</p> <p>WATCH VIDEO</p>	 <p>sky</p> <p>eBPF & Cilium at Sky</p> <p>WATCH VIDEO</p>	 <p>sky BET</p> <p>Skybet uses Cilium as their CNI</p> <p>READ BLOG</p>	 <p>The New York Times</p> <p>The New York Times uses Cilium on EKS for multi-region multi-tenant shared clusters</p> <p>READ USER STORY</p>	 <p>Trip.com</p> <p>Trip.com uses Cilium both on-premise and in AWS</p> <p>BLOG 1 - BLOG 2 - VIDEO</p>	 <p>ACCUKNOX</p> <p>Accuknox uses Cilium for network visibility and network policy enforcement</p>	 <p>ACOSS</p> <p>Acoss uses Cilium as their main CNI plugin for self-hosted Kubernetes</p>	 <p>ArangoDB</p> <p>ArangoDB Oasis uses Cilium to separate database deployments in a multi-tenant cloud environment</p>	 <p>ayedo</p> <p>Ayedo builds and operates cloud native platforms using Cilium</p>
 <p>证通股份</p> <p>ec2BB uses Cilium as their CNI and for load balancing</p> <p>READ BLOG</p>	 <p>FORM3</p> <p>Form3 is using Cilium in their production clusters (self-hosted, bare-metal, private cloud)</p> <p>WATCH VIDEO</p>	 <p>GitLab</p> <p>Kubernetes Network Policies in Action with Cilium</p> <p>VIDEO</p>	 <p>Google</p> <p>Google chooses Cilium for Google Kubernetes Engine (GKE) networking</p> <p>BLOG - VIDEO</p>	 <p>Melenion Inc</p> <p>Melenion uses Cilium as the CNI for its on-premise production clusters</p>	 <p>MUX</p> <p>Mux uses Cilium on self-hosted clusters in GCP and AWS to run its video streaming/analytics platform</p>	 <p>myfitnesspal</p> <p>MyFitnessPal trusts Cilium with high volume user traffic on AWS and GKE</p>	 <p>Northflank</p> <p>Northflank uses Cilium as its CNI plugin across GCP, Azure, AWS and bare metal</p>	 <p>ByteDance</p> <p>ByteDance uses Cilium as their CNI for self-hosted Kubernetes clusters</p>	 <p>CANONICAL</p> <p>Canonical's Kubernetes distribution microk8s uses Cilium as CNI plugin</p>	 <p>CIVO</p> <p>Civo is offering Cilium as the CNI option for Civo users to choose if for their Civo Kubernetes clusters</p>	 <p>COGNITE</p> <p>Cognite uses Cilium as the CNI plugin for industrial DataOps</p>
 <p>cilium</p> <p>Homes furnishings retailer uses Cilium for their self-hosted bare-metal private cloud</p> <p>WATCH VIDEO</p>	 <p>MÁSMÓVIL</p> <p>Scaling a Multi-Tenant Kubernetes Clusters in a Telco</p> <p>WATCH VIDEO</p>	 <p>Meltwater</p> <p>Meltwater is using Cilium in AWS on self-hosted multi-tenant k8s clusters as the CNI plugin</p> <p>WATCH VIDEO</p>	 <p>MOBILAB</p> <p>Mobilabs uses Cilium as the CNI for their internal cloud</p> <p>READ BLOG</p>	 <p>overstock.com</p> <p>Overstock uses Cilium as their CNI for self-hosted bare metal clusters</p>	 <p>Palantir</p> <p>Palantir is using Cilium as their main CNI plugin in AWS (self-hosted k8s)</p>	 <p>PLAID</p> <p>Plaid uses Cilium as the CNI for its serverless database platform</p>	 <p>PlanetScale</p> <p>PlanetScale uses Cilium as their CNI plugin in self-hosted Kubernetes on AWS</p>	 <p>radiofrance</p> <p>Radio France uses Cilium on AWS</p>	 <p>rappuuta robotics</p> <p>Rappuuta Robotics uses Cilium as their main CNI plugin for self-hosted clusters</p>	 <p>SAP</p> <p>SAP uses Cilium for projects across AWS, Azure, GCP, and OpenStack</p>	 <p>Sapian</p> <p>Sapian uses Cilium as the default CNI in their product DataBox Cloud for low latency in real-time communications environments on the edge</p>
 <p>Snapp!</p> <p>Snapp uses Cilium for its on-premise OpenShift clusters</p>	 <p>solo.io</p> <p>Cilium is part of Solo.io's Gloo Application Networking platform</p>	 <p>SPHERITY</p> <p>Spherity uses Cilium on AWS EKS for CNI, Hubble, and network policies</p>	 <p>sportradar</p> <p>Sportradar is using Cilium as their main CNI plugin in AWS (using topology)</p>	 <p>ISOVALENT</p> <p>Cilium is the platform that powers Isovalent's enterprise networking, observability, and security solutions</p>	 <p>JUMO</p> <p>JUMO uses Cilium as the CNI plugin for all of their AWS-hosted EKS clusters</p>	 <p>KRYPTOS LOGIC</p> <p>Kryptos uses Cilium as the CNI for their on-prem Kubernetes clusters</p>	 <p>Kube-OVN</p> <p>Kube-OVN uses Cilium to enhance the CNI service performance, security and monitoring</p>	 <p>Scaleway</p> <p>Scaleway uses Cilium as the default CNI for Kubernetes Kapsule</p>	 <p>SCHUBERG PHILIS</p> <p>Schuberg Philis uses Cilium as the CNI for mission critical Kubernetes clusters they run for their customers</p>	 <p>SIMPLE</p> <p>Simple uses Cilium as default CNI for EKS</p>	 <p>smile</p> <p>SmileDirectClub uses Cilium in self-hosted clusters vSphere and EC2 for manufacturing</p>
 <p>sproutfi</p> <p>Sproutfi uses Cilium as the CNI on its GKE based clusters</p>	 <p>SUPERORBITAL</p> <p>Superorbital uses Cilium in their customer engagements</p>	 <p>TAILOR BRANDS</p> <p>Tailor Brands uses Cilium in their EKS clusters</p>	 <p>T Systems</p> <p>TSI uses Cilium for it's Open Sovereign Cloud product</p>	 <p>KUBERMATIC</p> <p>Kubermatic uses Cilium as the CNI for their Kubernetes installer and platform</p>	 <p>KUBESPHERE</p> <p>KubeSphere is an open-source lightweight tool for deploying Kubernetes clusters and addons</p>	 <p>REPLY</p> <p>Liquid Reply is a consulting firm that uses Cilium in client projects</p>	 <p>MagicLeap</p> <p>Magic Leap uses Hubble for observability</p>	 <p>ungleich</p> <p>ungleich uses Cilium for IPv6-only Kubernetes deployments</p>	 <p>yahoo!</p> <p>Yahoo is using Cilium for L4 North-South Load Balancing for Kubernetes Services</p>		



Real World Applications

Networking in the Data Center

Katran: Layer 4 Load Balancer



Load Balancing Challenges

- Performance: High CPU with IP Virtual Server (IPVS)
 - So, run dedicated instances (new deployment challenges)
 - Difficult to manage a fair balance of Layer 4, Layer 7 LB
- Availability
 - Adjusting for incident or planned outage
- Management
 - Kernel changes to adjust algorithm

How eBPF enabled Katran

Performance

Reduce the cost per packet by using XDP

Deploy both L4 and L7 on the same hosts

Availability

Explore newer algorithms, mix of algorithms

Consistent Hashing (Maglev)

Management

Dynamically update without reboot

Iterate more quickly on solutions

XDP: eXpress Data Plane

- Programmable layer in the network driver
 - Inspired by Intel Data Plane Development Kit (DPDK)
 - Minimize the CPU instructions executed per packet
 - Avoid unnecessary memory allocations, copies
- Don't assume all packets reach a local application
 - Default path: Pass through
 - Support one-armed routers, firewall, denial of service
- Define network behavior using eBPF

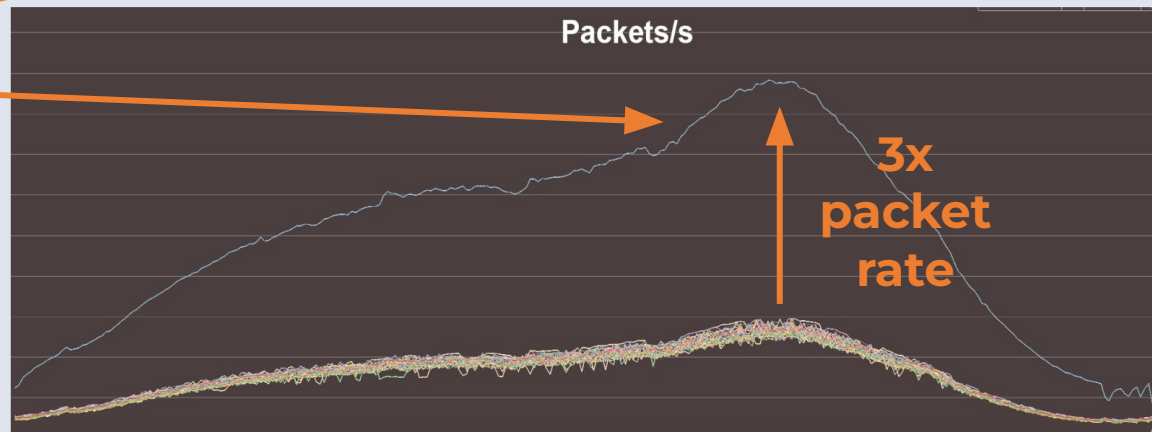
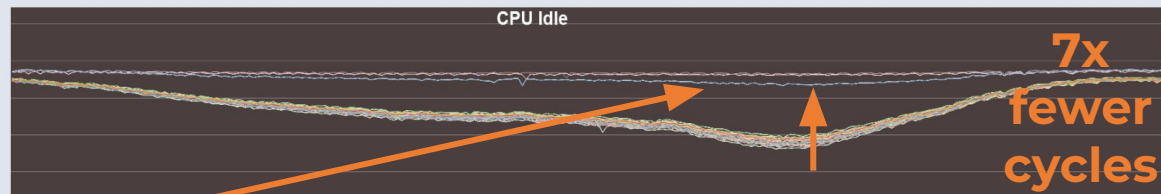
Load Balancing at Meta

Baseline:

- IPVS

Katran:

- eBPF runtime
- XDP attach



Cilium

- **Networking & Load-Balancing**

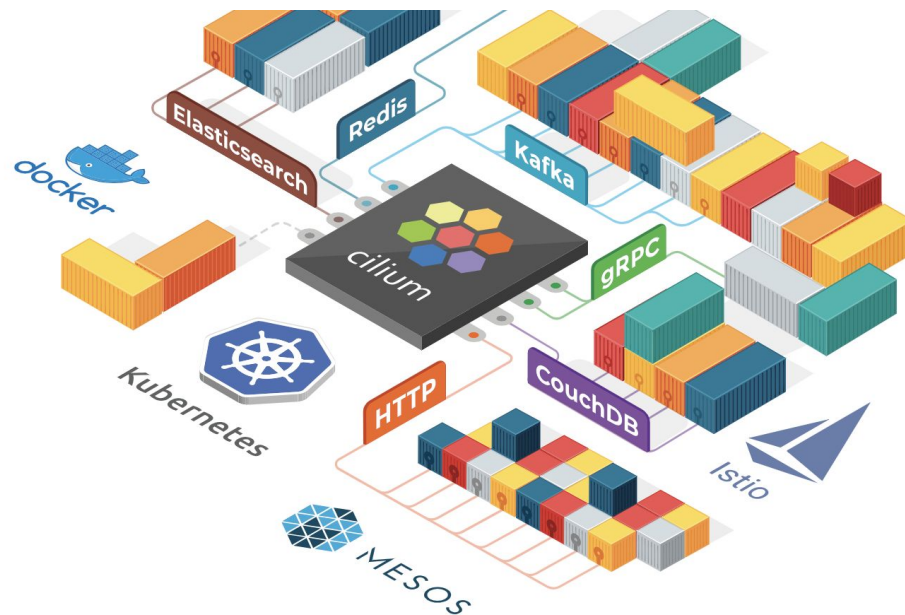
- Container Network Interface (CNI)
- Kubernetes Services
- Multi-cluster
- Virtual Machine Gateway

- **Network Security**

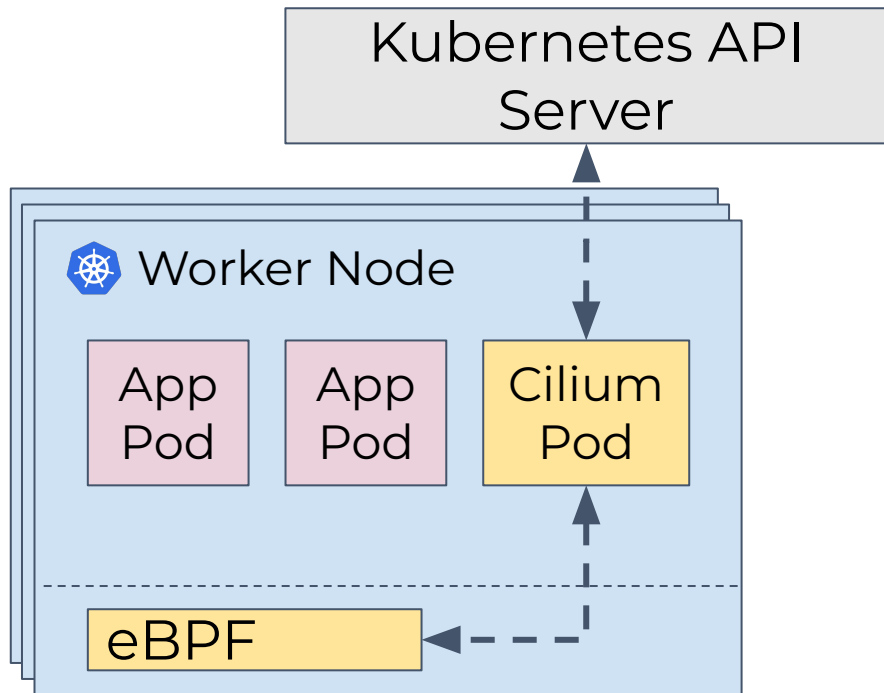
- Network Policy
- Identity-based
- Encryption

- **Observability**

- Metrics
- Flow Visibility
- Service Dependency



Networking for Cloud Native Apps



Declarative Intent

- Nodes, Pods
- Network Policies
- Services, Endpoints

Synchronize intent to workers

- Connectivity (IPAM, routing)
- External connectivity
- Service-based Load Balancing
- Identity-aware Network policy
- Flow visibility & metrics
- Transparent Encryption
- Multi-cluster Routing & Security

Rough Numbers

- x0 deployments per minute
- x00 workloads per node
- x,000 nodes per cluster
- x0,000 services
- x00,000 containers

Optimizing Linux Networking

Load Balancing

Avoid linear iteration for Virtual IP resolution

Per connection rather than per packet LB

Routing

Skip costly upper stack operations

Query native Linux routing verdict from kernel tables

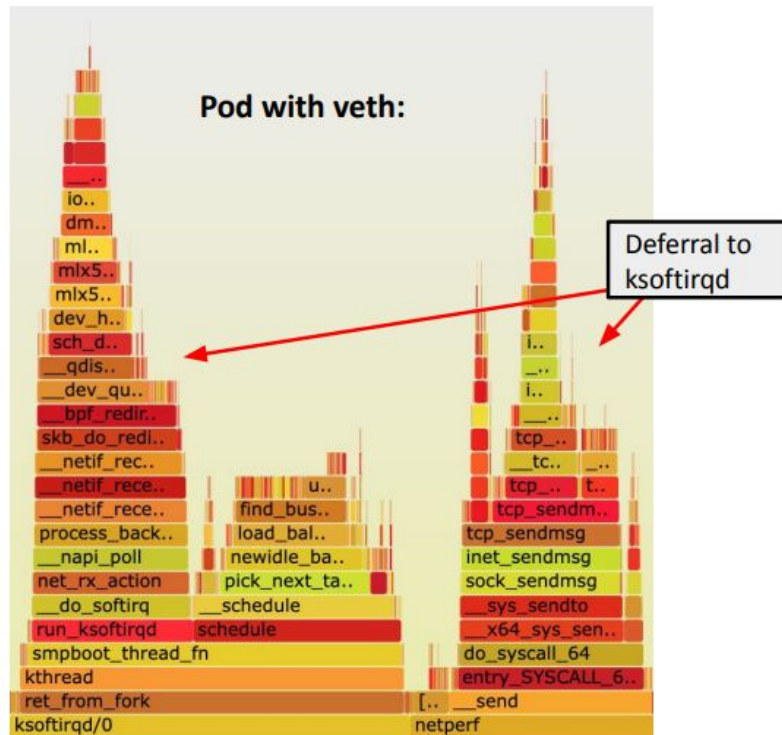
Scheduling

Optimize packet handover between logical networks

Participate in packet departure scheduling (EDT)

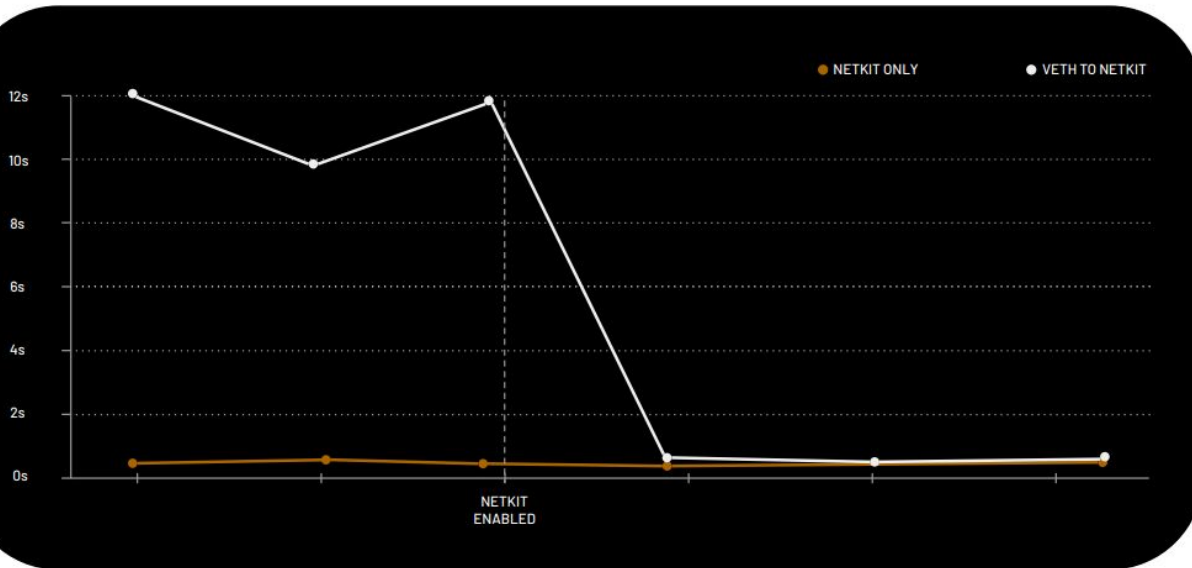
Limitations in Virtual Ethernet

- Observation: Containers slower than the base host
 - Inefficient scheduling can have a big impact
 - What if we try a "run to completion" model?
 - Combine with other optimizations - Early Departure Time (EDT) scheduling; BIG TCP etc.



Netkit driver

- 100GiB single stream TCP @ 8K MTU
- Meta: P99 Reduction from 12s to 0.1s



Remains in process context all the way, leading to better process scheduler decisions.



eBPF Beyond Networking

Active areas of development

eBPF Use Cases

Microsoft Proposes "Hornet" Security Module For The Linux Kernel

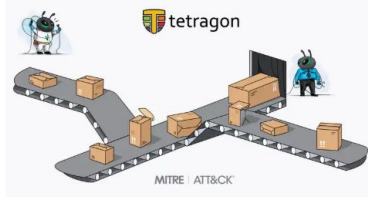
Written by [Michael Larabel](#) in [Microsoft](#) on 21 March 2025 at 02:21 PM EDT. [28 Comments](#)



Microsoft's newest open-source contribution to the Linux kernel being proposed is... Hornet, a Linux security module (LSM) for providing signature verification of eBPF programs.

Using sched_ext to improve frame rates on the SteamDeck

Ideas behind the LAVD scheduler

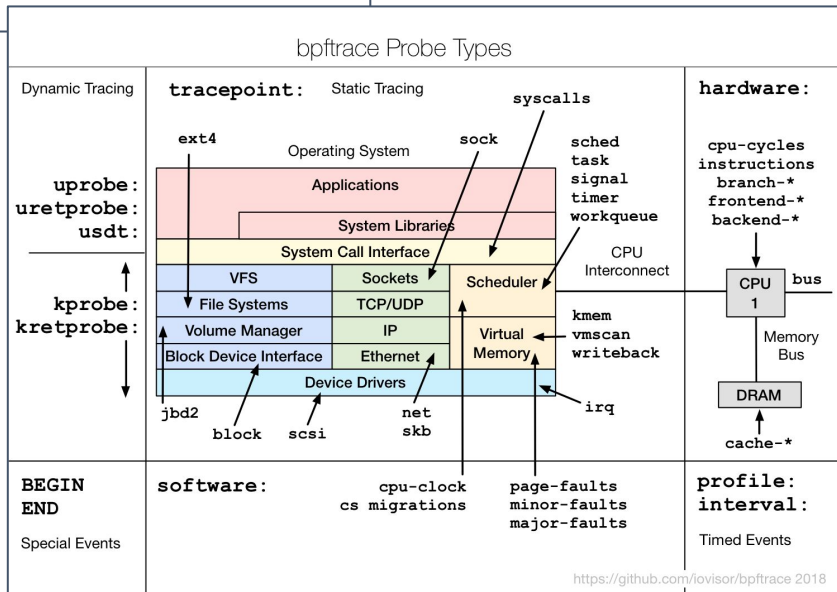


Tetragon · Dec 12, 2024

Telemetry to Tactics: Tetragon Through the Lens of the MITRE ATT&CK Framework

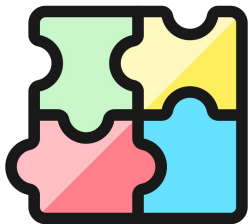
Use the MITRE ATT&CK framework to transform the deep telemetry from Tetragon into clear, actionable insights about adversary behavior.

By Paul Arah



<https://github.com/iovisor/bpftrace> 2018

Cross-Platform Standardization



eBPF for Windows

I E T F Datatracker Groups Documents Meetings Other User		
BPF/eBPF (bpf)		
About Documents Meetings History Photos Email expansions List archive »		
WG	Name	BPF/eBPF
	Acronym	bpf
	Area	Internet Area (int)
	State	Active
	Charter	charter-ietf-bpf-01 Approved
	Document dependencies	Show
	Additional resources	GitHub Organization
Personnel	Chairs	David Vernet , Suresh Krishnan
	Area Director	Erik Kline
	Tech Advisors	Alexei Starovoitov , Christoph Hellwig , Dave Thaler
Mailing list	Address	bpf@ietf.org
	To subscribe	https://www.ietf.org/mailman/listinfo/bpf
	Archive	https://mailarchive.ietf.org/arch/browse/bpf/
Chat	Room address	https://zulip.ietf.org/#narrow/stream/bpf

<https://datatracker.ietf.org/wg/bpf/about>

eBPF Foundation

Platinum



Silver



eBPF Foundation Announces \$250,000 in Grant Awards for Five eBPF Academic Research Projects

By Dan Brown | August 29, 2024 | 7 min read

Projects will advance eBPF's open source technology by improving scalability, static analysis, verifier, virtual memory and more

eBPF is unlocking systems innovation



ISOVALENT



Get Started Today

ebpf.io

ebpf.io/slack