

# Recent Linux Kernel Networking improvements

Hannes Frederic Sowa  
<[hannes@stressinduktion.org](mailto:hannes@stressinduktion.org)>  
<[hannes@redhat.com](mailto:hannes@redhat.com)>  
Red Hat Inc.

# Types of Performance Improvements

- Reducing cycles in common code paths
- New features enabling the kernel to handle certain networking scenarios more intelligent/faster

# “Recent” Networking performance improvements

- dma\_rmb, dma\_wmb
- fib\_trie cleanup
- napi\_alloc\_skb
- tx-nocache-copy
- xmit\_more
- TCP improvements
- Non-Vlan offloaded GRO
- AF\_UNIX zerocopy
- IPv6 dst\_entry caching improvements
- iptables rules per NUMA-node
- unconditional netfilter hook registration in bridge

# Future

- ECMP for IPv4, aligned with IPv6
- Low level optimizations
  - improvements in sk\_buff memory allocation
  - memcpy/memset
- TCP friends
- using for xmit\_more in forwarding