

云智未来⁹th

第九届中国系统架构师大会
SYSTEM ARCHITECT CONFERENCE CHINA 2017

全用户态服务开发套件F-Stack

演讲者：姜凤波

SACC
2017

北京·新云南皇冠假日酒店

IT168.com

ChinaUnix

ITPUB


Explosive growth of data services

▶ Challenges


- ▶ service traffic:CDN, live streaming(RTMP),etc.
- ▶ DDoS, CC

▶ Solution


- ▶ NIC: 25GbE → 40GbE → 100GbE
- ▶ CPU: 24Cores → 48Cores → 56Cores
- ▶ L2/L3: Kernel bypass(DPDK)
- ▶ L4/L7: ?



The problem - Kernel



Kernel Bypass

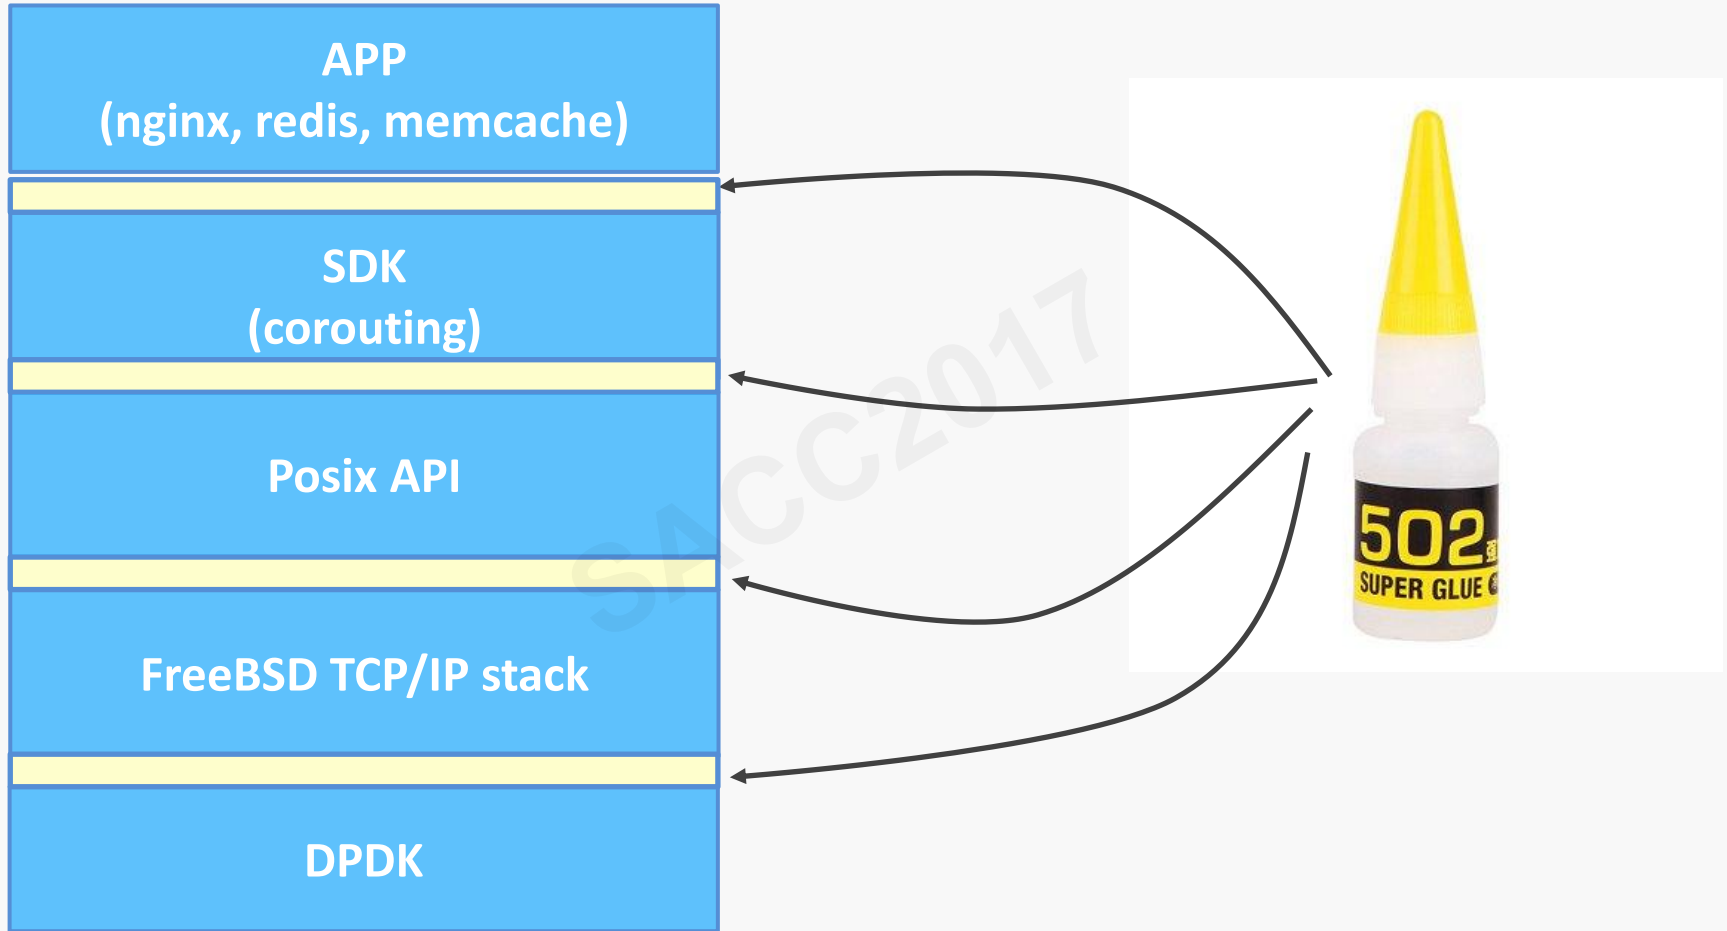


User space TCP/IP stack

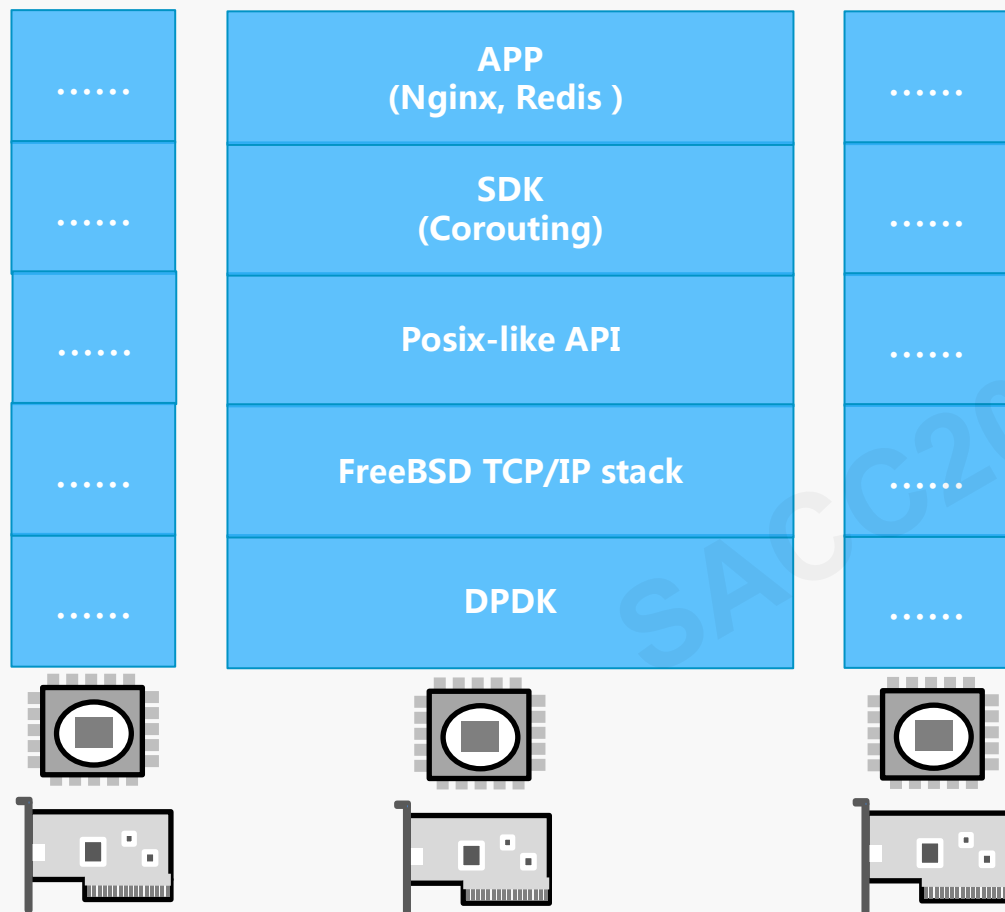
- ▶ Seastar's native stack
 - ▶ problems in WAN
 - ▶ incompatible with existing programs
- ▶ mTCP's user-level stack
 - ▶ tcp only.
 - ▶ too simple to be a production ready stack

VLAN/VxLAN/Tunnel/Bonding/Network Tools are not supported

What does F-Stack do?



F-Stack



1、 Full user-space

- ▶ No content switch
- ▶ Zero copy
- ▶ No hirqs and sirqs

2、 Shared-nothing architecture

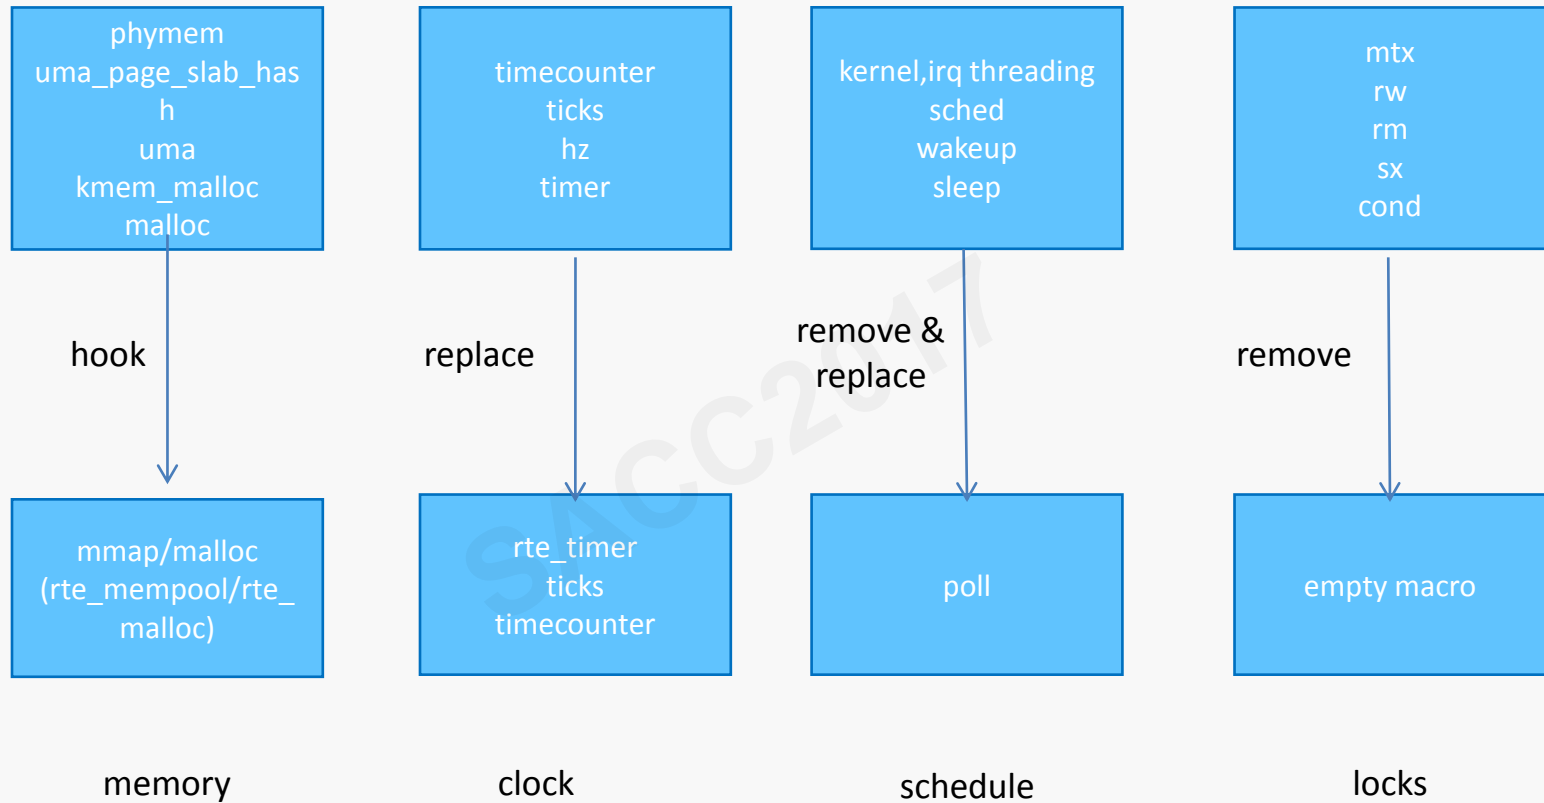
- ▶ Linear scalability
- ▶ No schedule
- ▶ No lock
- ▶ No cache locality miss

Why FreeBSD?

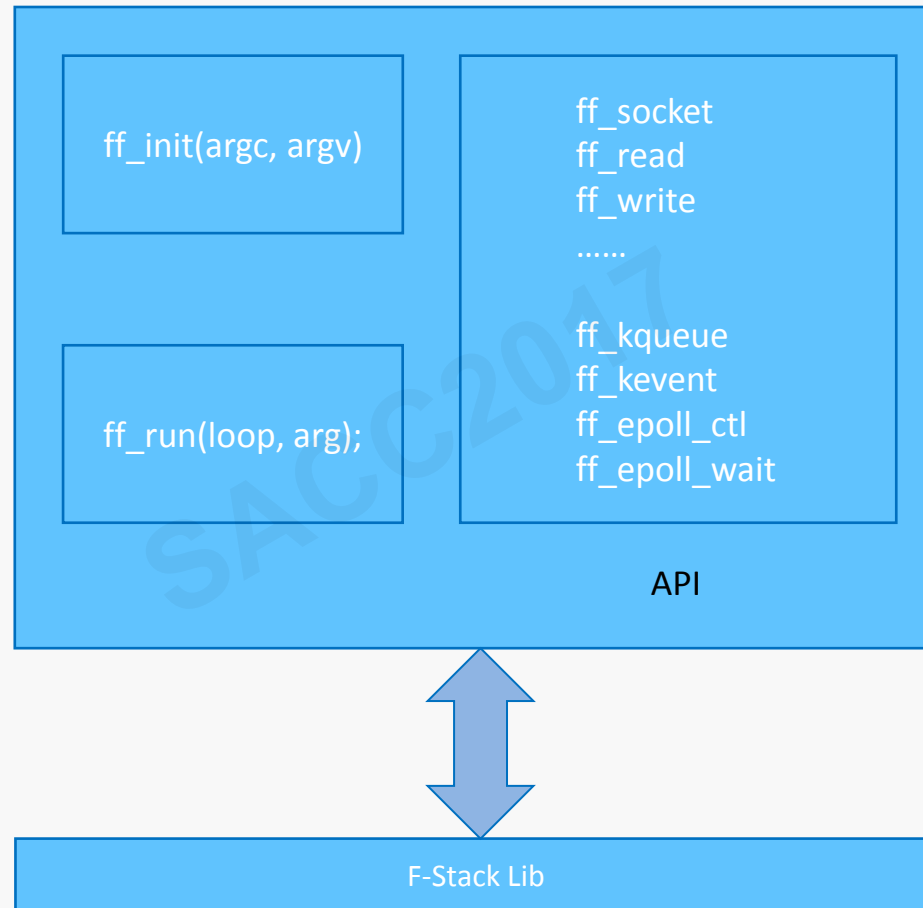
- ▶ Why full TCP/IP stack ?
 - ▶ Advantageous functional, production ready stack, IGMP/IPv6/VLAN/VxLAN/Tunnel/Bonding/Net Tools
 - ▶ Stable.
- ▶ Why not Linux ?
 - ▶ complicated logic.
 - ▶ GPL

SACC2017

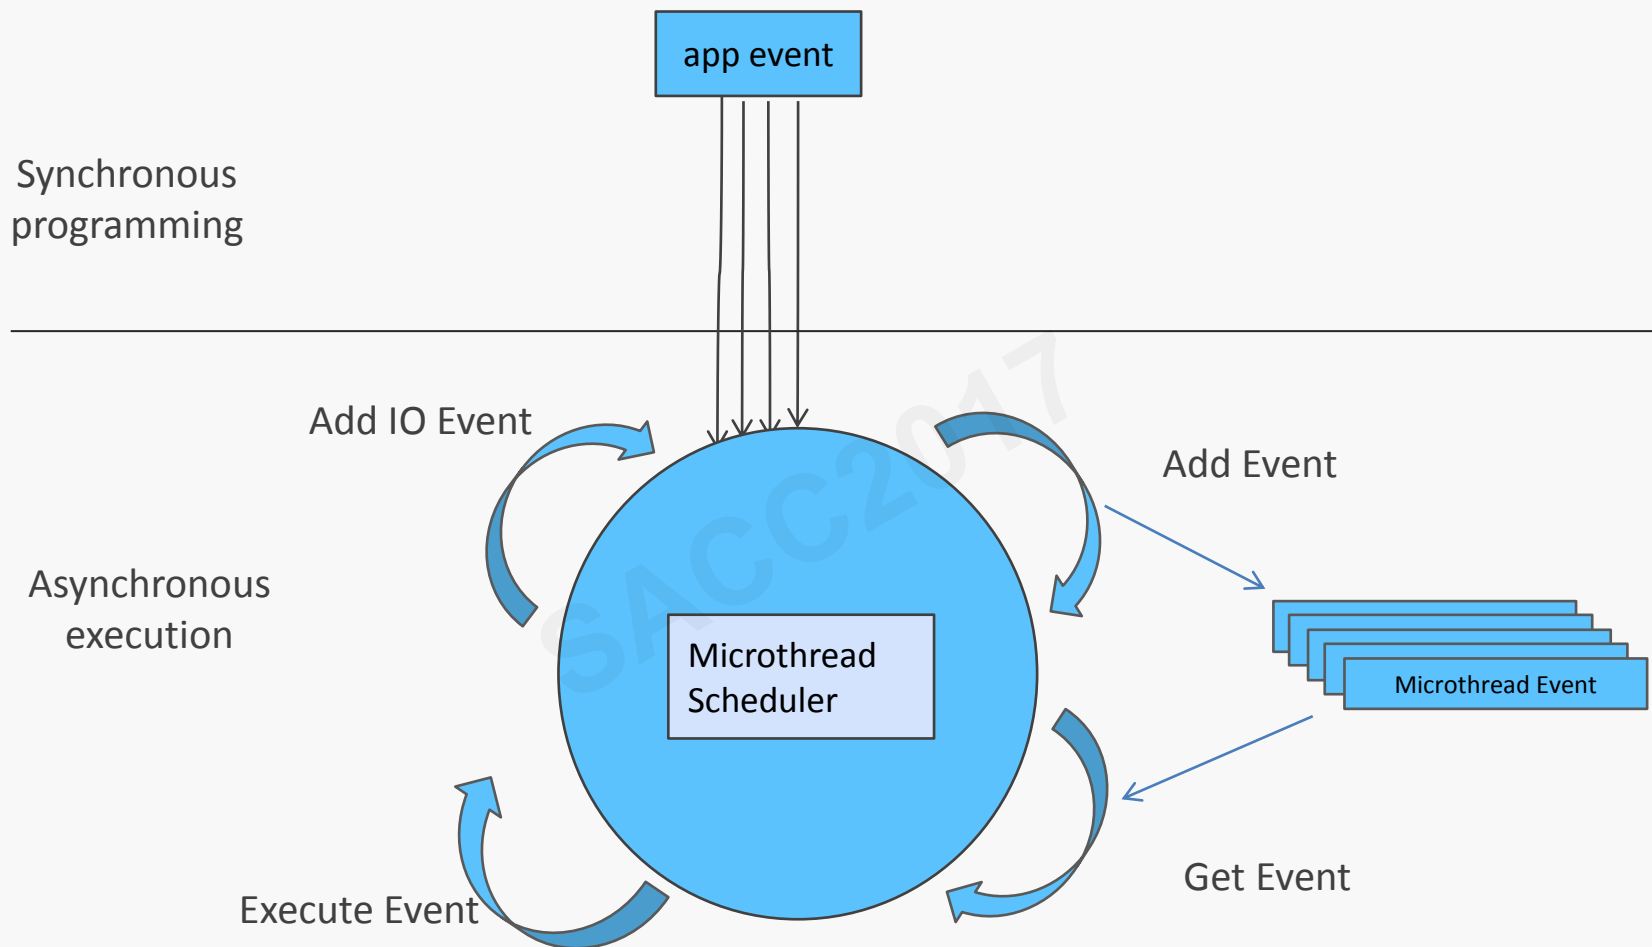
User space FreeBSD TCP/IP stack



Posix-like API



Coroutine



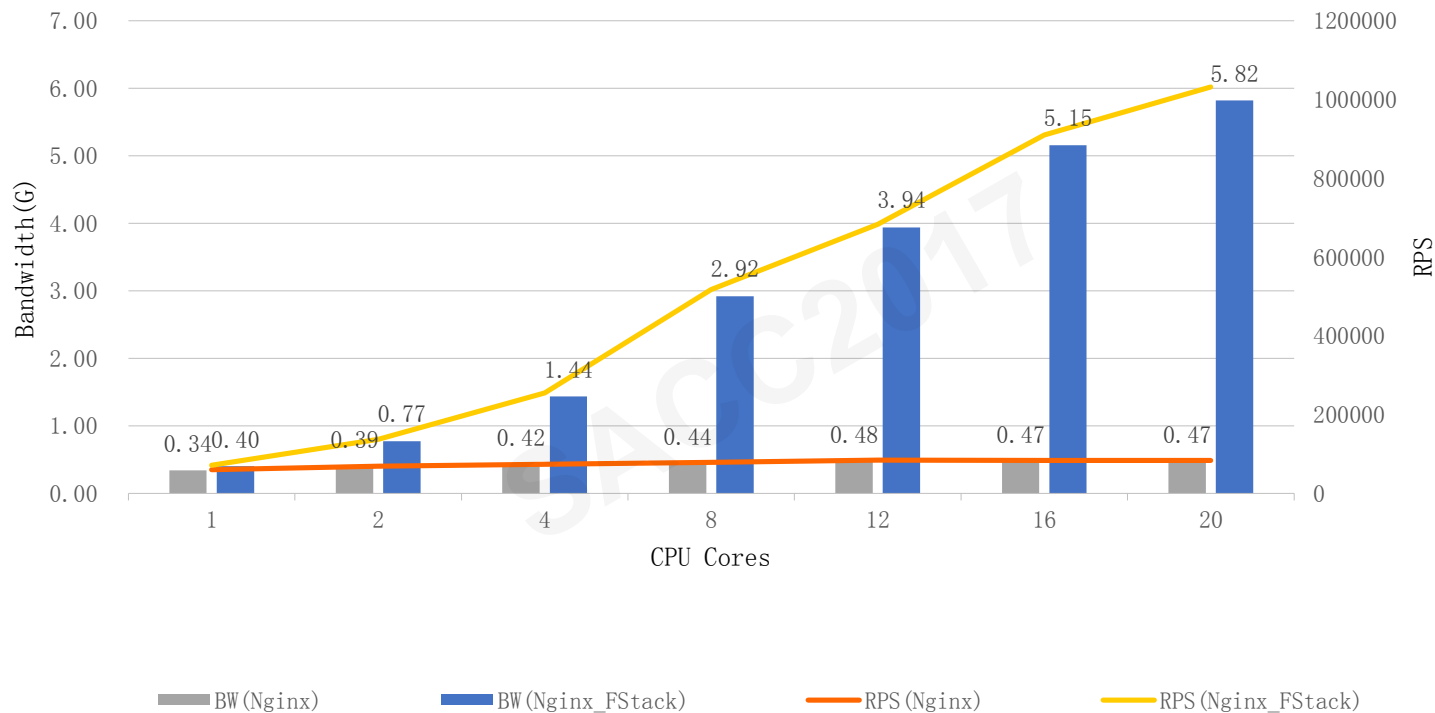
Tools

- ▶ sysctl
- ▶ ifconfig
- ▶ route
- ▶ netstat
- ▶ top
- ▶ etc..

SACC2017

Nginx_F-Stack

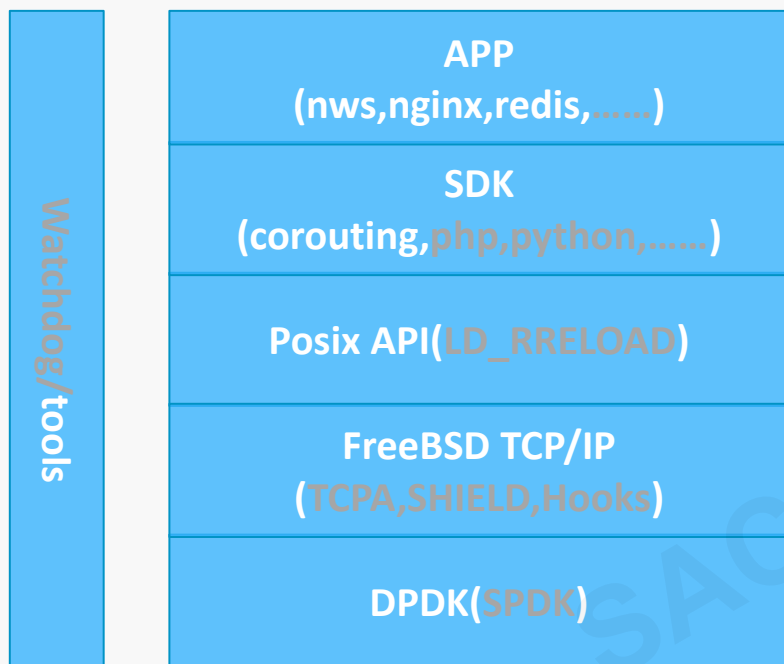
cache body RPS test(No Keep Alive)



F-Stack in Tencent Cloud

- ▶ 2013 UDP DNS server based on DPDK
- ▶ 2014 TCP(user-space TCP/IP stack) DNS server
- ▶ 2015 L4/L7 gateway application
- ▶ 2016 Used on HttpDNS,CDN,COS(Cloud Object Storage)
- ▶ 2017 Port FreeBSD network stack to F-Stack and open source

Roadmap





<https://github.com/f-stack>



<http://www.f-stack.org>



F-Stack WeChat subscription



THANKS

The background features a dark, almost black, space filled with numerous small, bright blue particles. These particles are arranged in several distinct, curved paths that sweep across the frame from the bottom left towards the top right. A bright, white-to-yellow light source is positioned near the center of the image, slightly to the left of the word 'THANKS', creating a strong lens flare effect that illuminates the surrounding particles and the text.