

# OpenResty XRay 简介

对您的在线应用进行全天候、无死角的分析 and 诊断



试用 OpenResty XRay 产品

<https://openresty.com.cn/cn/xray/>

info@openresty.com

# 软件世界的挑战

业务快速迭代

团队成员水平不一

测试不充分

业务复杂度不断提高

CPU 资源占用过高

内存资源占用过多 (含内存泄漏)

硬盘 IO 资源不足

存在长延时响应

很难线下复现的异常和错误 (含进程崩溃)

# K8s/Docker 容器时代 的挑战

---

很多容器，很多应用，很多发行版，很多技术栈

---

容器最小化，缺乏最基本的调试工具

---

容器权限集合最小化

---

出问题时自动丢弃和重启容器，软件 Bug 容易被掩盖

---

容器虚拟化、微服务——软件复杂度进一步提高

## 传统方法的缺点

侵入式需要修改应用

响应需求慢

需要大数据存储和分析

仅限表面指标

只有现象，没有原因

缺少深度全技术栈分析和诊断

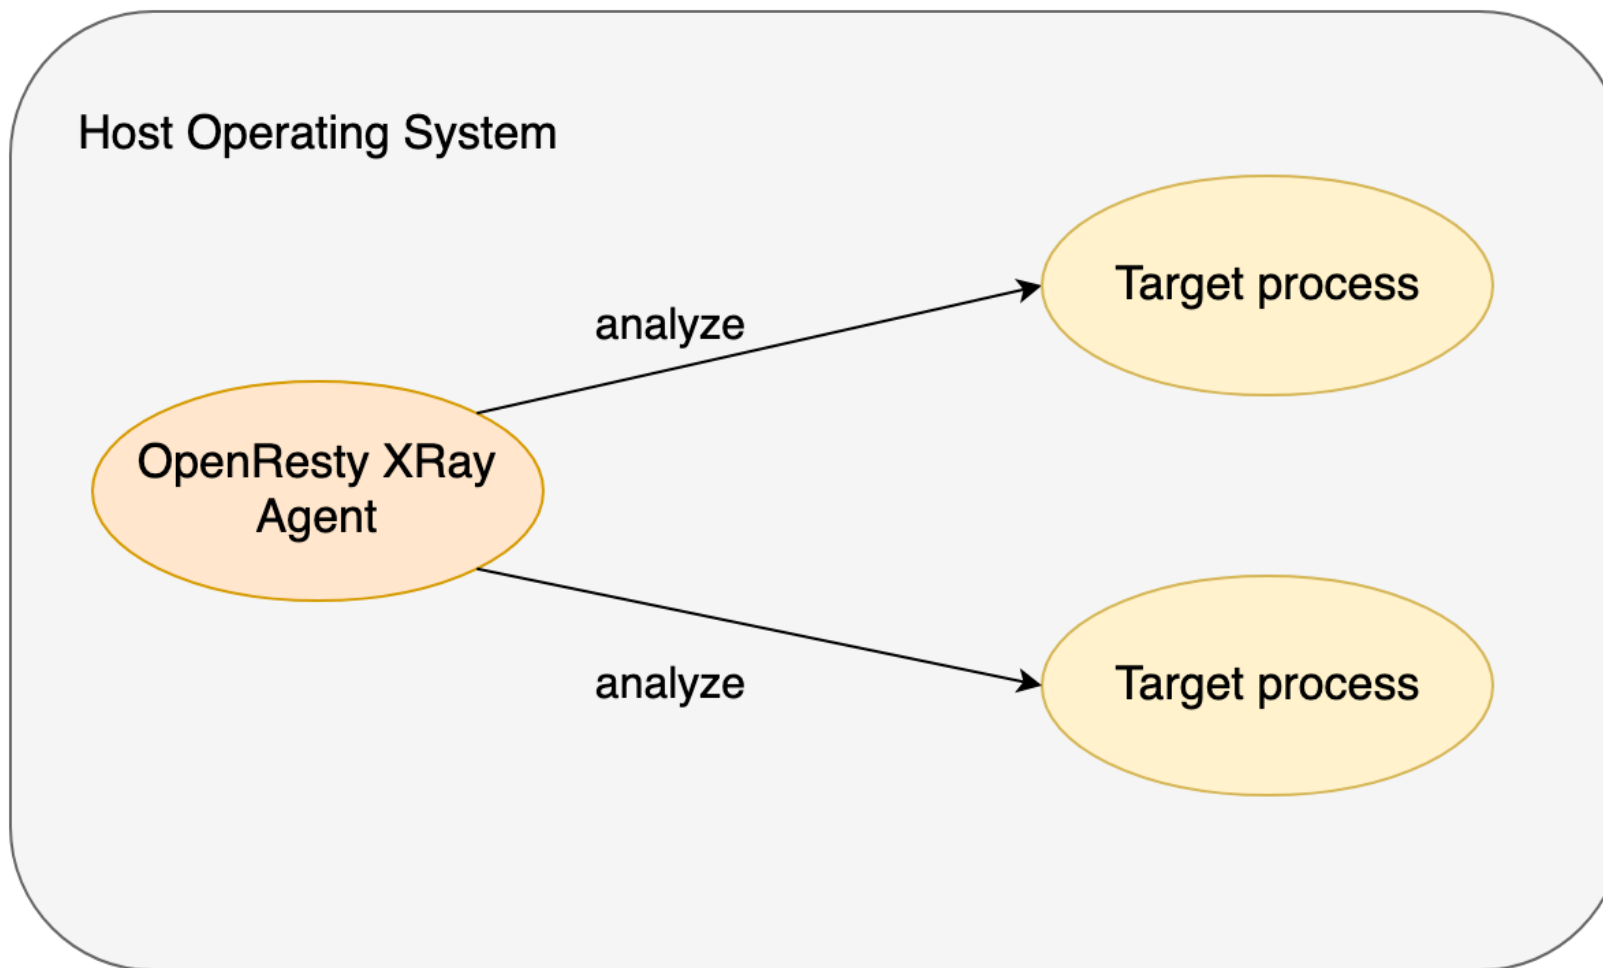
数据采集和处理流程复杂，开销大，易出错

# OpenResty XRay

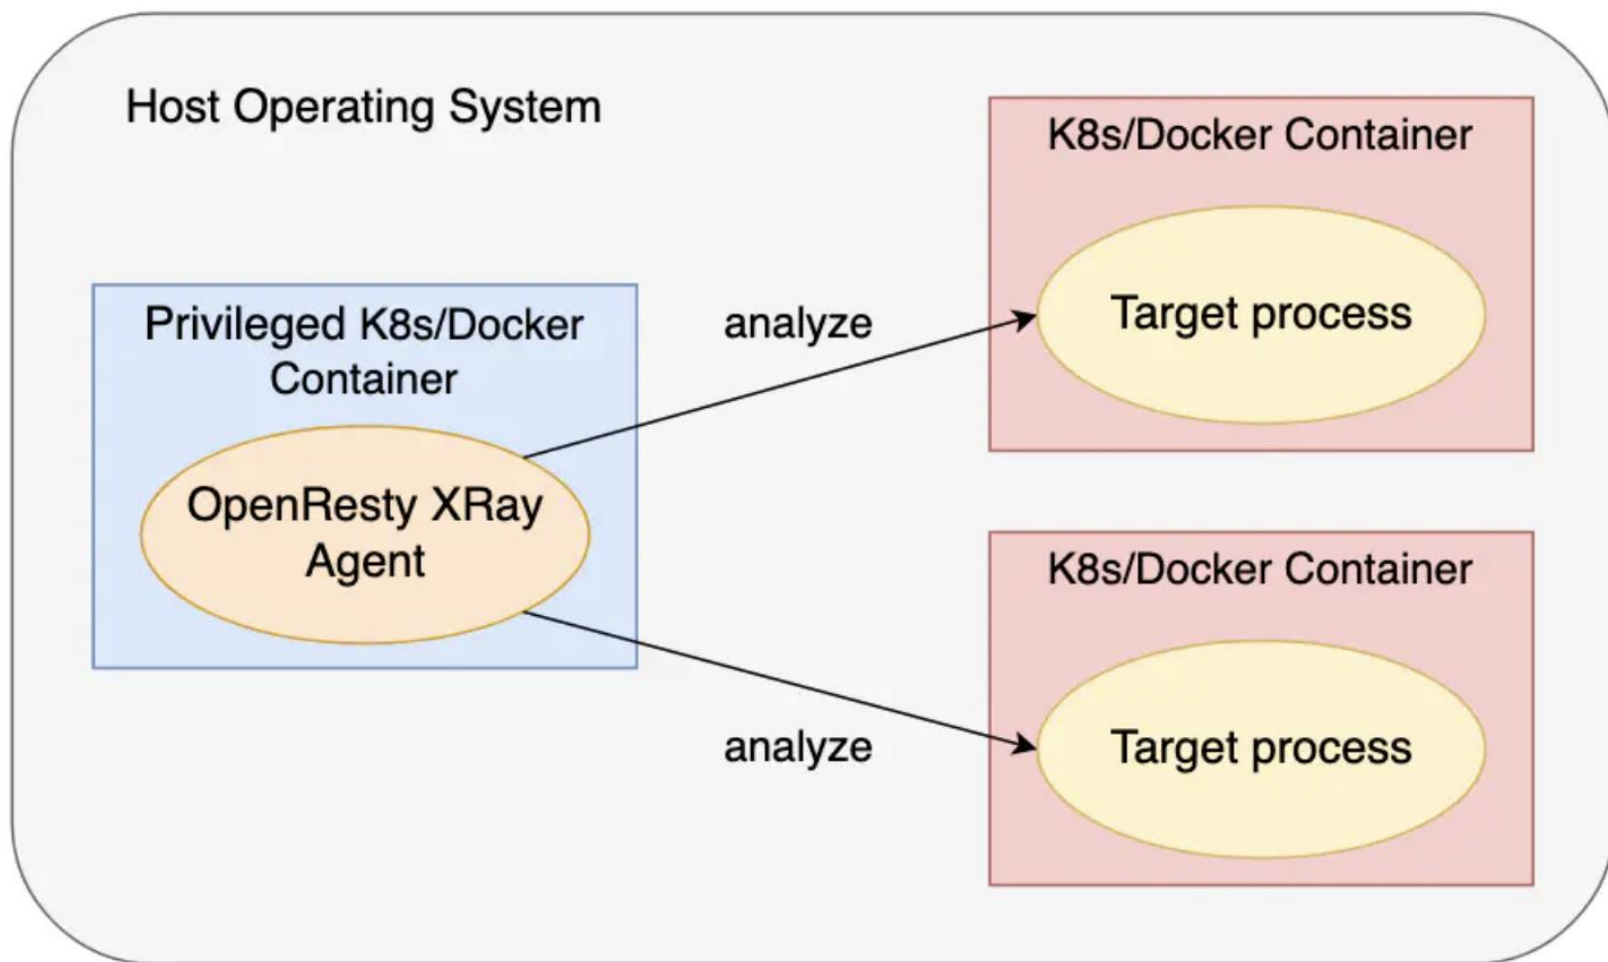


- OpenResty XRay 是一款动态追踪产品
  - 可以实时分析各种云和服务器应用程序
  - 将运行中的进程和容器视为只读数据库，并提取必要的信息来解决性能问题、异常、错误和安全漏洞
  - 拥有知识库、推理引擎和数百个高级分析器
  - 可以在不改变或影响目标应用程序的情况下诊断和缩小深层问题的根源
-

# OpenResty XRay 直接分析非容器应用进程




# OpenResty XRay 穿透容器分析应用



# 100% 非侵入式

---


- 无需修改您的应用
- 无需加装新的插件、模块或补丁到您的应用
- 无需注入任何代码到您的应用
- 无需重启您的应用进程
- 无需在您的应用中使用特殊的启动或编译选项
- 无需重建您现有的应用容器或应用软件包





# OpenResty XRay 全自动采样 无人值守的使用模式

- 定时采样
- 事件驱动采样（CPU 变化，内存变化，IO 变化，异常错误）
- 链式推理驱动



# 极低的性能损耗

---

- 未采样时的性能开销严格为 0
- 采样时的性能开销通常不易觉察



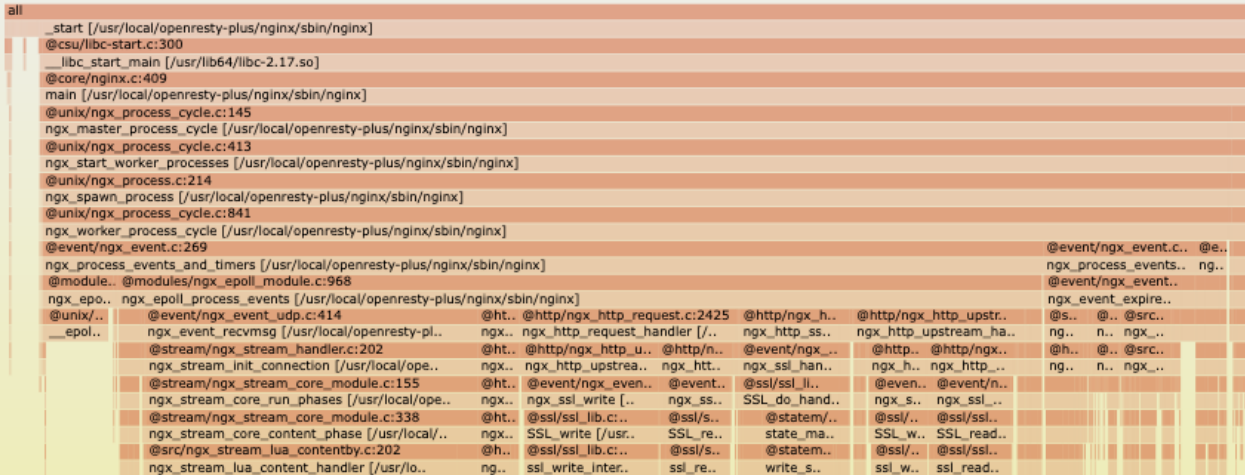
# OpenResty XRay CPU 性能分析

- 高 CPU 使用率会降低系统稳定性和服务质量，甚至导致服务不可用
- CPU 时间在不同场景下，是如何分布在不同代码路径上的（火焰图，火焰图自动解释器）
- 覆盖不同软件层面的代码路径：业务编程语言层面（Lua/Python/PHP/Perl/Go/等等），系统编程语言层面（C/C++/Rust），操作系统内核层面（网络协议栈/进程调度器/内存管理/系统调用）
- 常见的 CPU 瓶颈举例：重复计算（缺少缓存）、SSL 握手相关、垃圾回收（GC）开销，动态内存分配开销、序列化与反序列化、意外的密集系统调用、死循环、病态正则表达式匹配、实现低效的（第三方）软件库、自旋锁竞争

### C-Land CPU Flame Graph for LuaJIT

Search


whole application



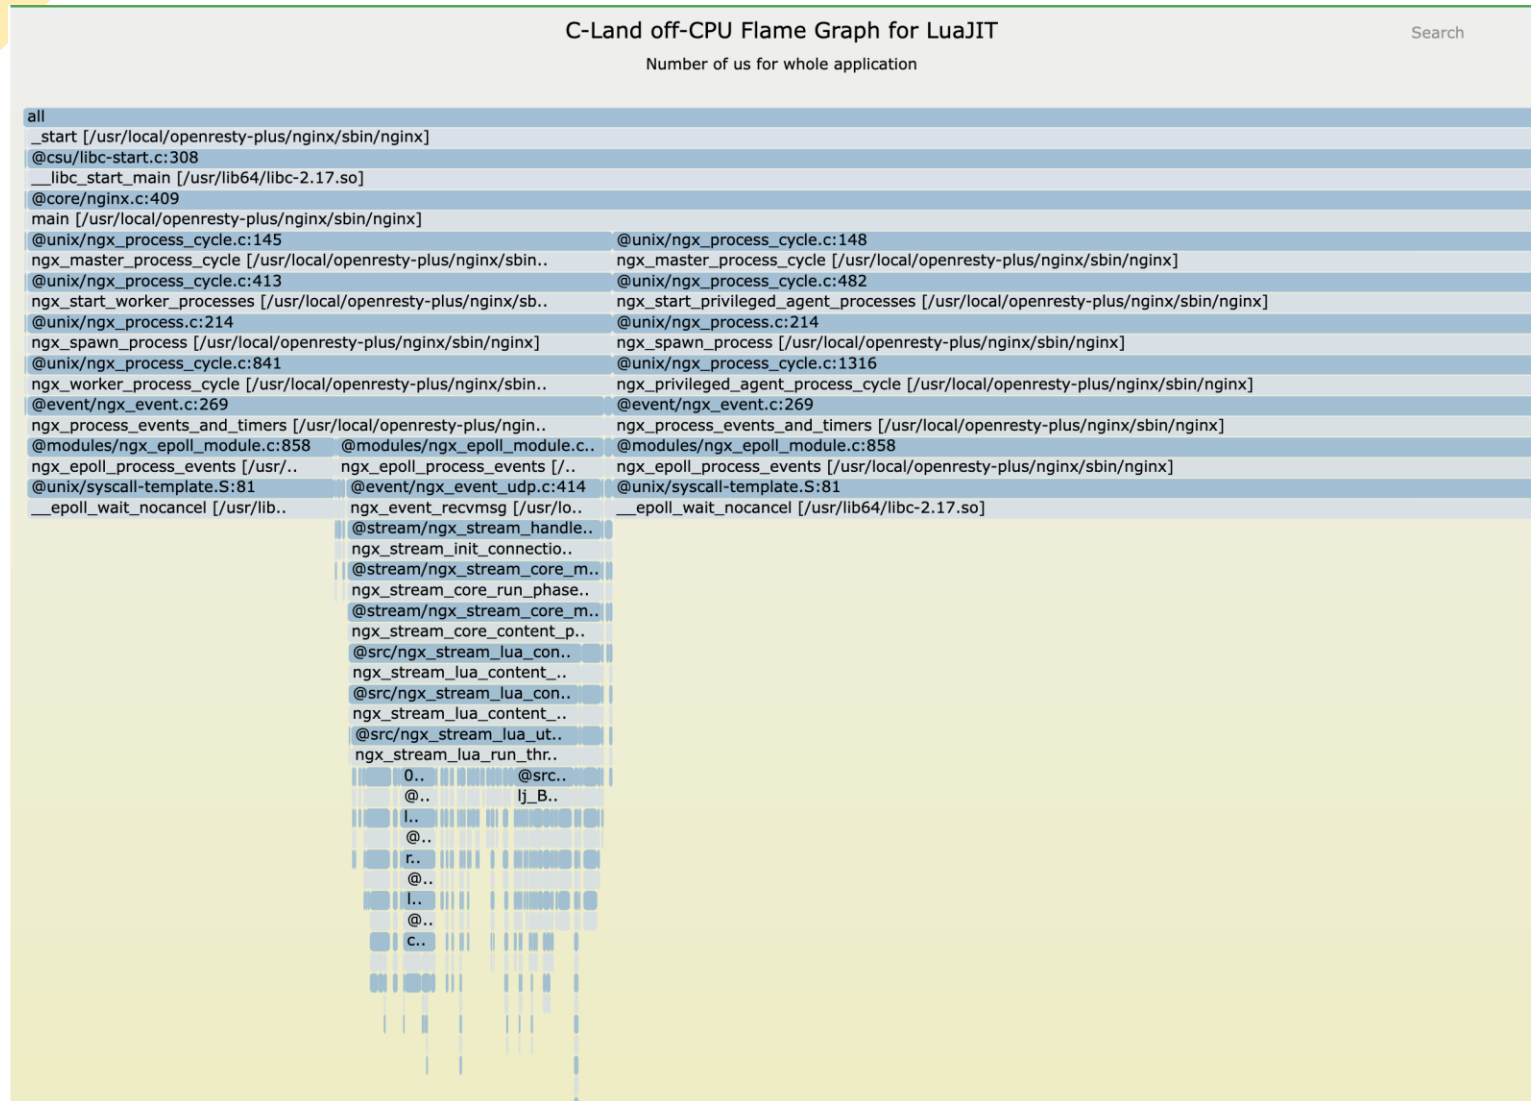
### Lua-Land CPU Flame Graph

Search

whole application



# OpenResty XRay CPU 阻塞 (off-CPU) 分析




# OpenResty XRay 请求延时分析

- 分解延时到不同应用操作和处理阶段
- 精准抓包，只抓实际有问题的连接上的网络包（问题包括延时较高、超时、连接错误、上层应用报错等等）
- 异步非阻塞 IO 的延时统计（如 Lua 协程 yield 的时间在 Lua 代码路径上的分布）

## Nginx Request Latency

In 0 seconds, total 10 requests, matched 10 requests




# OpenResty XRay 内存使用分析

- Glibc/Jemalloc 等 C 内存分配器的内存使用（含 Glibc 内存碎片）
- 内存如何定量分布在所有的 GC 对象上（比如 Lua 对象、Python 对象、PHP 对象等等），按 GC 对象之间的引用关系。
- 内存泄漏，内存碎片，还是延迟释放？


## Application-Level Memory Usage Breakdown

02/04/2023 20:45, Total: 82.78MB



# GC 对象引用关系火焰图


定量显示内存是如何在所有的对象引用路径上分布的






# Nginx memory pool allocations via the brk syscall

When the old and new nginx cycles coexist




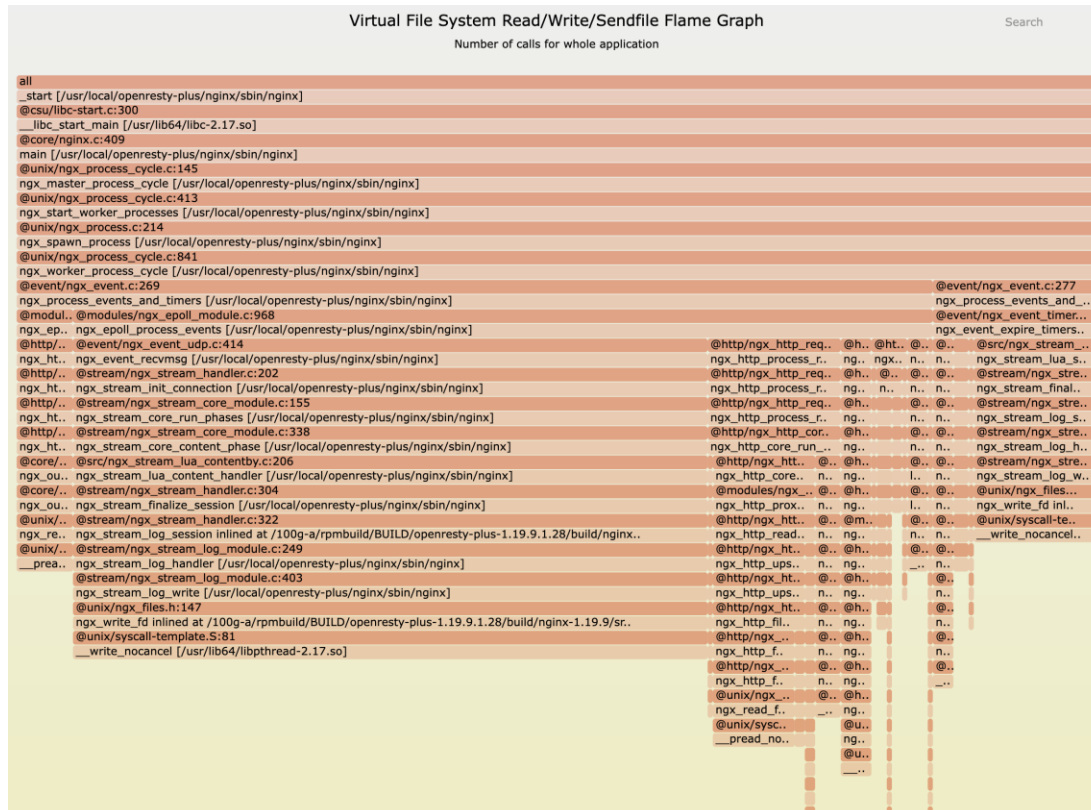
# Nginx memory pool allocatons via the mmap syscall

For total 35.00 MB in 35 memory mappings with 49,464 chunks




free                      other libc alloc  
old cycle pool            new cycle pool

# OpenResty XRay 文件 IO 性能分析



# 在线智能网络抓包

只抓真正有异常的网络连接上的包



# OpenResty XRay 自动分析诊断报告

阅读自动分析和诊断报告的博客文章

Daily report Weekly report 2023-01-12

← Previous day

## OpenResty XRay Analysis Report for Demo



See the full report →

Agent(s): 1

Online time: 23 hours

Watch time: 1 day

2023-01-04 16:00:00 ~ 2023-01-05 15:40:59

OpenResty ⓘ

CPU 25

off-CPU 1

Errors & Exceptions 4

Memory 16

CPU Usage: min: 0%, avg: 0.78%, max: 102%

Command line: nginx: master process /usr/local/openresty/nginx/sbin/nginx

Exe path: /usr/local/openresty/nginx/sbin/nginx

### ⓘ CPU 25

▶ Lua code execution takes up to **100.00%** of the CPU time of the target processes. [More ↓](#)

**Suggestions:** To see how the CPU time is distributed across all the hot Lua code paths, we can analyze the *Lua-land CPU ...* [More ↓](#)

▶ **NEW** [10.1%] #1 of the hottest Lua code paths for CPU time: `lj_err_argt` ← `lj_lib_checkstr[10]` ← `lj_fff_fallback` ← `[builtin#strin` [More ↓](#)

▶ **NEW** [11.9%] #2 of the hottest Lua code paths for CPU time: `gc_onestep` ← `lj_gc_step` ← `lj_gc_step_jit` [More ↓](#)

▶ **NEW** [45.5%] #3 of the hottest Lua code paths for CPU time: `lj_BC_FUNCC` ← `print` ← `C:ngx_http_lua_ngx_print` [More ↓](#)

▶ The `string.gmatch()` Lua function calls takes up to **89.80%** of the CPU time of the target processes. [More ↓](#)

**Suggestions:** Try optimizing the CPU time usage on the code paths highlighted.

▶ **NEW** [29.2%] #1 of the hottest C-land code paths for CPU time: `__GI__writev` ← `ngx_writev` ← `ngx_linux_sendfile_chain` ← `ng` [More ↓](#)

# 自动内存问题诊断报告

## i Memory 15

- NEW Glibc allocator takes up to **72.68 MB** in a target process. [More ↓](#)
- NEW Glibc arena takes up to **72.68 MB** in a target process. [More ↓](#)
- NEW In-use total memory in glibc arena takes up to **67.06 MB** in a target process. [More ↓](#)
- NEW Reserved free memory in glibc arena takes up to **5.67 MB** in a target process. [More ↓](#)
- NEW Free memory reserved by the LuaJIT allocator takes up to **12.19 MB** in a target process. [More ↓](#)
- NEW In-use total memory by the LuaJIT allocator takes up to **4.26 MB** in a target process. [More ↓](#)
- NEW Lua GC size of all types takes up to **1.45 MB** in a target process. [More ↓](#)
- NEW [**10.7%**] #1 of the hottest reference paths for LuaJIT GC object: `table 0x7f950c0a0828 (584B 524.40KB)` ← `[light userdata]` [More ↓](#)
- NEW [**12.4%**] #2 of the hottest reference paths for LuaJIT GC object: `trace 0x7f950a93a8e8 (2.36KB 161.13KB)` ← `next side` ← [More ↓](#)

# 自动延时分析与诊断

## Latency 5

- ▼ ↑ 22.2% [100%] #1 of the hottest Lua code paths for Newly Created CoSocket: C:ngx\_http\_lua\_socket\_tcp\_connect ← connect ← C:ngx\_http\_lua\_socket\_tcp\_connect ← check\_peer ← spawn\_checker ← check\_peers ← pcall ← [builtin#pcall]  
See [Job 4418885005](#) for more details.  
[Collapse ↑](#)
- ▶ [100%] #2 of the hottest Lua code paths for Newly Created CoSocket: C:ngx\_http\_lua\_socket\_tcp\_connect ← connect ← C:ngx\_http\_lua\_socket\_tcp\_connect ← \_request ← request\_admin ← pcall ← [builtin#pcall] [More ↓](#)
- ▶ ↑ 105.67 ms [106.97 ms] #1 of the hottest Lua code paths for Request Yield Latency: lua\_yield ← lj\_BC\_FUNCC ← ngx\_sleep ← C:ngx\_http\_lua\_ngx\_sleep ← limit\_req\_rate ← helper\_1 ← xpcall ← [builtin#xpcall] ← run\_rewrite\_phase ← access ← access\_by\_lua(nginx.conf:586) [More ↓](#)
- ▶ ↑ 4.00 ms [16.00 ms] #2 of the hottest Lua code paths for Request Yield Latency: lua\_yield ← ngx\_stream\_lua\_socket\_tcp\_receive ← lj\_BC\_FUNCC ← receive ← C:ngx\_stream\_lua\_socket\_tcp\_receive ← go ← content\_by\_lua(nginx.conf:140) [More ↓](#)
- ▶ [1.58 ms] #2 of the hottest Lua code paths for Request Yield Latency: lua\_yield ← lj\_BC\_FUNCC ← ngx\_sleep ← C:ngx\_stream\_lua\_ngx\_sleep ← limit\_req ← process\_req ← go ← content\_by\_lua(nginx.conf:132) [More ↓](#)

# off-CPU 自动诊断报告

## i off-CPU 8

- ▶ NEW Lua code execution takes up to **99.99%** of the off-CPU time of the target processes. [More ↓](#)

**Suggestions:** Try optimizing the off-CPU time usage on the code paths highlighted.

- ▶ NEW [21.5%] #1 of the hottest Lua code paths for off-CPU time: `__read_nocancel` ← `_IO_file_underflow@@GLIBC_2.2.5[2]` ← `_IO_default_xsgetn[2]` ← `lj_BC_FUNCC` ← `read` ← `[builtin#` [More ↓](#)
- ▶ NEW [31.7%] #2 of the hottest Lua code paths for off-CPU time: `__read_nocancel` ← `fread[2]` ← `lj_BC_FUNCC` ← `read` ← `[builtin#io.method.read]` ← `_getAccessLog` ← `_getIfNumber` [More ↓](#)
- ▶ NEW [32.6%] #3 of the hottest Lua code paths for off-CPU time: `__read_nocancel` ← `lj_BC_FUNCC` ← `read` ← `[builtin#io.method.read]` ← `_getProxyIgnoreHeaderInfo` ← `_getIfNumber` [More ↓](#)



# 异常错误自动诊断报告

## i Errors & Exceptions 2

- ▶ NEW [100%] #1 of the hottest Lua code paths throwing out Lua exceptions: 71) ← no field package.preload['test'] ← no file '/' [More ↓](#)
- ▶ NEW [100%] #2 of the hottest Lua code paths throwing out Lua exceptions: 166) ← no field package.preload['resty.http'] ← r [More ↓](#)

# 安全问题自动分析

自动检查和报告未开启 TLS 加密的连接

动态扫描未开启证书来源校验的 TLS 连接

检查使用了非安全的 SSL 协议版本

扫描远程 shell 命令执行的事件与代码上下文



# Core Dump 进程遗骸分析 (进程崩溃)

All

Core Dump Analysis

Application Type \*

OpenResty

Analyzer \*

openresty-core-dump-analysis (OpenResty Core Dump Analysis)

Core File \*

/tmp/core.3859164

Executable File Path \*

/usr/local/openresty/nginx/sbin/nginx

► Advanced Settings

## Core File Meta Data

**File Name:** /tmp/core.3859164

**Executable File Path:** /usr/local/openresty/nginx/sbin/nginx

**Size:** 14.72MB

Analyze

# 从 Core Dump 文件提取深层信息

Analysis 7173425 [🔗](#)

## (gdb) lbt

```
C:ngx_md5_body
trace#1:access.lua:4
check_token
@/usr/local/openresty/lualib/access.lua:3
auth
@/usr/local/openresty/lualib/access.lua:21
@access_by_lua(nginx.conf:51):2
```

## (gdb) full\_lbt

```
C:ngx_md5_body
trace#1:access.lua:4
check_token
@/usr/local/openresty/lualib/access.lua:3
auth
@/usr/local/openresty/lualib/access.lua:21
headers = "access"
token = "hello"
@access_by_lua(nginx.conf:51):2
```

Compiler Output

Analyzer Output

Graphs

## (gdb) ngx\_process\_info

```
parent: 3859163
process: worker 0
```

## (gdb) cur\_http\_req

```
current phase: access
schema: http, req_size: 52, resp_size: 0GET / HTTP/1.1
Host: localhost:80
TOKEN: hello
```

## (gdb) ubt

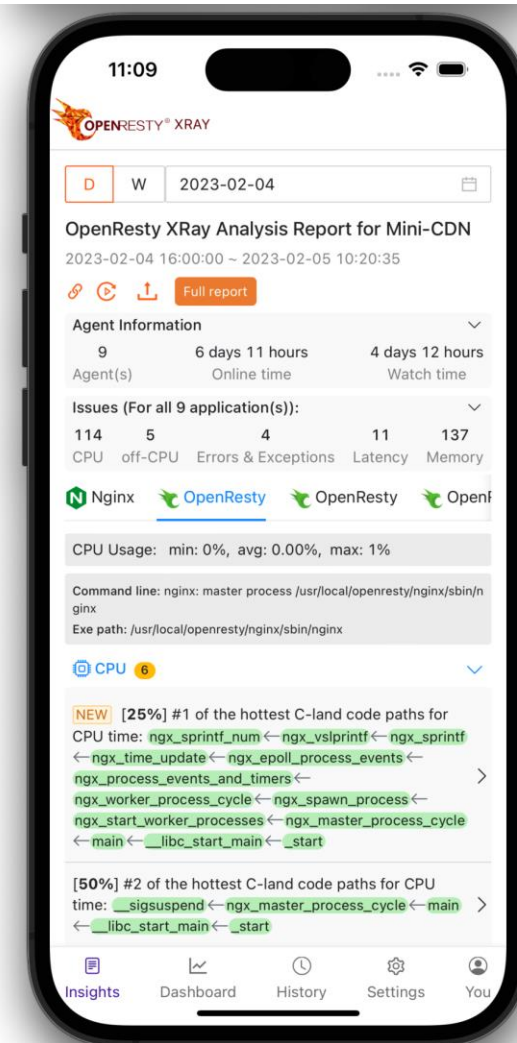
```
0x42e958 ngx_md5_body [/usr/src/debug/openresty-1.21.4.1/build/nginx-1.21.4/src/core/nginx_md5.c:199]
0x42f1be ngx_md5_final [/usr/src/debug/openresty-1.21.4.1/build/nginx-1.21.4/src/core/nginx_md5.c:91]
0x4ea997 ngx_http_lua_ffi_md5 [/usr/src/debug/openresty-1.21.4.1/build/nginx-1.21.4/./ngx_lua-0.10.21/s
7f763447ffd3: []
0x4f86a1 ngx_http_lua_run_thread [/usr/src/debug/openresty-1.21.4.1/build/nginx-1.21.4/./ngx_lua-0.10.2
```

# OpenResty XRay

## 移动端 App

任何时候从任何地方察看您的在线应用

- Android (Google Play)
- iOS (苹果商店)




## OpenResty XRay 并非只针对 OpenResty 应用

- Nginx, LuaJIT, OpenResty
- 初步支持:  
Python, PHP, Perl, Redis, PostgreSQL
- 即将发布:  
Go, Ruby, NodeJS, Java



# 支持绝大多数主流 Linux 发行版和容器部署方式



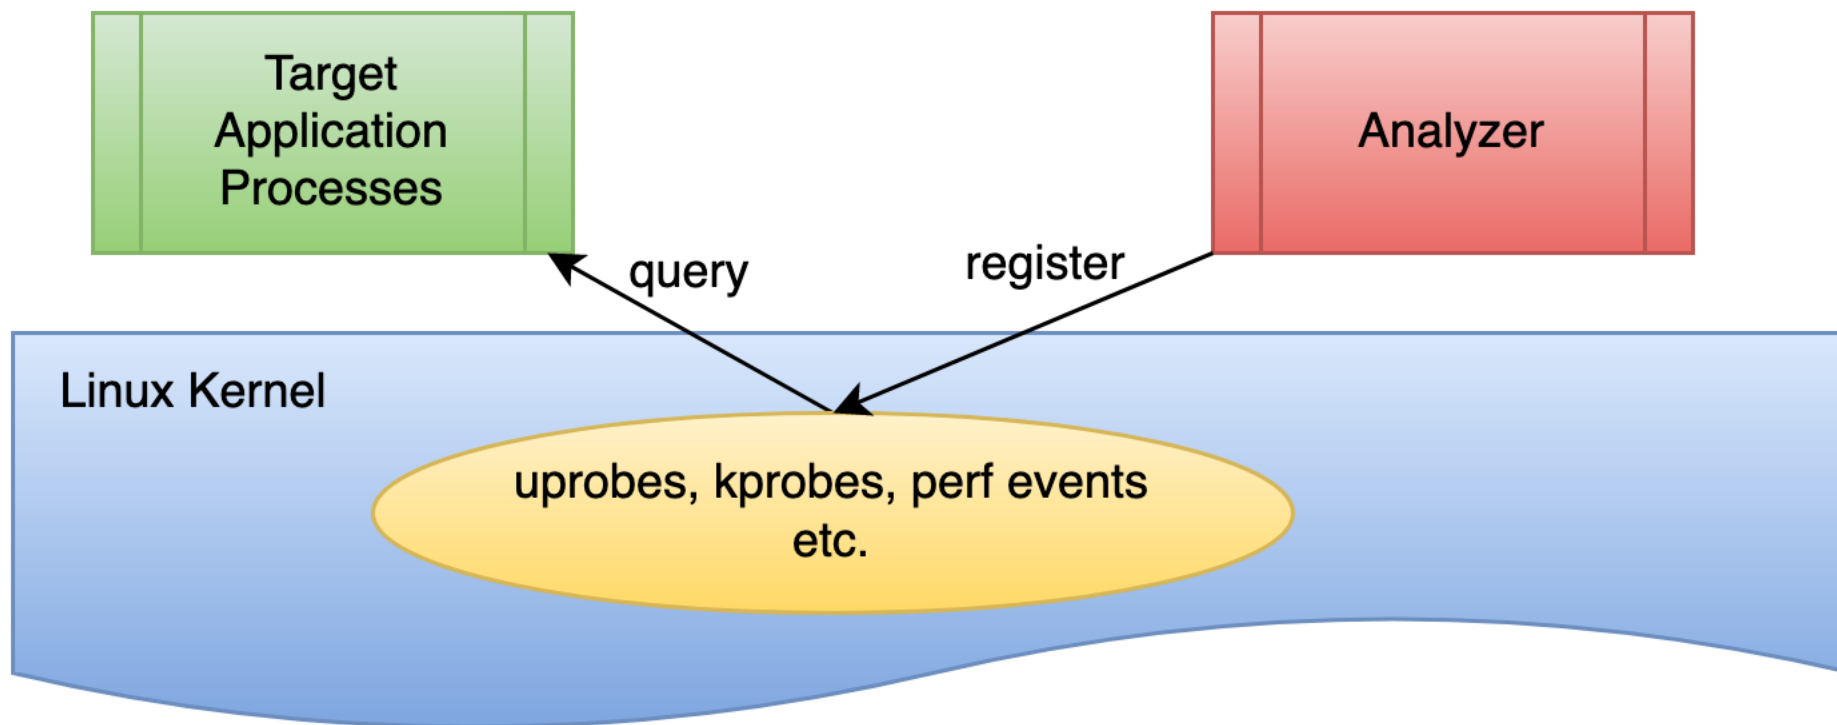
Ubuntu



Debian



# OpenResty XRay 基于先进的动态追踪技术





# 动态追踪的优点

- 非侵入式，无需修改应用
- 热插拔，通常无需应用配合（很多开源动态追踪工具有时还是要求应用配合）
- 开销通常较低，可以在数据源进行聚合汇总
- 事发后的在线实时调试能力
- 全技术栈分析无死角
- 按需采样
- 不采样时严格 0 损耗

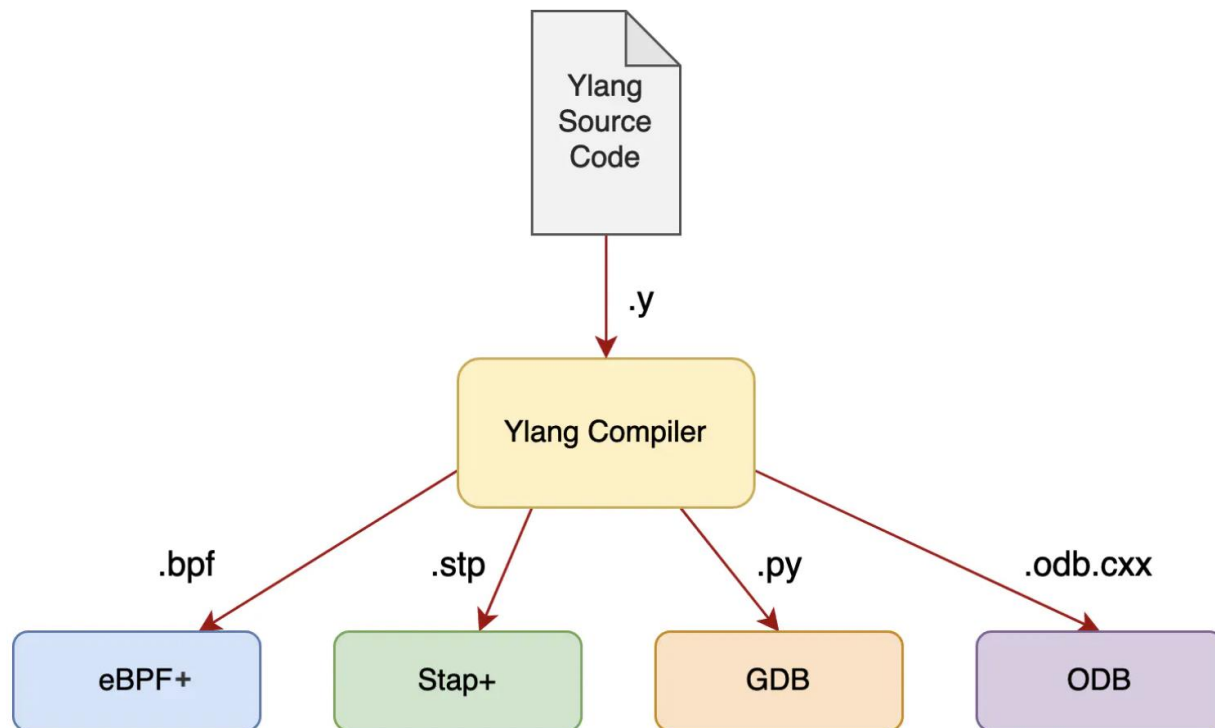
# OpenResty XRay 的全新一代动态追踪技术

---


- Y 语言 (Ylang) 编译器 (支持 GNU C 和标准 C 语言的绝大部分语法)
- Ylua 语言
- YSQL 语言
- Stap+ 显著改进了 SystemTap
- eBPF+ 显著改进了 eBPF (同时 LLVM+ 也大大改进了开源 LLVM)
- ODB 是超轻量版的 GDB
- Ylang 编译器也能生成高度优化的 GDB 的 Python 扩展代码
- 针对线上生产环境的严格性能损耗控制

# “一次编写，到处运行”

从 Y 语言代码到各种不同运行时的代码



# Y 语言之上更抽象的语言



# 在 OpenResty XRay 界面上编写和调试 Ylang/YLua/YSQL 等语言的分析工具

The screenshot displays the OpenResty XRay web interface. The top navigation bar includes the OpenResty XRay logo, a version indicator (844), a dropdown menu for the current instance (prd.openresty.com), and user information (yichun@openresty.com). The left sidebar contains a navigation menu with options like Overview, CPU, Memory, Disk, Network, Latency, Analyzers (highlighted), Load Gen, History, Events, Advanced, Custom Apps, Settings, and Upload.

The main content area is titled "All analyzers / Edit analyzer". It features a form for editing an analyzer. The "Analyzer name" is "process-exit" and the "Analyzer description" is "C-land CPU Flame Graph". Below this, there are tabs for "YSQL", "YLua", and "YLang", with a "Learn YLang" link and "Run" and "Save" buttons. The code editor shows the following YLang code:

```
1 // c-cpu: C-land CPU flame graph sampling tool.
2 // Copyright (C) OpenResty Inc.
3 // All rights reserved.
4
5 _target sig_atomic_t ngx_quit;
6 _target sig_atomic_t ngx_debug_quit;
7 _target ngx_uint_t ngx_exiting;
8 _target sig_atomic_t ngx_reconfigure;
9 _target sig_atomic_t ngx_reopen;
10
11 _probe _process.begin
12 {
13     printf("ngx_quit %ld\n", ngx_quit);
14     printf("ngx_exiting %ld\n", ngx_exiting);
15     printf("ngx_reconfigure %ld\n", ngx_reconfigure);
16     _exit();
17 }
18
19
```

At the bottom of the editor, there are "Run" and "Clear the editor" buttons. Below the editor, there is a link for "Analysis history for this analyzer".

On the right side, there are several configuration sections:

- Try this analyzer against:** A section with a "Target" dropdown (set to "By Applications") and buttons for "By Processes", "By Executables", and "Whole System".
- Application:** A dropdown menu showing "openresty (PGID 1471, master proc)" and an "Update" button.
- Target Processes:** A dropdown menu showing "PID 2782 (worker process) CPU: 1%" and an "Update" button.
- Advanced Settings:** A section with a "YLang Settings" dropdown.
- Dependent debug data to compile the analyzer:** A section with an information icon and text: "We have a huge package archive database for public software and we will not install any debuginfo packages on the target machine."
- Kernel debug info:** A section with a dropdown menu showing "Not required", "Required", and "Good to have".
- Unwind data:** A section with a dropdown menu showing "NO NEED" and an information icon: "Need unwind data when you need the backtrace."

# OpenResty XRay 数百种标准分析器

## ▼ OpenResty XRay Standard Analyzers

- c-alloc-fgraph
- c-count-alloc-free
- c-memory
- c-memory-leak-fgraph
- c-off-cpu
- c-on-cpu
- collect-luajit-ffnames
- cpu-hogs
- epoll-loop-blocking-distr
- epoll-sched-latency-distr
- epoll-wait-ret-distr
- epoll-wait-timers
- epoll-wait-timers-fgraph
- file-system-fgraph
- func-latency-distr
- glibc-chunks
- jemalloc-bins

- kernel-on-cpu
- lj-add-timer-lua-fgraph
- lj-alloc-stats
- lj-c-memory-leak-fgraph
- lj-c-off-cpu
- lj-c-on-cpu
- lj-config
- lj-dump-loaded-mods
- lj-err-mem
- lj-excep-lua-fgraph
- lj-free-stats
- lj-gc-step-calls
- lj-lua-exception
- lj-gco-ref
- lj-gco-stat
- lj-lua-err-msg
- lj-lua-new-timer-errors
- lj-lua-newcdata
- lj-lua-newfunc
- lj-lua-newgco

- ngx-add-timer-event-fgraph
- lj-trace-stats
- ngx-add-timer-event-timer-distr
- lj-vm-states
- mmap-leaks
- musl-libc-chunks
- ngx-access-log-buffer-size
- ngx-config
- ngx-config-servers
- ngx-cpu-hottest-hosts
- ngx-cpu-hottest-uris
- ngx-downstream-keepalive-stats
- ngx-dump-req
- ngx-dump-timers
- ngx-epoll-wait-timers
- ngx-err-log-lvl-distr

# 调试符号

- OpenResty XRay 有一个中央包数据库，索引了几百 TB 的公开包的调试符号，仍在快速增长
- 目标机器无需安装或保存调试符号，只要调试符号曾被 OpenResty XRay 中央包数据库索引
- 无法找到调试符号或编译时未生成调试符号的程序，OpenResty XRay 将有能力自动重建调试符号（已有可工作的机器学习算法的原型实现）



为众多企业客户所信任



**Qunar.Com**

zoom



# 了解更多

OpenResty XRay 常见问题

访问 OpenResty XRay 官方博客

访问 OpenResty XRay 官方主页

试用 OpenResty XRay 产品