


# Perfetto



Platform-wide performance instrumentation and tracing  
for Android and Chrome

Tracing Summit 2018 - Edinburgh

[primiano@google.com](mailto:primiano@google.com)

# What is Perfetto about?

1. Record traces

Tracing library & daemons

*on-device*

2. Analyze traces

Perfetto trace processor

*offline*

3. Visualize traces

Perfetto UI

# What is Perfetto about?

1. An open source (AOSP / Apache2 license) project for recording, processing and visualizing traces.
2. A production C++11 codebase for secure and efficient (zero-copy\*, zero-malloc\*) userspace-to-userspace tracing.
3. Integration with ftrace, /proc/{stat,vmstat,pid/\*} and soon perf\_event\_open.
4. A SQLite-based codebase for analyzing and processing traces.
5. A UI frontend.

\* Some copies / allocations are involved, once every ~4KB.



Linux ftrace

/proc/...  
interfaces

Heap profiler  
Coming soon!

Any app / process

Platform probes services

Shared memory

Shared memory

Producer endpoint  
Perfetto tracing service  
Consumer endpoint

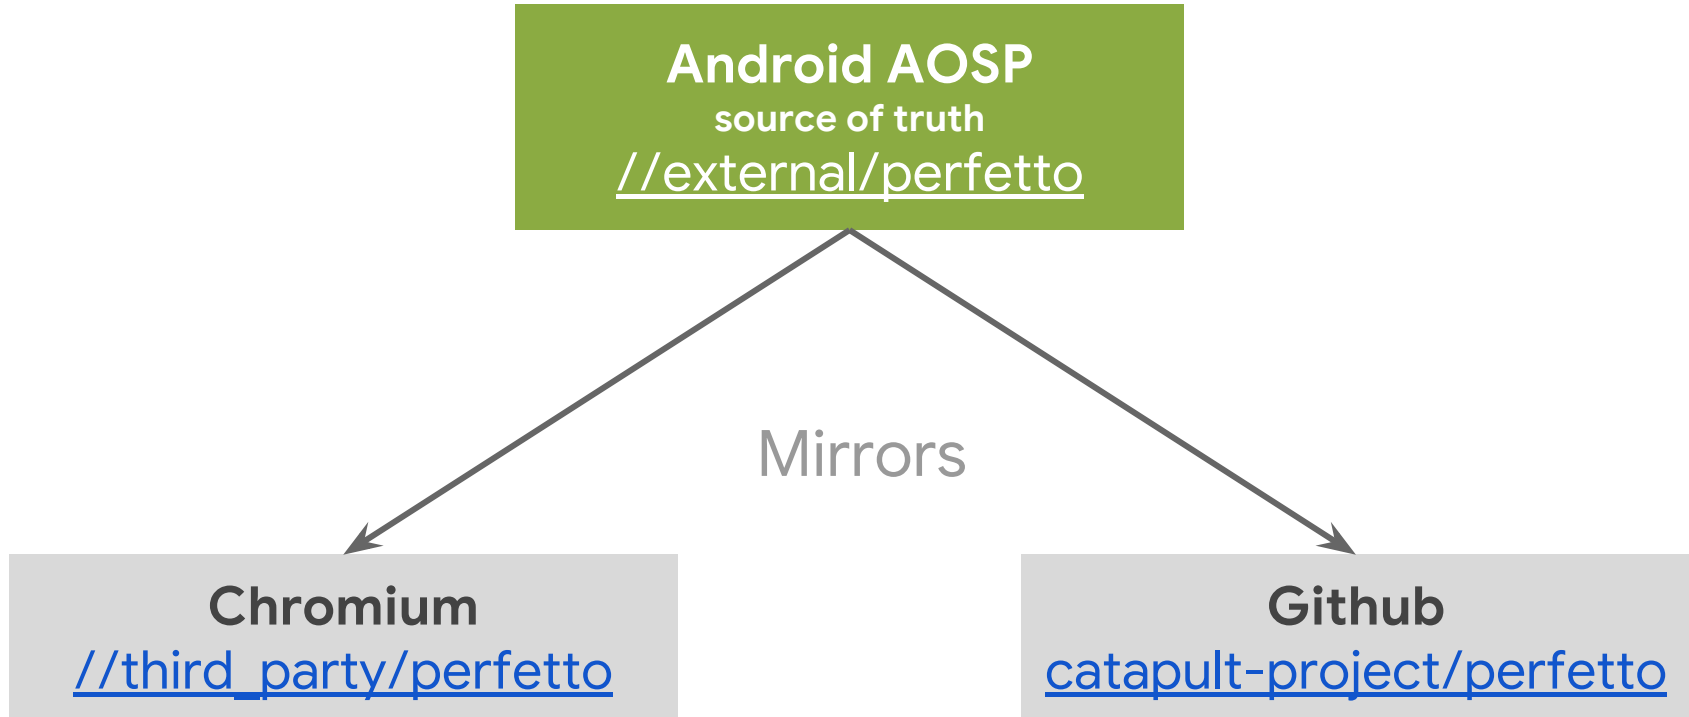
Protobuf-based  
config

Protobuf-based  
trace

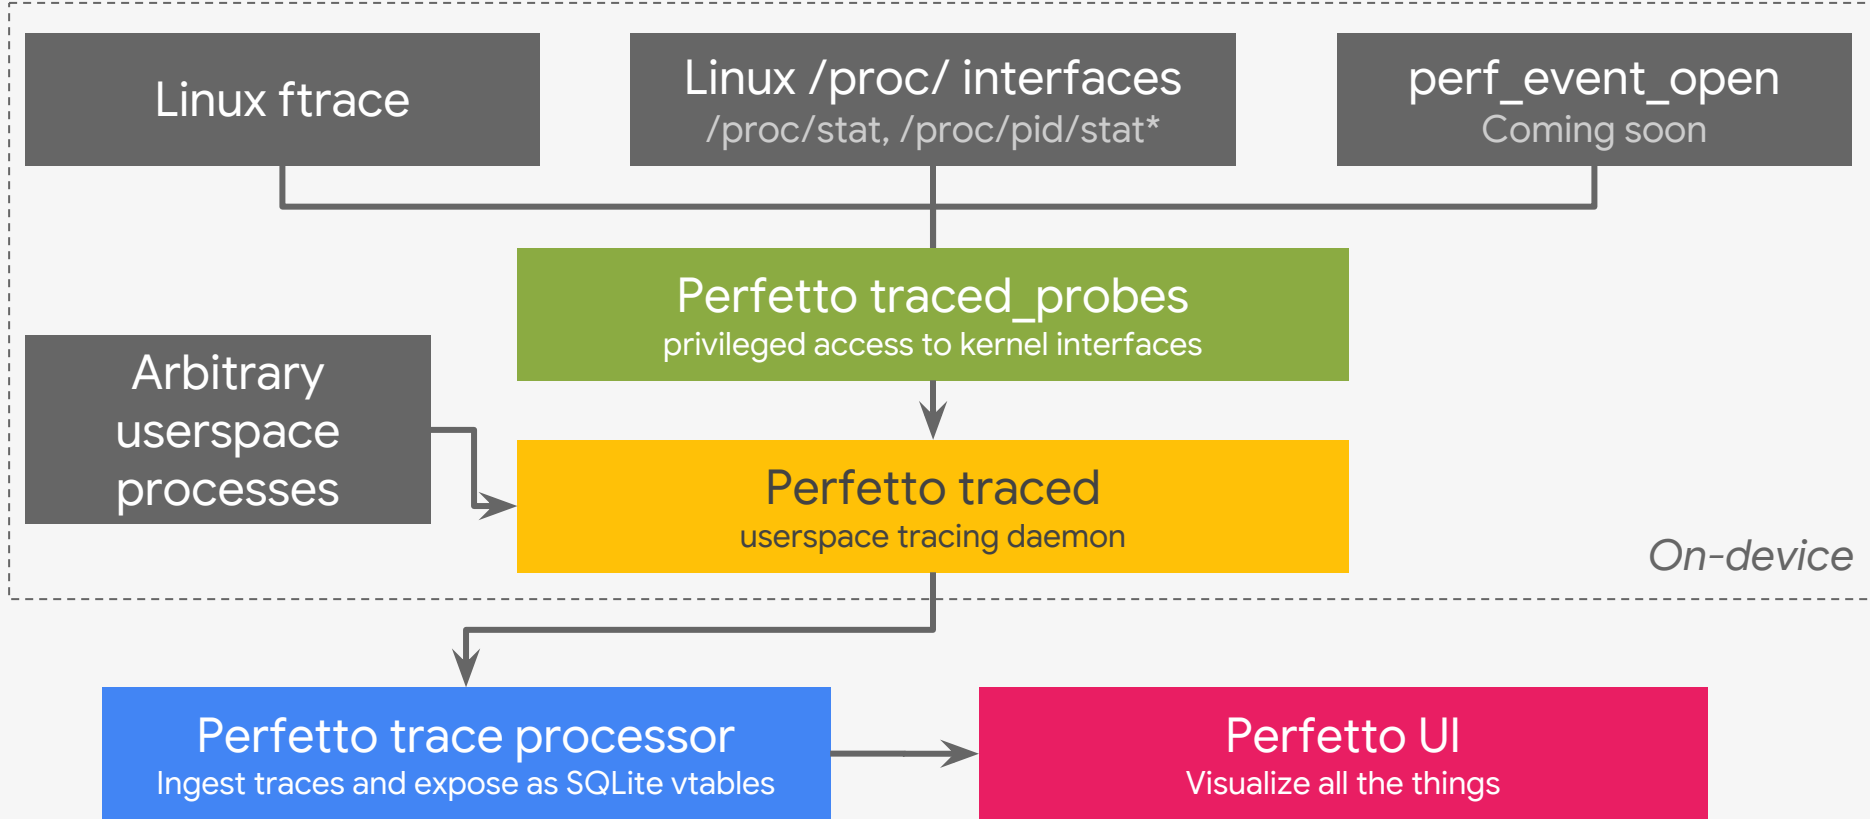
Perfetto UI  
Visualize all the things

Perfetto trace processor  
Ingest traces and expose as SQLite vtables

# Where to find the code?



# What is Perfetto about?



# Userspace tracing library

# Key concepts

## Producers

- The thing that writes protobufs into the trace buffers
- Untrusted. Potentially malicious. Everything can be a Perfetto producer.
- On startup advertises its capabilities to the tracing service.
- At some point the tracing service asks it to start collecting data

## Tracing service

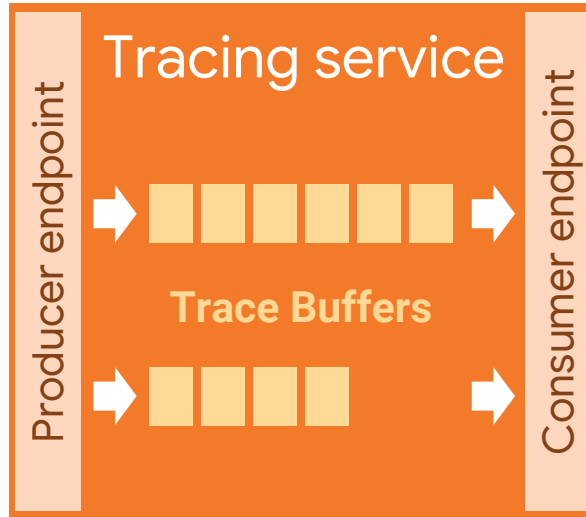
- The thing that owns the log buffers (there is one\* buffer for the all system / browser)
- Acts as registry and handles handshakes between producers and consumer(s)
- In chrome: a /services service
- In android: a system service (traced)

## Consumer(s)

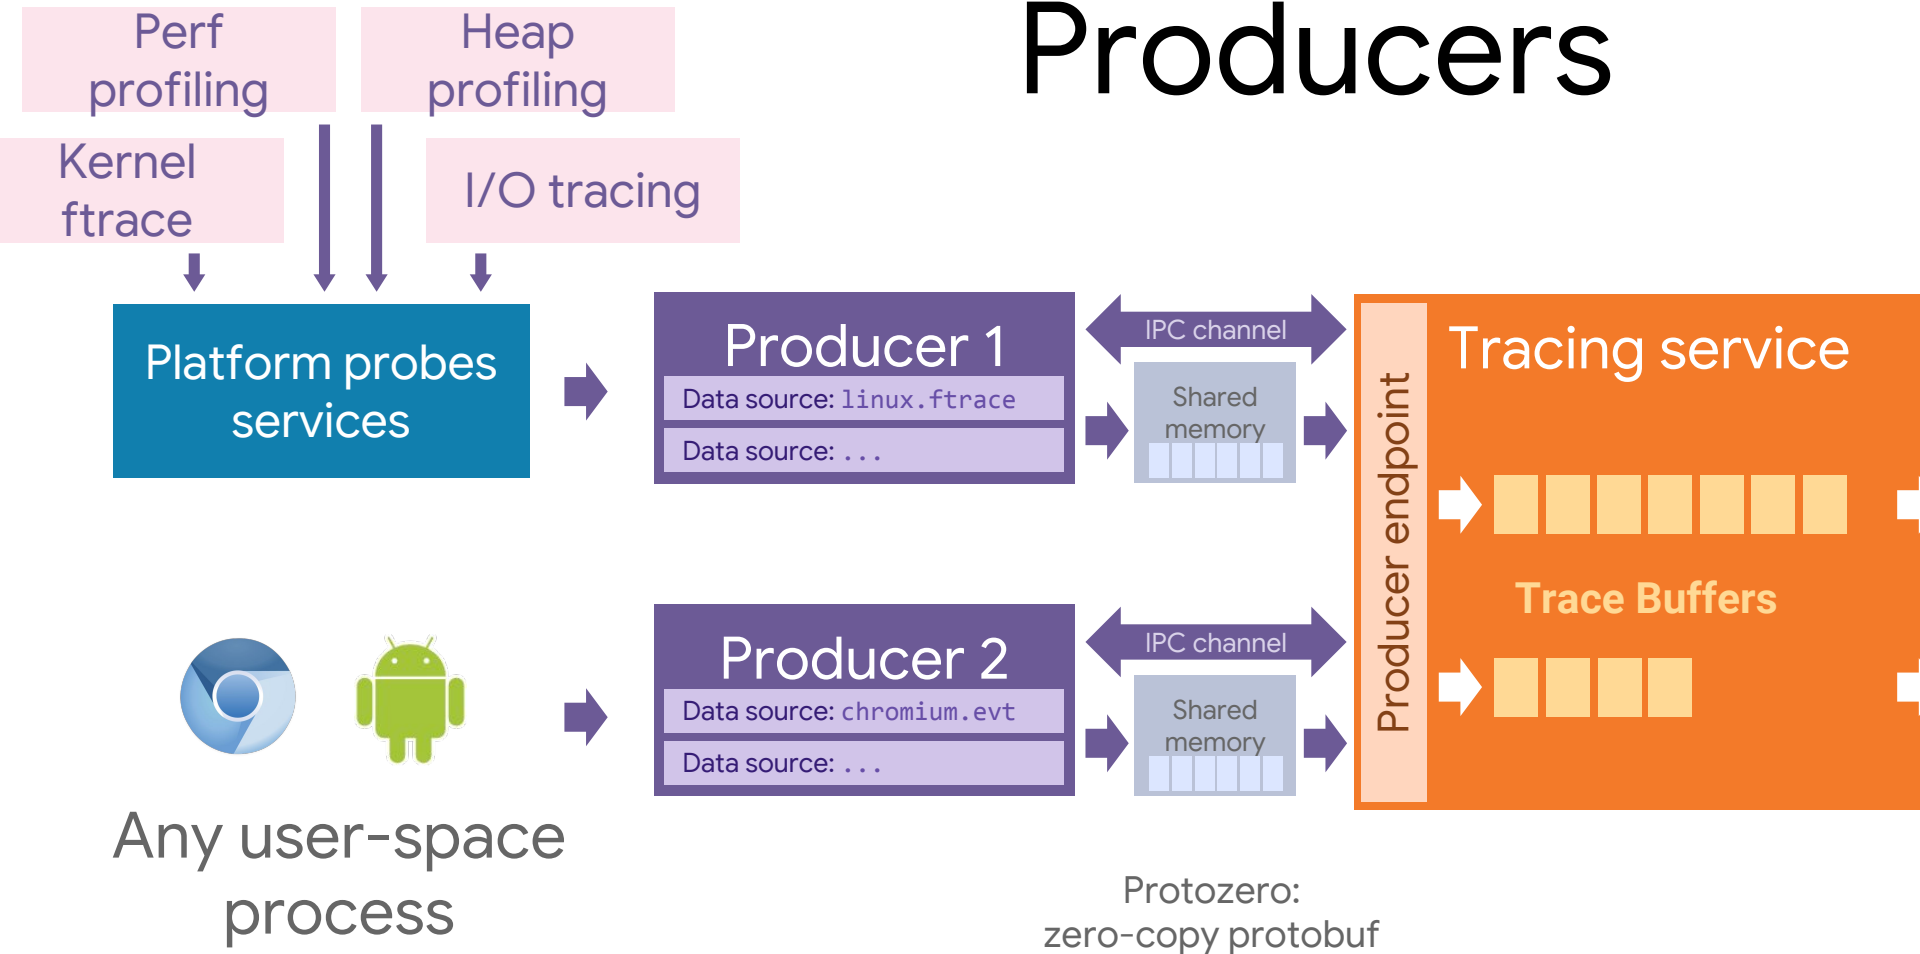
- The thing that configures all the tracing session and decides who should trace and what.
- Is allowed to configure the tracing service and read back the trace data
- Trusted / privileged
- In chrome: the thing that exposes data to the UI
- In android: shell (for the UI) and Android Metrics services



# Tracing Service

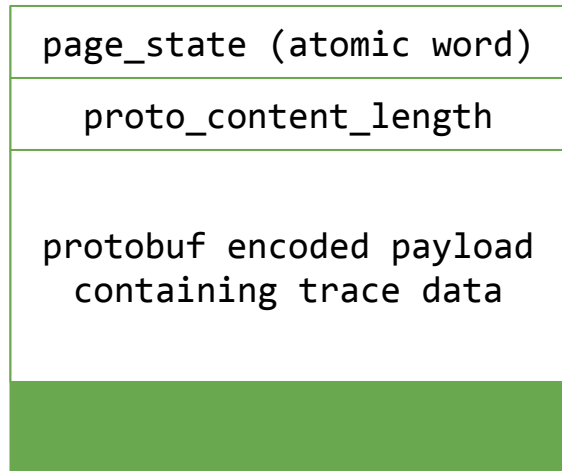


# Producers



# Shmem buffer format

Per-process shared memory buffer

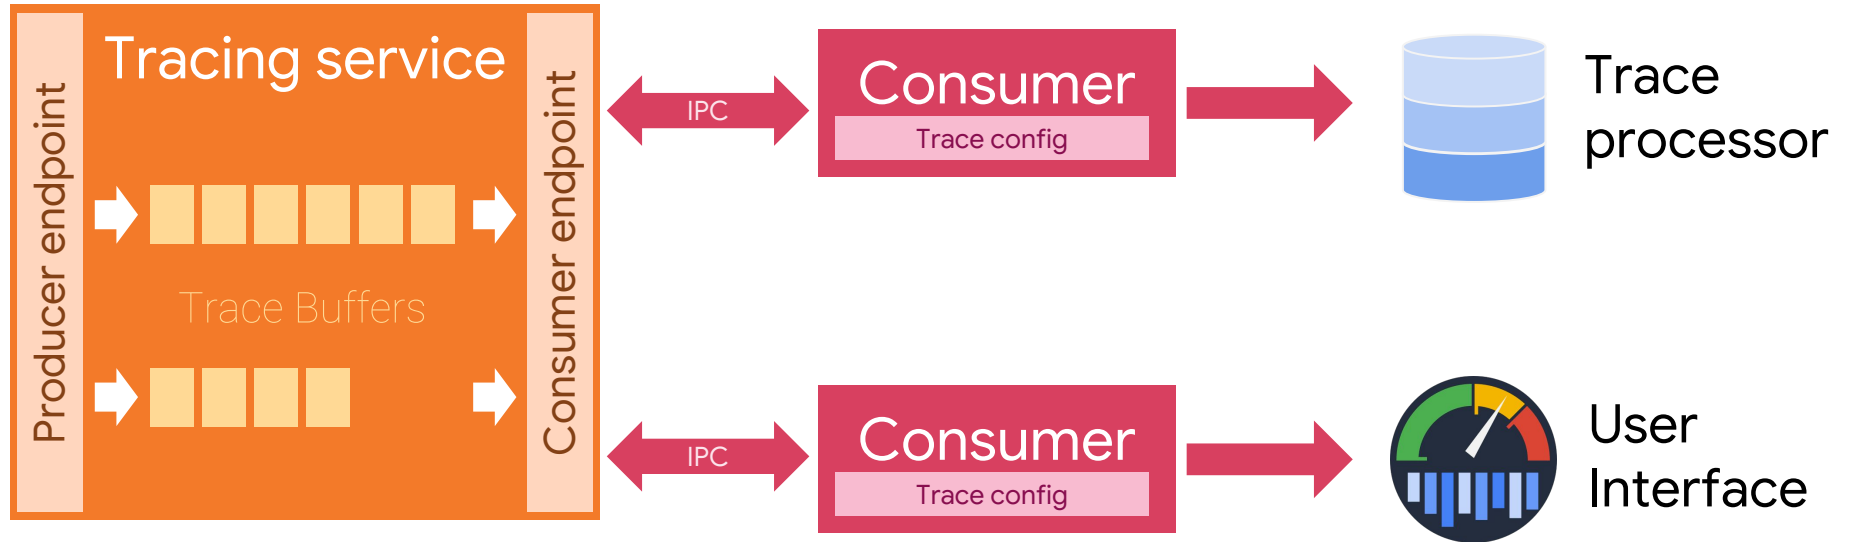


...

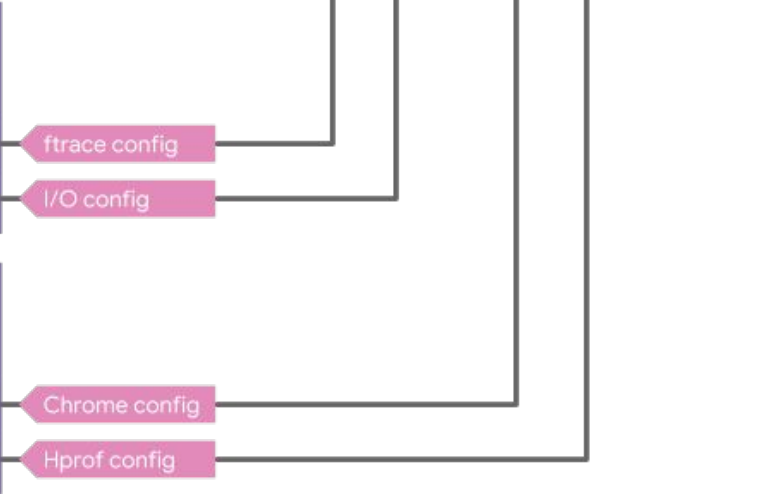

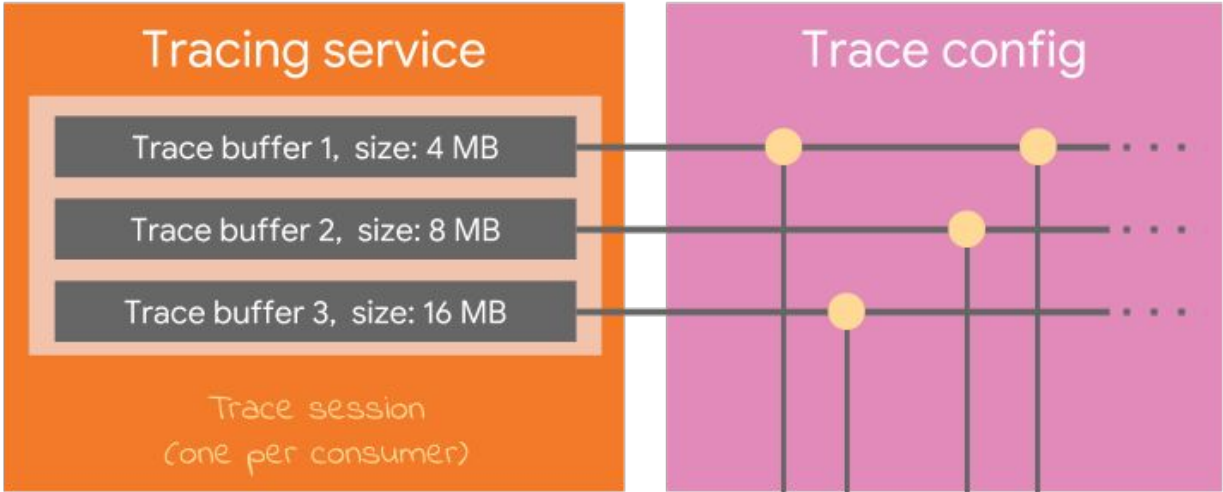


Page

# Consumers

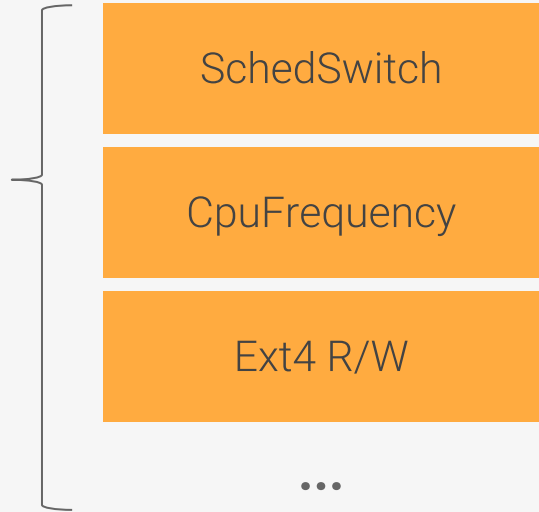
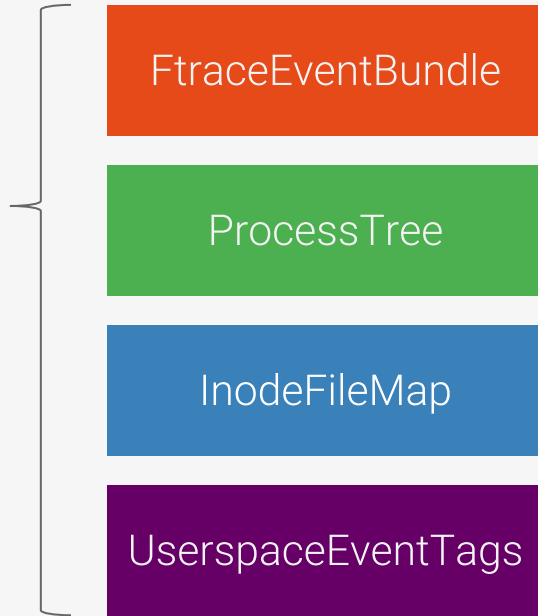
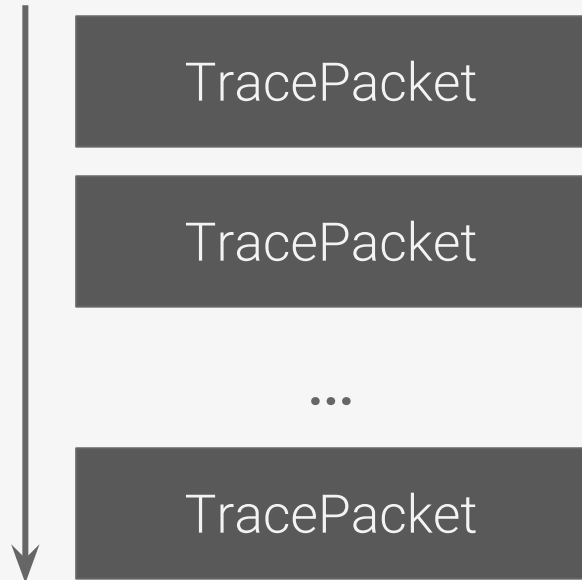


# Trace config



# What is a trace?

A sequence of TracePacket(s)  
protobuf messages



**A trace can be large (10 GB)**

# Trace Processor

# Trace processor

C++11 + SQLite codebase

Ingests traces of various formats (for now our .proto and Chrome's JSON, in future also ftrace text)

Builds an in-memory columnar database from trace contents.

Exposes the storage to SQLite through vtable hooks

Adds some trace-specific constructs on top of conventional SQLite ones.





```
$ out/mac_release/trace_processor_shell ~/Downloads/1gb-trace-truncated.proto
```

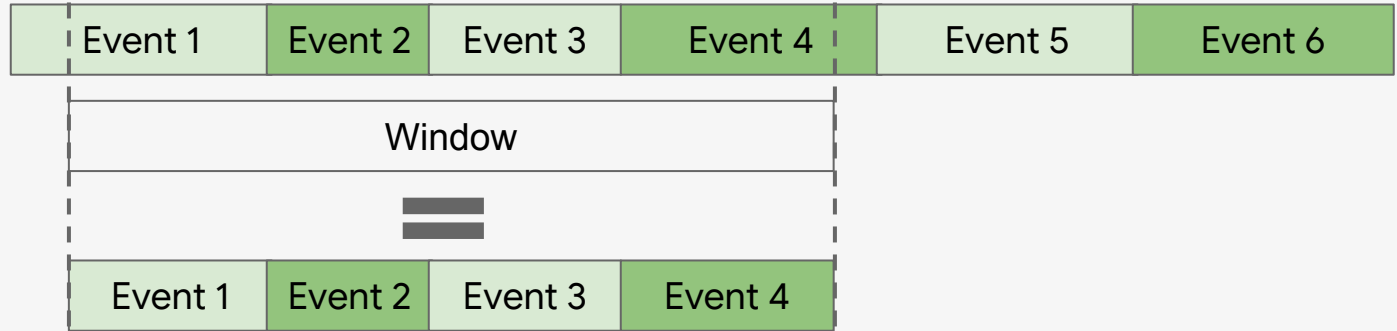
```
trace_processor_shell.cc Trace loaded: 1048.58 MB (184.9 MB/s)
```

```
> select proc_name, cpu, cpu_sec from (select process.name as proc_name, upid, cpu, cpu_sec from (select cpu, utid, sum(dur)/1e9 as cpu_sec from sched group by utid) left join thread using(utid) left join process using(upid)) group by upid, cpu order by cpu_sec desc limit 100
```

proc_name	cpu	cpu_sec
migration/2	2	2532.212882
migration/3	3	2529.064936
migration/1	1	2527.338100
migration/4	4	2526.877703
migration/5	5	2524.508852
migration/6	6	2523.372052
migration/7	7	2522.564051
/system/bin/surfacef	3	22.770180
rcu_preempt	7	16.257903
irq/760-synapti	4	14.566679
smem_native_rpm	7	11.273782
kswapd0	3	10.327598
ksoftirqd/0	0	10.231438
kworker/u16:2	7	9.276288
migration/0	0	8.302623
/vendor/bin/msm_irqb	3	8.256403
kworker/u16:4	7	7.876912
rcuop/0	7	6.730403
rcuos/0	7	6.469543
sugov:0	3	6.113958
/vendor/bin/hw/andro	3	5.919216

# New constructs

## Windowing



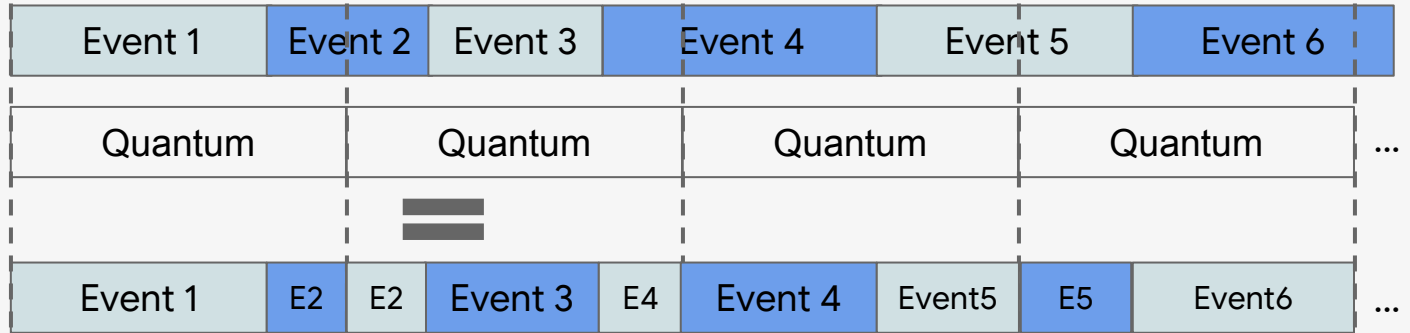
```
CREATE VIRTUAL TABLE bounds USING window;  
UPDATE bounds SET window_start=X, window_dur=Y where 1  
CREATE VIRTUAL TABLE clipped USING span(sched, bounds)
```

↓  
Orig table

↓  
Clip mask

# New constructs

## Quantization



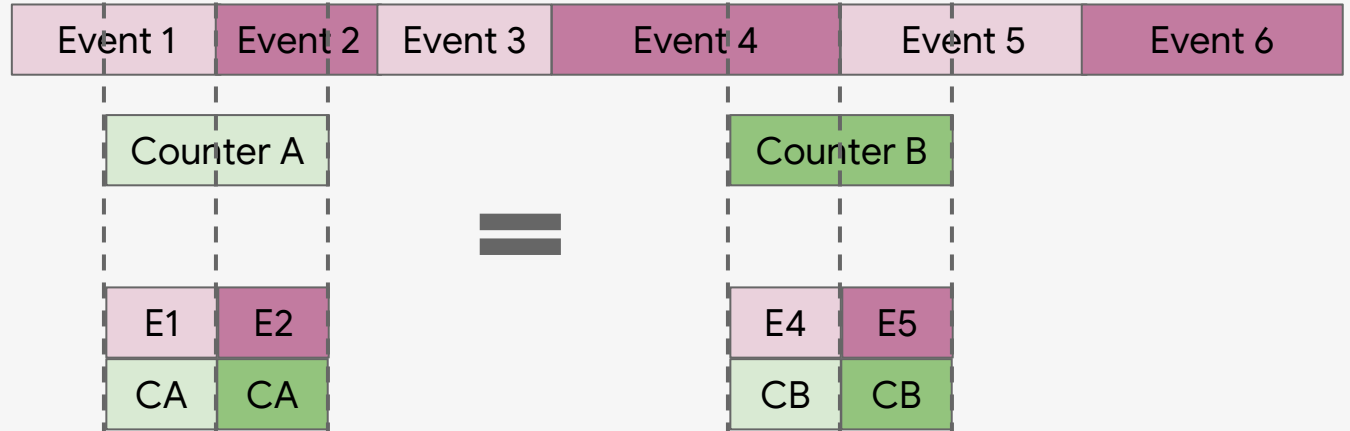
```
CREATE VIRTUAL TABLE bounds USING window;
```

```
UPDATE bounds SET quantum=Z where 1
```

```
CREATE VIRTUAL TABLE quantized USING span(sched, bounds)
```

# New constructs

## Harmonization



```
CREATE VIRTUAL TABLE quantized USING span(sched, counters, cpu)
```



Join key

UI

# Perfetto UI

Re-written from scratch from the ashes of chrome://tracing

Web-based: TypeScript + WebAssembly running in a worker

All the processing / analysis engine is based on the Trace Processor

Supports ~5 GB traces (limited by browser renderer limit)

URL: <https://ui.perfetto.dev>

Or just build it from sources and run locally.

Traces

- Open trace file
- Open example trace
- Record new trace
- Share current trace

Workspaces

Custom and pre-arranged views

Tracks and views

Add new tracks to the workspace

Metrics and auditors

Add new tracks to the workspace

Tracks



Flame Graph



# Thanks for your attention

For docs / links:

[www.perfetto.dev](http://www.perfetto.dev)

primiano@google.com