Containers and Performance Co-Pilot

Nathan Scott
Performance Tools, Red Hat
January 2015

Outline

- Performance Co-Pilot (PCP)
  - Overview
  - PCP Basics
- Kernel Instrumentation
  - namespaces
  - cgroups
- Containers in PCP

Overview

- What is PCP?
  - Open source **toolkit**
  - System-level analysis
  - Live and historical
  - Extensible (monitors, collectors)
  - Distributed

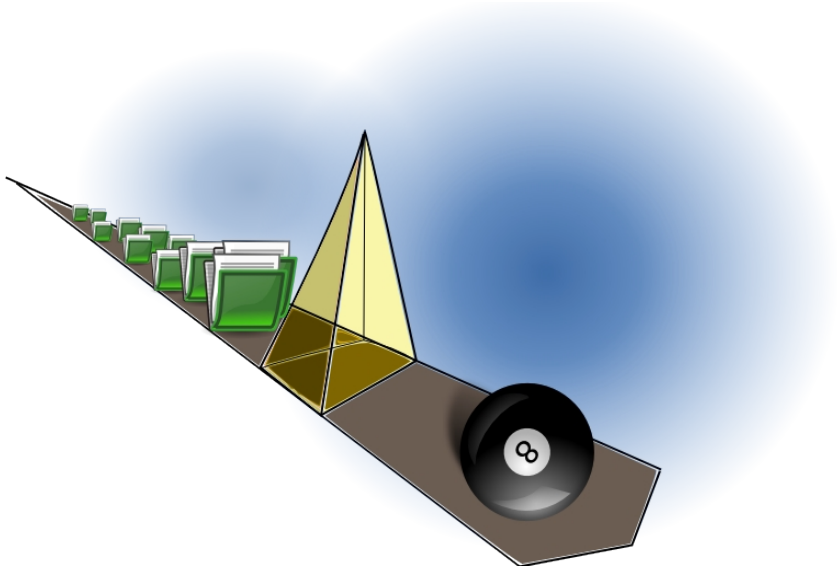

Architecture

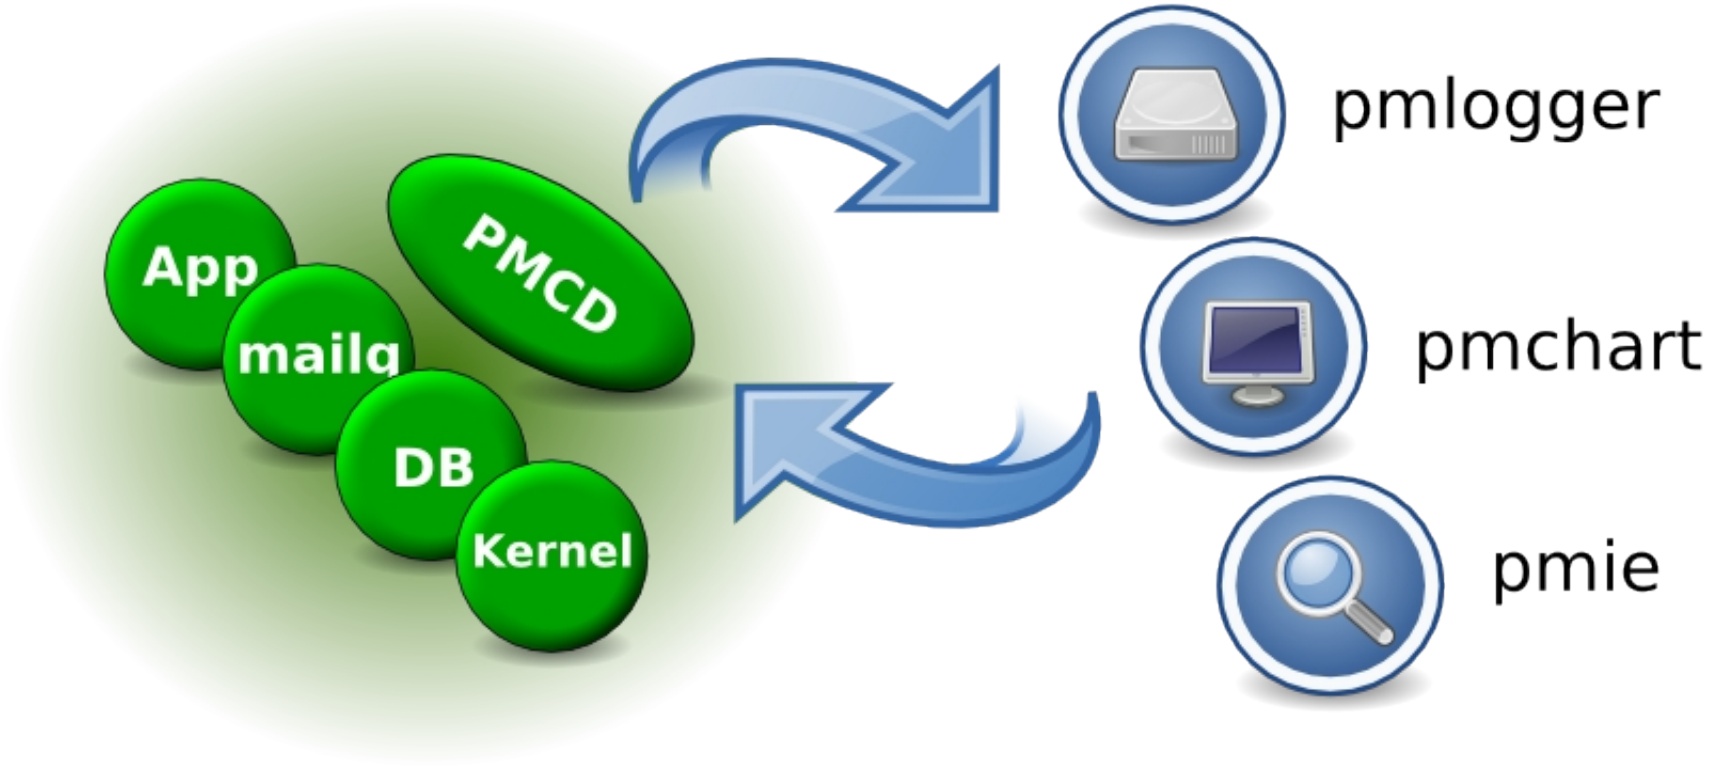

Metrics

- **pminfo** --desc -tT --fetch ***disk.dev.read***

  disk.dev.read [*per-disk read operations*]

  Data Type: *32-bit unsigned int* InDom: *60.1*

  Semantics: *counter* Units: *count*

  Help: *Cumulative count of disk reads since boot time*

  Values:

  inst [0 or "***sda***"] value ***3382299***

  inst [1 or "***sdb***"] value ***178421***

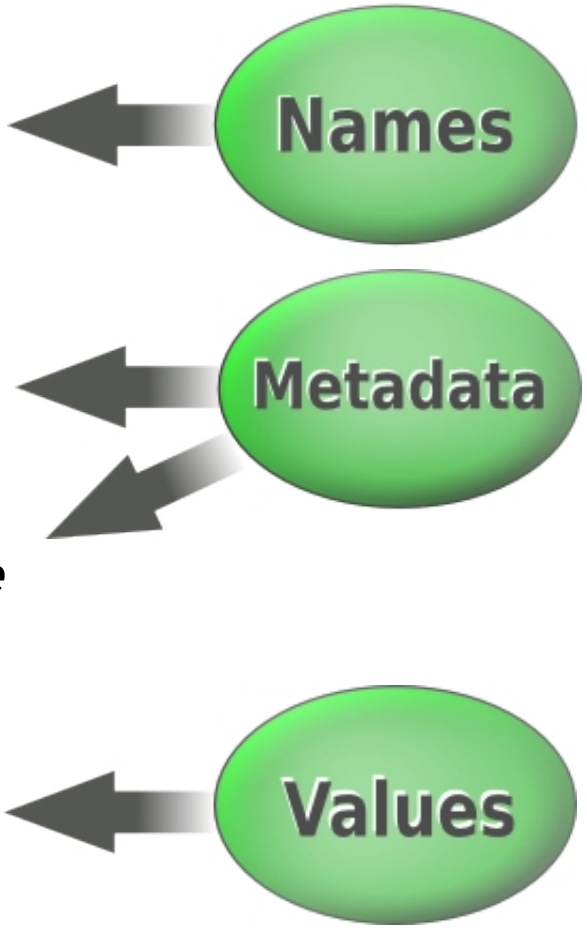


Kernel Instrumentation
for Containers

cgroup accounting

- [subsys].stat files below /sys/fs/cgroup
  - individual cgroup or summed over children
- **blkio**
  - IOPs/bytes, service/wait time – aggregate/per-dev
  - Split up by read/write, sync/async
- **cpuacct**
  - Processor use per-cgroup - aggregate/per-CPU
- **memory**
  - mapped anon pages, page cache, writeback, swap, active/inactive LRU state

namespaces

- Example: **cat /proc/net/dev**
  - Contents differ inside vs outside a container
- Processes (e.g. **cat**) in containers run in different network, ipc, process, uts, mount namespaces
- Namespaces are inherited across fork/clone
  - Processes within a container share common view

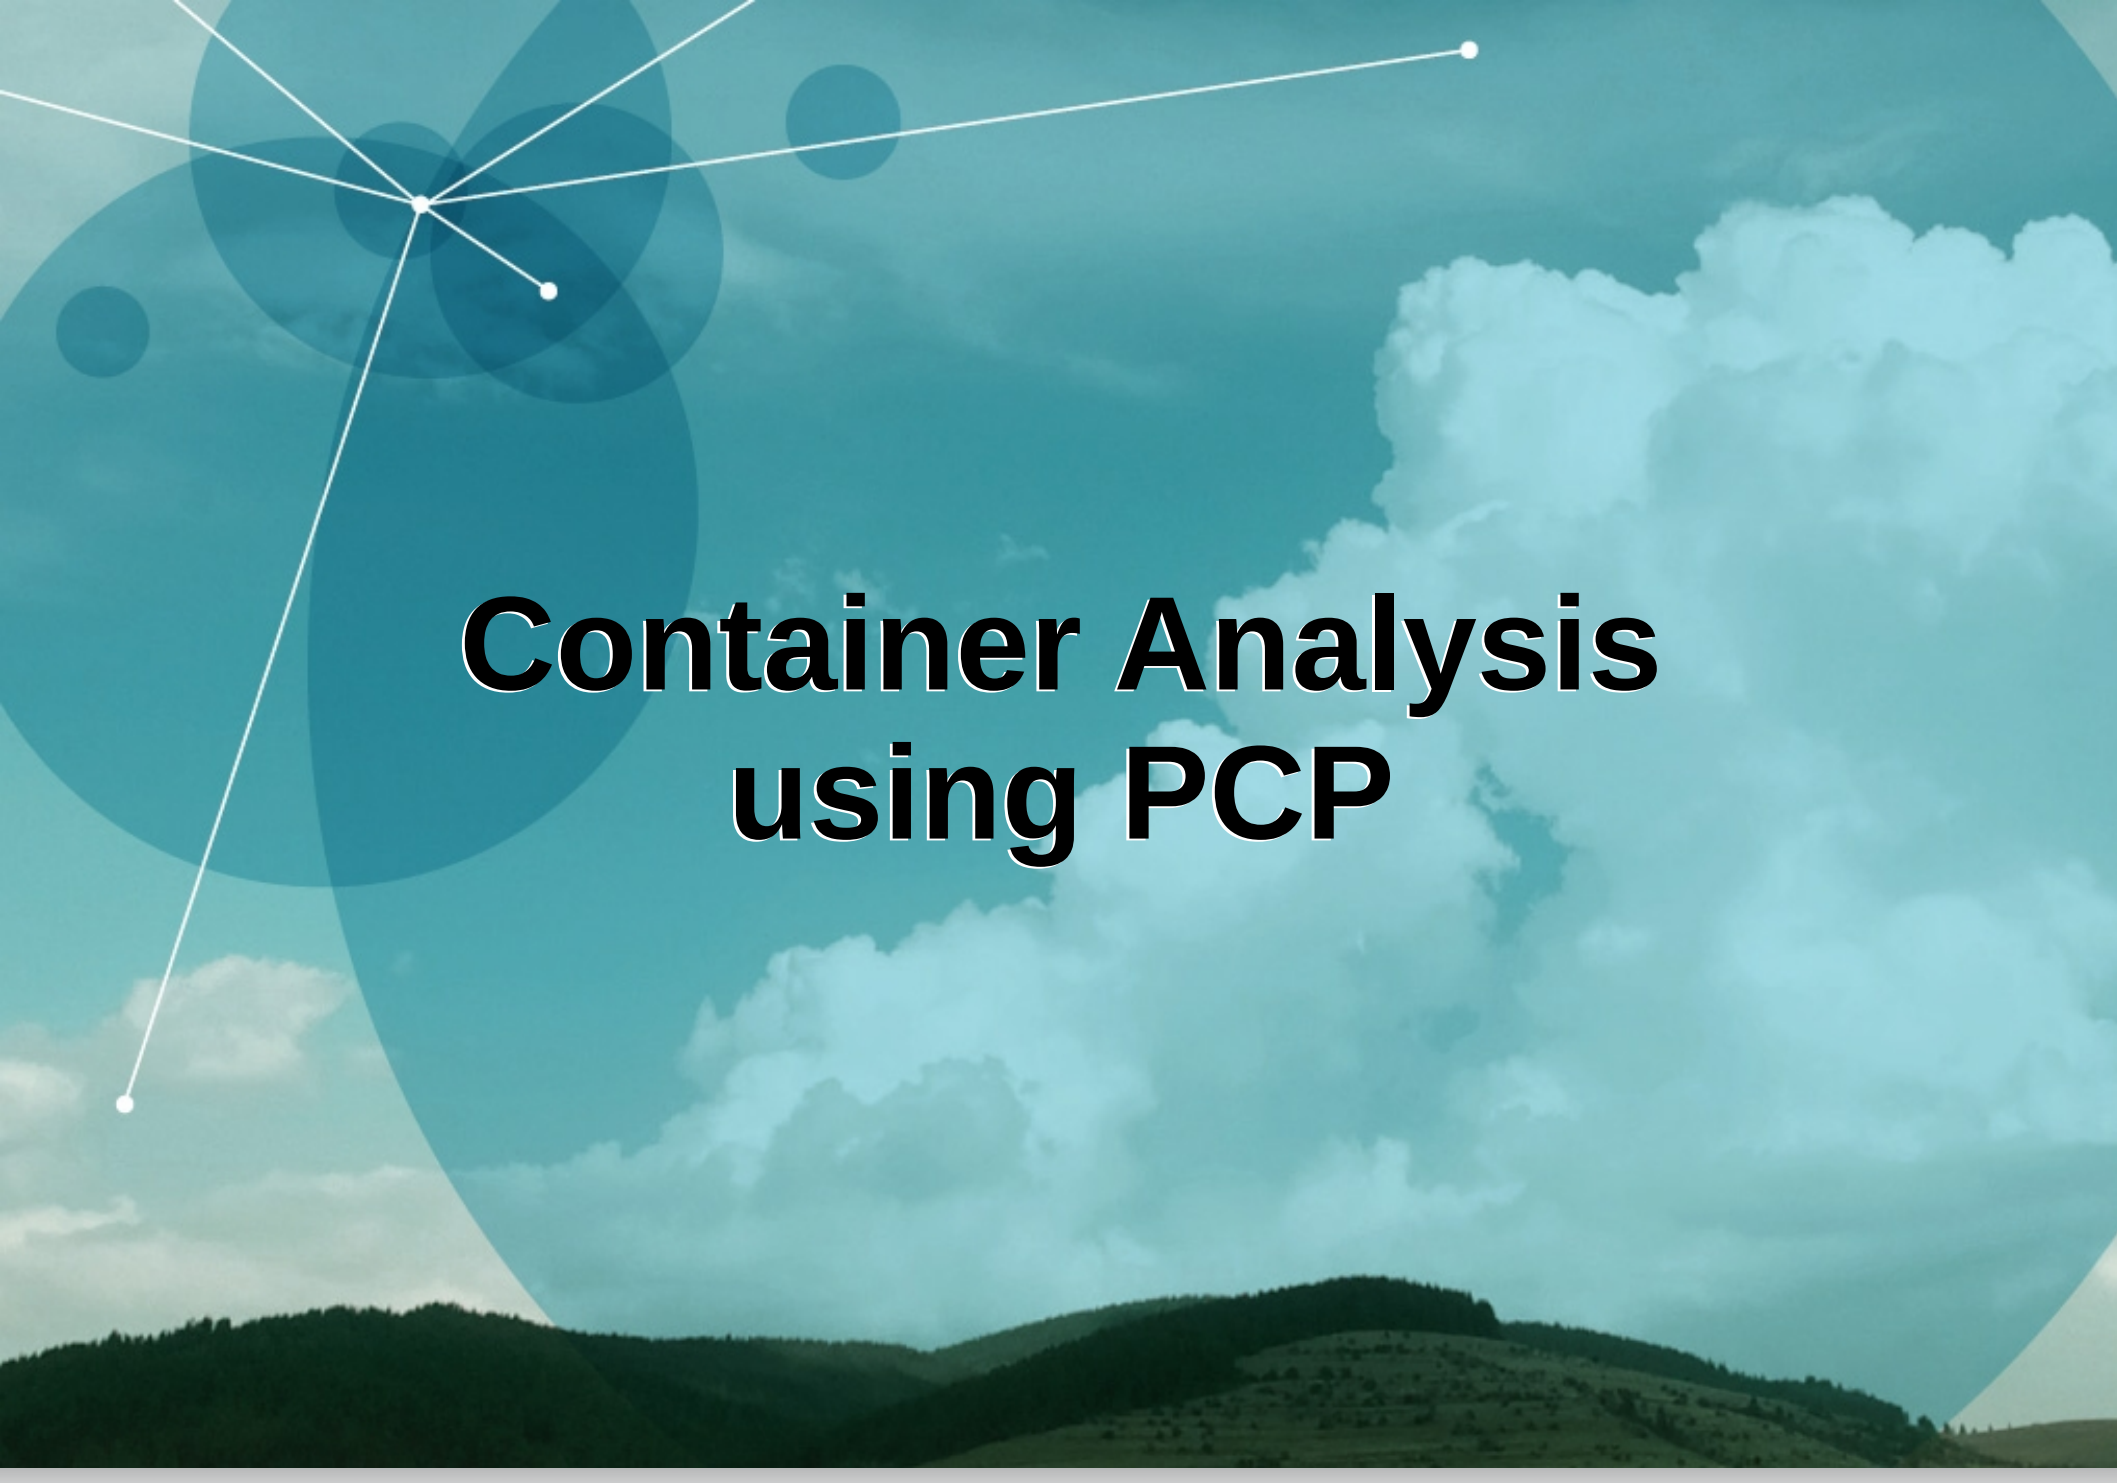
Container Analysis
using PCP

Goals

- Allow targeting of individual containers
- e.g. ***/proc/net/dev***
  - pminfo --fetch network

    vs
  - pminfo –fetch –container=crank network

- Zero installation inside containers required
- Simplify your life (dev_t auto-mapping)
- Data reduction (proc.*, cgroup.*)

Mechanisms

- pminfo -f –host=**acme.com** –container=**crank** network
  - Wire protocol extension
  - Inform interested PCP collector agents
  - Resolving container names, mapping names to cgroups, PIDs, etc.
  - setns(2)
- Runs on the board, plenty of work remains
- New monitor tools with container awareness

Resources

http://www.pcp.io

Source, downloads, books, FAQ, mailing lists

git://git.pcp.io/pcp

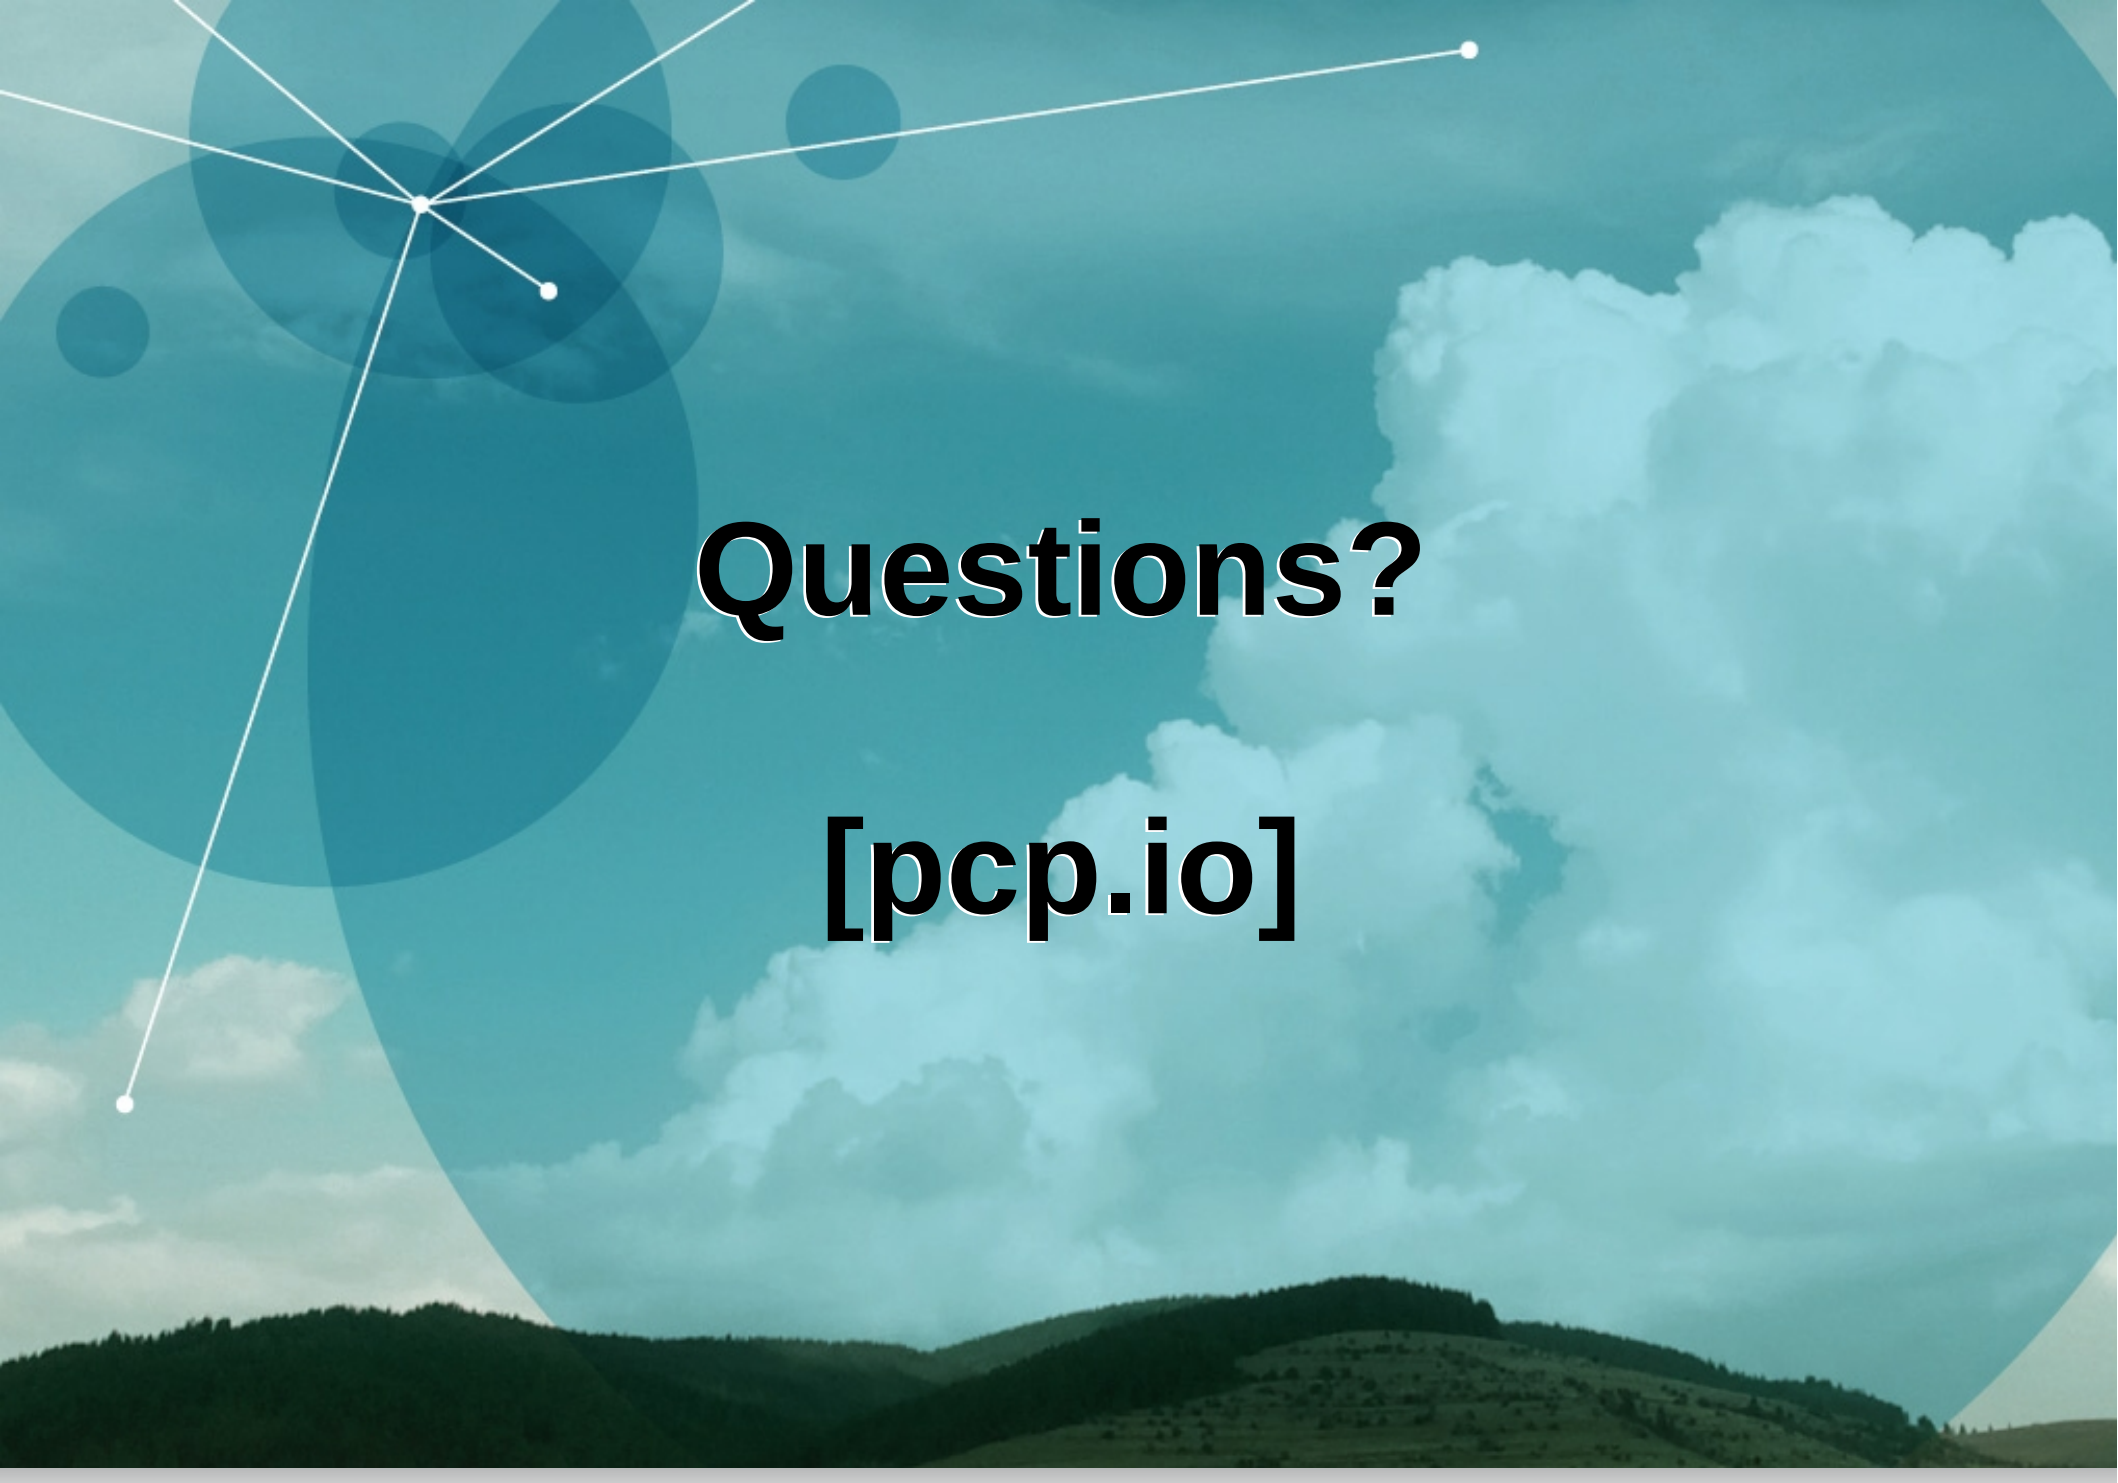
Questions?
[pcp.io]