



Running Your Datacenter with PostgreSQL



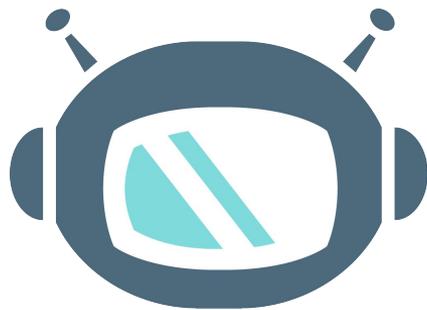
Presented by: Tony Perez



Tony Perez

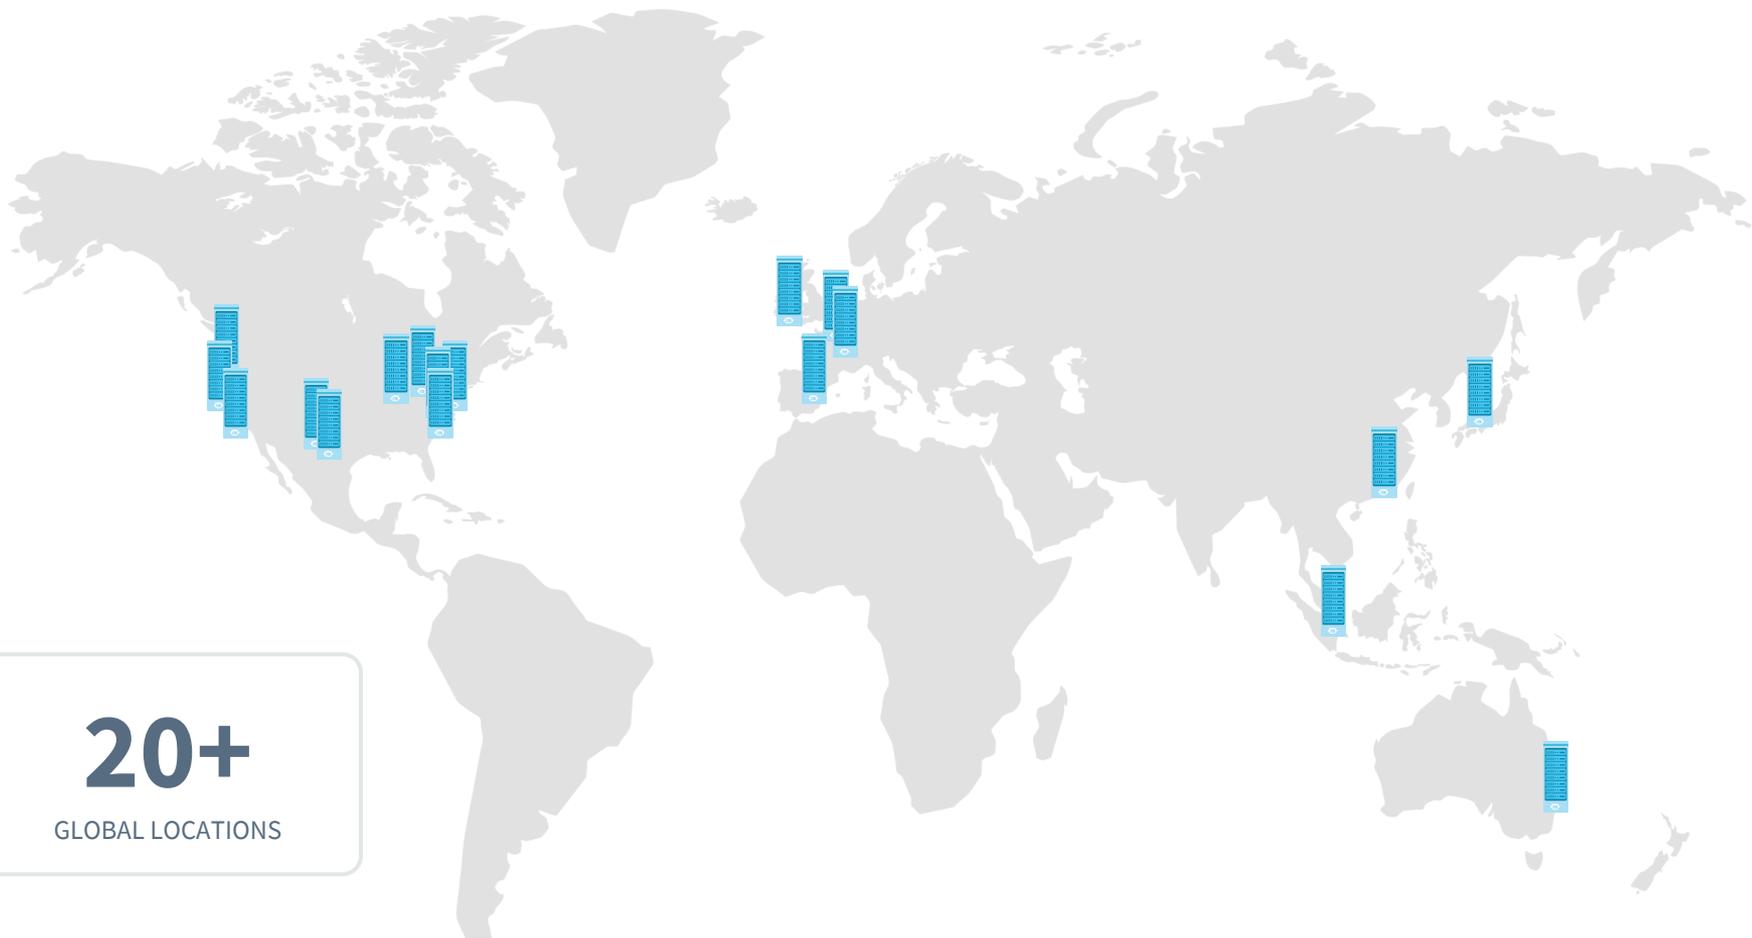
- Infrastructure Engineer at Packet
- Network and Infrastructure Nerd

Fun fact: *Tony's has been at Packet for ~3 years and it's his first startup – it's also his first full time job!*



packet

Built for Developers, Loved by Enterprise

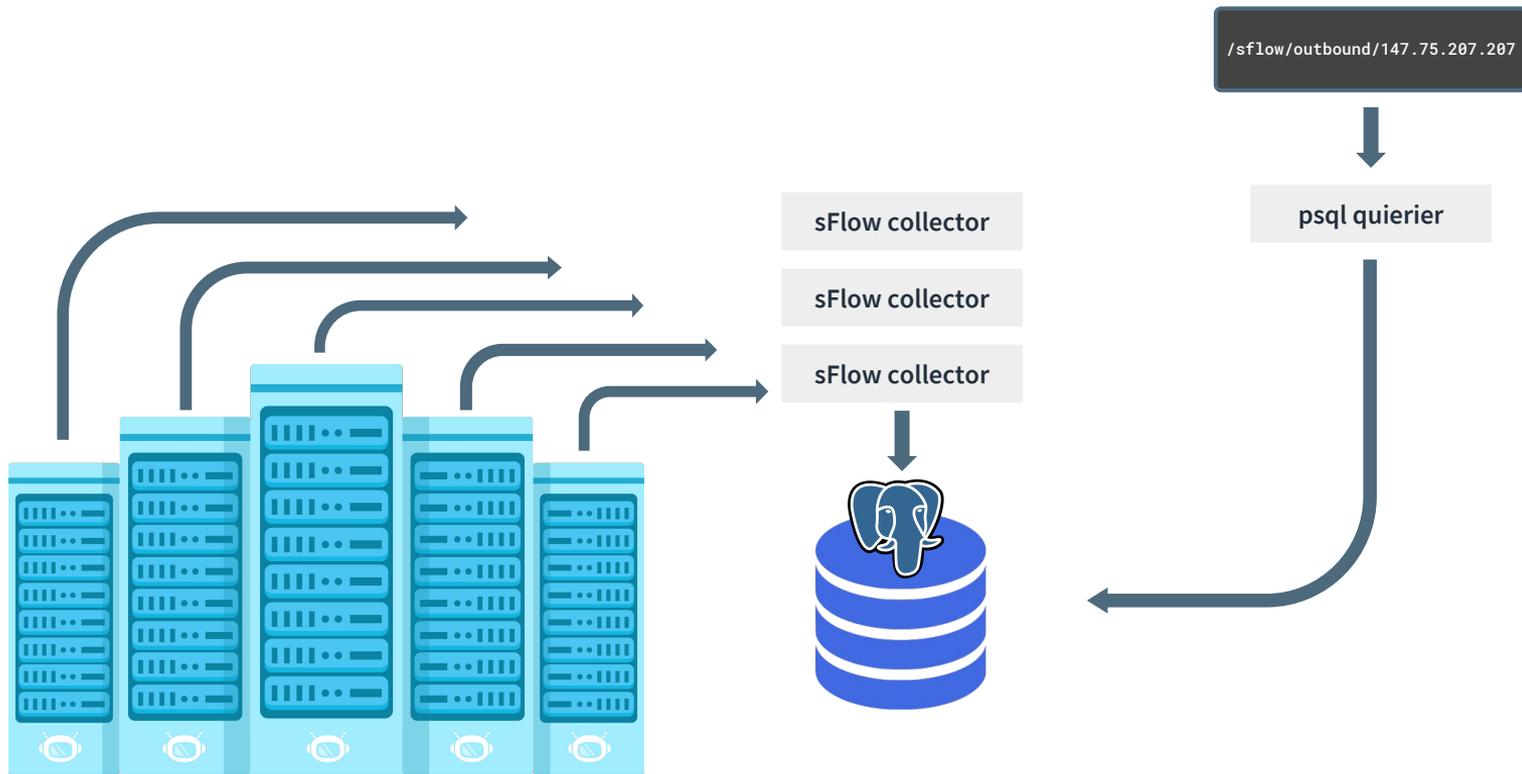


20+
GLOBAL LOCATIONS

How do we use PostgreSQL?

- We use PostgreSQL for all of our microservices:
 - Packet API
 - IPAM(IP Address Management)
 - Bandwidth Billing
 - etc.
- For monitoring we use Prometheus(the [wrouesnel/postgres_exporter](#))
- Dashboarding is done thru Grafana

Flow Data Pipeline



What is sFlow?



ToR Switch



mac_src		ac:1f:6b:82:e1:10
mac_dst		dc:38:e1:50:7b:89
ip_src		10.99.252.199
ip_dst		10.100.238.23
port_src		80
port_dst		19596
packets		2048
bytes		3108864
stamp_inserted		2016-06-08 01:30:00+00

sFlow Data



SFlow Parser

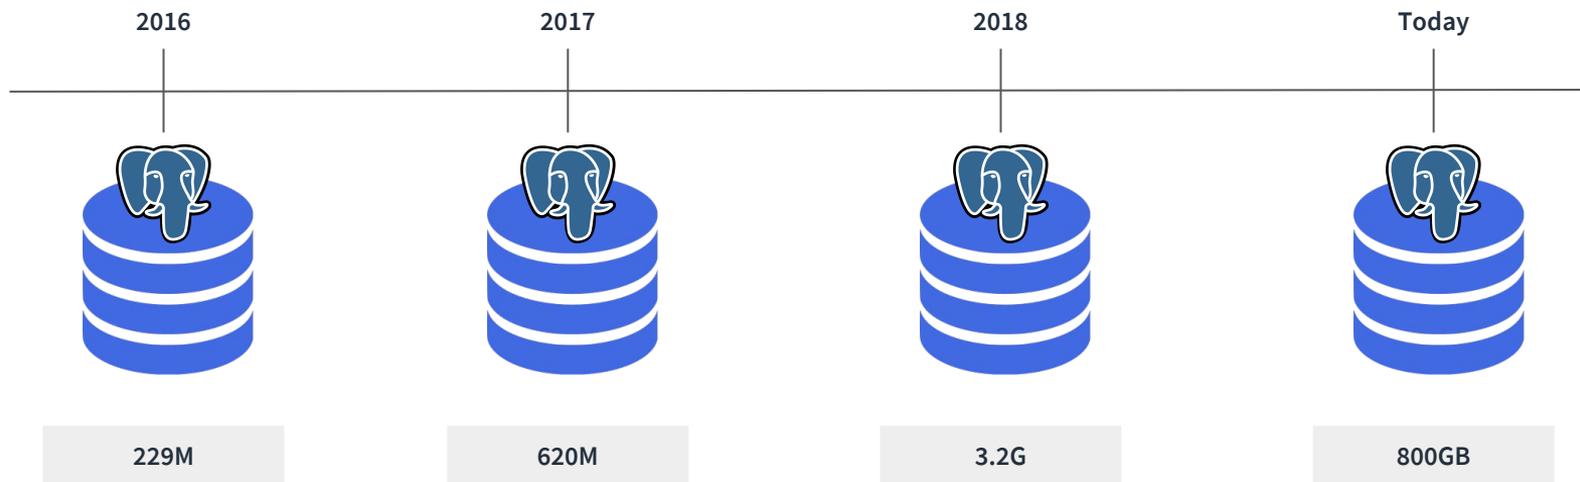


Data Store

pmacctd schema

Column	Type	Collation	Nullable	Default	Storage	Stats target	Description
tag	bigint		not null	0	plain		
class_id	character(16)		not null	' '::bpchar	extended		
mac_src	macaddr		not null	'00:00:00:00:00:00'::macaddr	plain		
mac_dst	macaddr		not null	'00:00:00:00:00:00'::macaddr	plain		
vlan	integer		not null	0	plain		
as_src	bigint		not null	0	plain		
as_dst	bigint		not null	0	plain		
ip_src	inet		not null	'0.0.0.0'::inet	main		
ip_dst	inet		not null	'0.0.0.0'::inet	main		
port_src	integer		not null	0	plain		
port_dst	integer		not null	0	plain		
tcp_flags	smallint		not null	0	plain		
ip_proto	smallint		not null	0	plain		
tos	integer		not null	0	plain		
packets	integer		not null		plain		
bytes	bigint		not null		plain		
flows	integer		not null	0	plain		
stamp_inserted	timestamp with time zone		not null	'0001-01-01 00:00:00'::timestamp without time zone	plain		
stamp_updated	timestamp with time zone				plain		

PostgreSQL Data Growth



Improvements

- Stay up to date with the latest stable releases of PostgreSQL
- Take advantage of new features like partitioning.
- Orchestrating 40+ postgres server requires your database to be part of your platform and not a one off server.
- Large tables get harder and harder to debug when there are query issues and so figuring out a way to aggregate is key.

Questions?

tony@packet.com /  packethost /  <https://slack.packet.com>