

Traffic analysis with NetFlow

Paolo Lucente

<paolo at pmacct dot net>

<http://www.pmacct.net/>

Traffic analysis with NetFlow

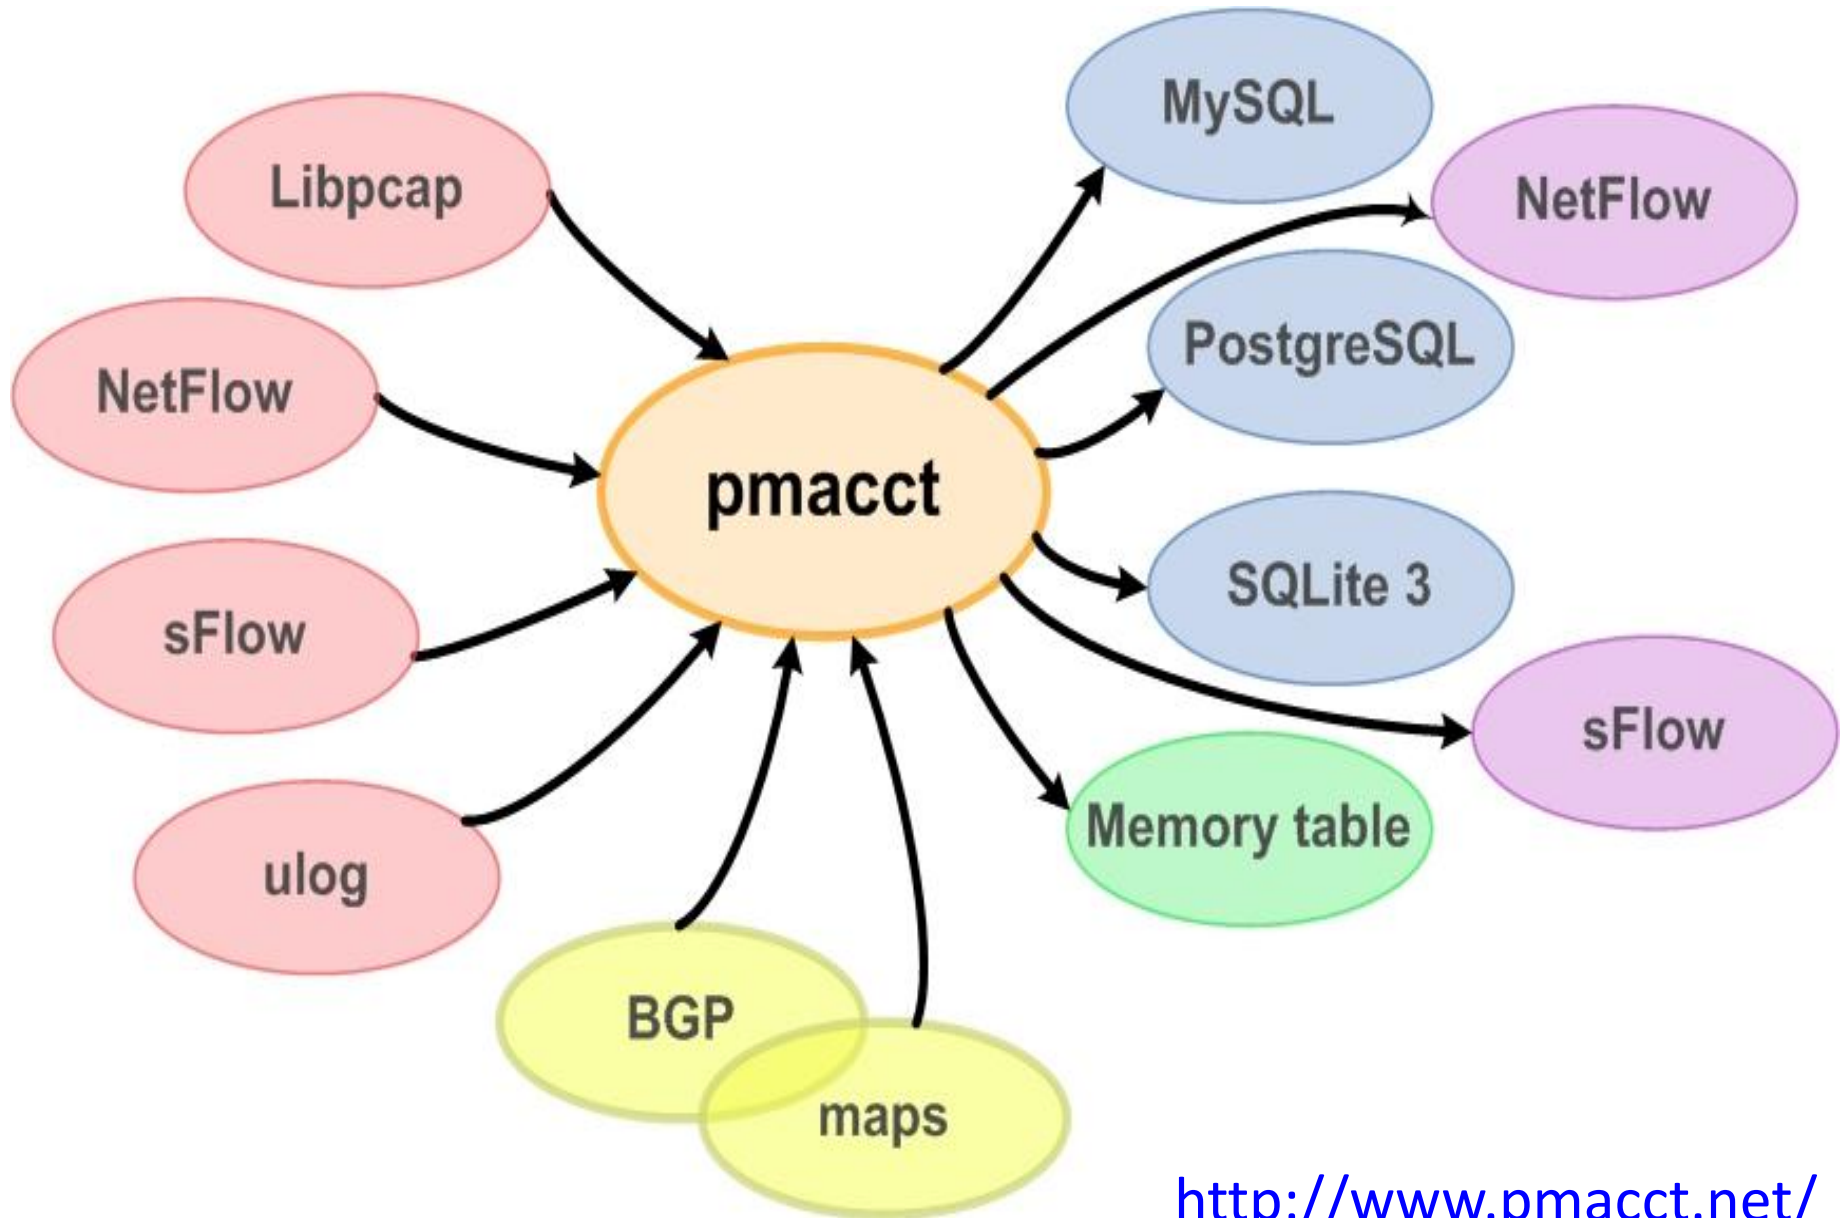
Agenda

- **whoami: Paolo & pmacct**
- Ramblings: Square 0 to reporting

whoami: Paolo

- Been working for operators for a while
- Been involved with IP accounting for a while
 - Hence I stumbled upon NetFlow a while ago 😊
- Within operators, NetFlow is beneficial in several contexts, ie.:
 - Traffic engineering
 - Capacity planning
 - Peering
 - ...
 - and also (ie. not only) security

pmacct is open-source, free, GPL'ed software



<http://www.pmacct.net/>

Traffic analysis with NetFlow

Agenda

- whoami: Paolo & pmacct
- **Ramblings: Square 0 to reporting**

Square 0

- NetFlow is not the only existing export protocol
- Current trend is to build extensible protocol formats by means of templates (NetFlow v9, IPFIX) or modules (sFlow)
- Export protocols, versions and templates (or modules) available are solely mandated by the underlying network hardware, ie.:
 - Cisco? NetFlow (and NetFlow-Lite on switches?)
 - Juniper? NetFlow and IPFIX (so baaaad as we speak!) on M/T/MX; sFlow on EX
 - Brocade? sFlow!

Square 1

- Essentially, top-down approach
- Do:
 - Start with a vision
 - Start with goals to achieve
- Don't:
 - Collect just for the sake of
 - Collect because one day raw data will be useful ..
- *Goals drive to requirements*

Requirements (1/2)

- Given goals, network topology, hardware, etc.
 - Define what devices (routers, switches, ..) need to export
 - Define what interfaces need to report
 - Define which direction, inbound or outbound or both?
 - Define whether to sample or not and if yes how much
- *Reality checks: don't just assume hardware can follow decisions; verify completeness of traffic sources*

Requirements (2/2)

- Is NetFlow (or similars) going to do the job alone?
 - Correlation with BGP benefits peering and IP transit analysis (profitability, costs, violations, etc.)
 - Correlation with IGP (plus estimation methods) benefits traffic engineering and capacity planning
- Is there any margin for use of data reduction techniques?
 - Micro-flow information and as much as possible 1:1 sampling rate benefits security applications and R&D
- *Requirements drive to choice of tooling and storage (ie. flat-files, RRD, RDBMS, etc.)*

Collector implementation - gotchas

- Organizational
 - 800 pound gorilla project spanning across different departments (including IT and Security where existing)
- Technical
 - Probes loosing traffic along the way, ie.:
 - Cisco 7600 TCAM space being overwhelmed
 - Juniper MX series requiring extra hardware, a MS-DPC (\$\$\$), to cope with sustained export load
 - Over-engineered protocol features (ie. NetFlow v9, IPFIX sampling) drive to bizarre, creative implementations
- *Verify expectations and perform consistency checks against other sources, ie. interface counters*

Collector implementation - scalability

- Divide-et-impera approach, ie.:
 - Assign probes to collectors
 - Assign collectors to storages, ie. RDBMS
- Work on data reduction:
 - Increase sampling rate
 - Strip down the aggregation method
 - Fit data into larger time-bins
- Remove from the equation interfaces which are not strictly essential to the task, ie. low-speed
- *Do take into account on larger deployments that a single collector might not fit all the bill due to, typically, CPU or memory constraints*

Storage method selection

- Plenty on the offer – both on the commercial and open-source side: flat-files, RDBMS, RRD, column-oriented DB, memory, etc.:
 - I do not come here with general statements
 - In fully integrated products (ie. backend + frontend) this is transparent to end-users
 - Otherwise, going open access (which btw some fully integrated products also offer):
 - Depends, ie. whether a well-known query language like SQL is a pro; an indexed storage is a pro; an “have it all” to the micro-flow level kind of storage is a pro; etc.
 - Often it’s matter of personal preferences, ie. what one feels most comfortable to query once in production

Maintenance (1/2)

- (Script to) verify the right set of data is hitting the collector:
 - Ingress + egress directions can lead to duplicates
 - Backbone + edge interfaces can lead to duplicates
 - Network infrastructure evolves: mistakes and forgetting is around the corner
- (Script to) track collector hardware resources:
 - Supports the divide-et-impera approach
 - It's good rule to do it anyway

Maintenance (2/2)

- (Script to) keep data-set size under control:
 - Valid for both memory and disk-based storage methods
 - Drop older data is less complex than consolidating; if going for inexpensive storage, ie. disk, consider collecting same data at different time resolutions, with different expiration periods

Reporting

- Reporting via email or web is popular
- Key is interaction between backend and frontend:
 - If open access is granted to data-set, great stuff!
 - Otherwise, make sure new reports can be (easily) generated as requirements emerge
 - Reporting can influence organization of the backend, ie. RDBMS re-indexing
- *Flexibility must be possible (if Sales, Purchase, Product Management, etc. find it useful – be prepared to handle their creativity!)*

Traffic analysis with NetFlow

Thanks for your attention!

Questions?

Paolo Lucente

<paolo at pmacct dot net>

<http://www.pmacct.net/>