

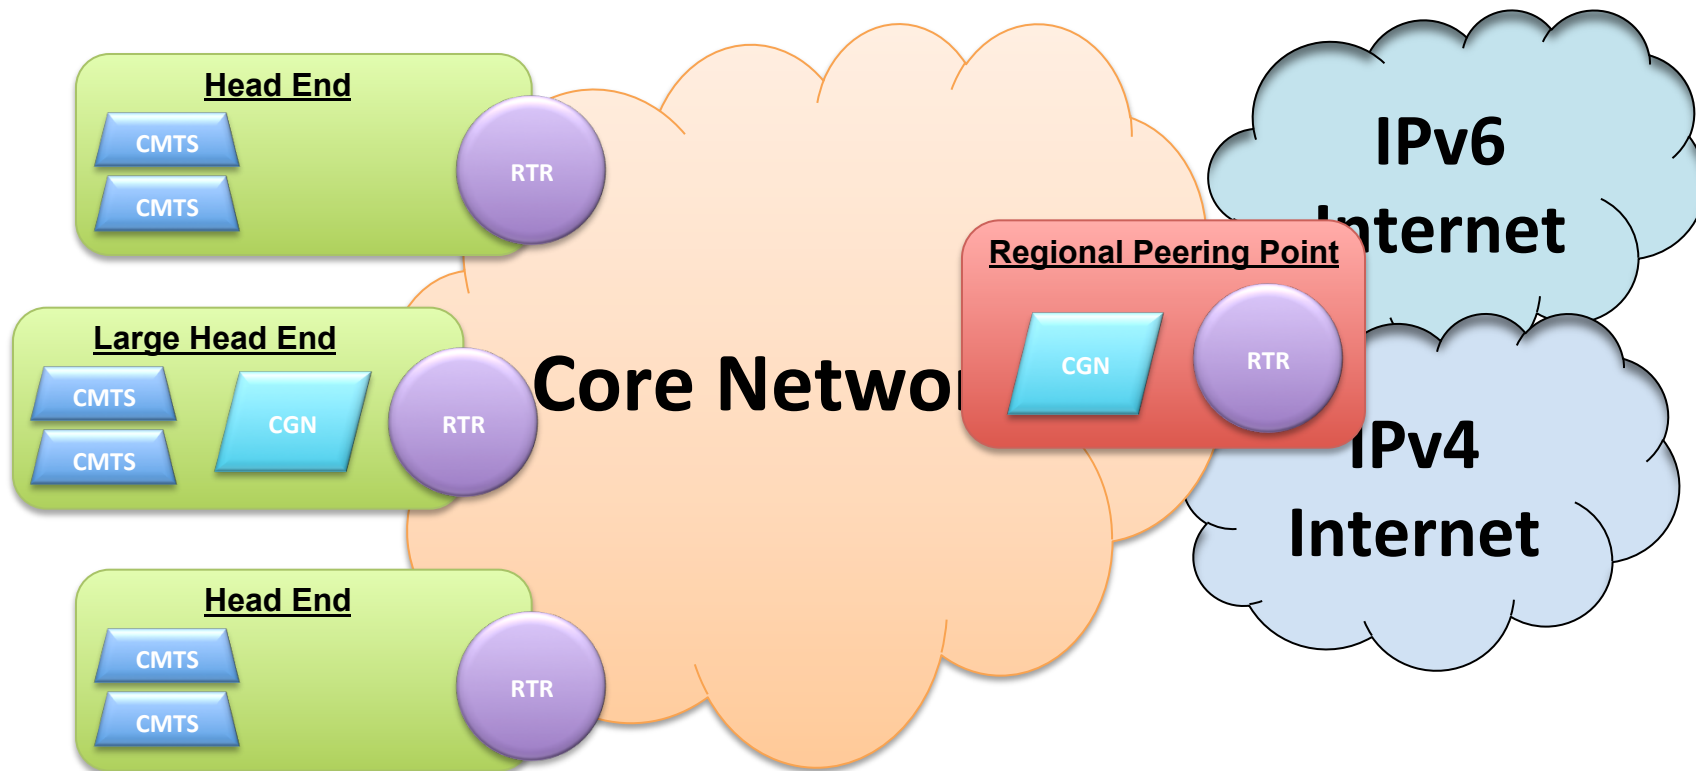
CGN Log Reduction Strategies

draft-donley-deterministic-cgn

Chris Donley | November 17, 2011

Background - CGN Deployment

- **CGNs may be deployed using a phased hybrid approach**
 - Deploy CGNs regionally to start
 - Add CGNs at large headends as needed
 - Rule-of-thumb: 50K subscribers/CGN



CGN Logging Generates Huge Data Volumes

- Subscribers generate around 33,000 connections/day
 - 150-byte log messages x 33,000 connections/day= 5MB/day/sub
 - 1M-sub ISP will generate 150 TB of data/month (1.8 PB/yr)
 - Assuming 50,000 subs/CGN => 23 Mbps of logging traffic
 - PER CGN!!!
 - Imagine DB search times...
- Conclusion: per-connection CGN logging is not scalable

Ports per subscriber

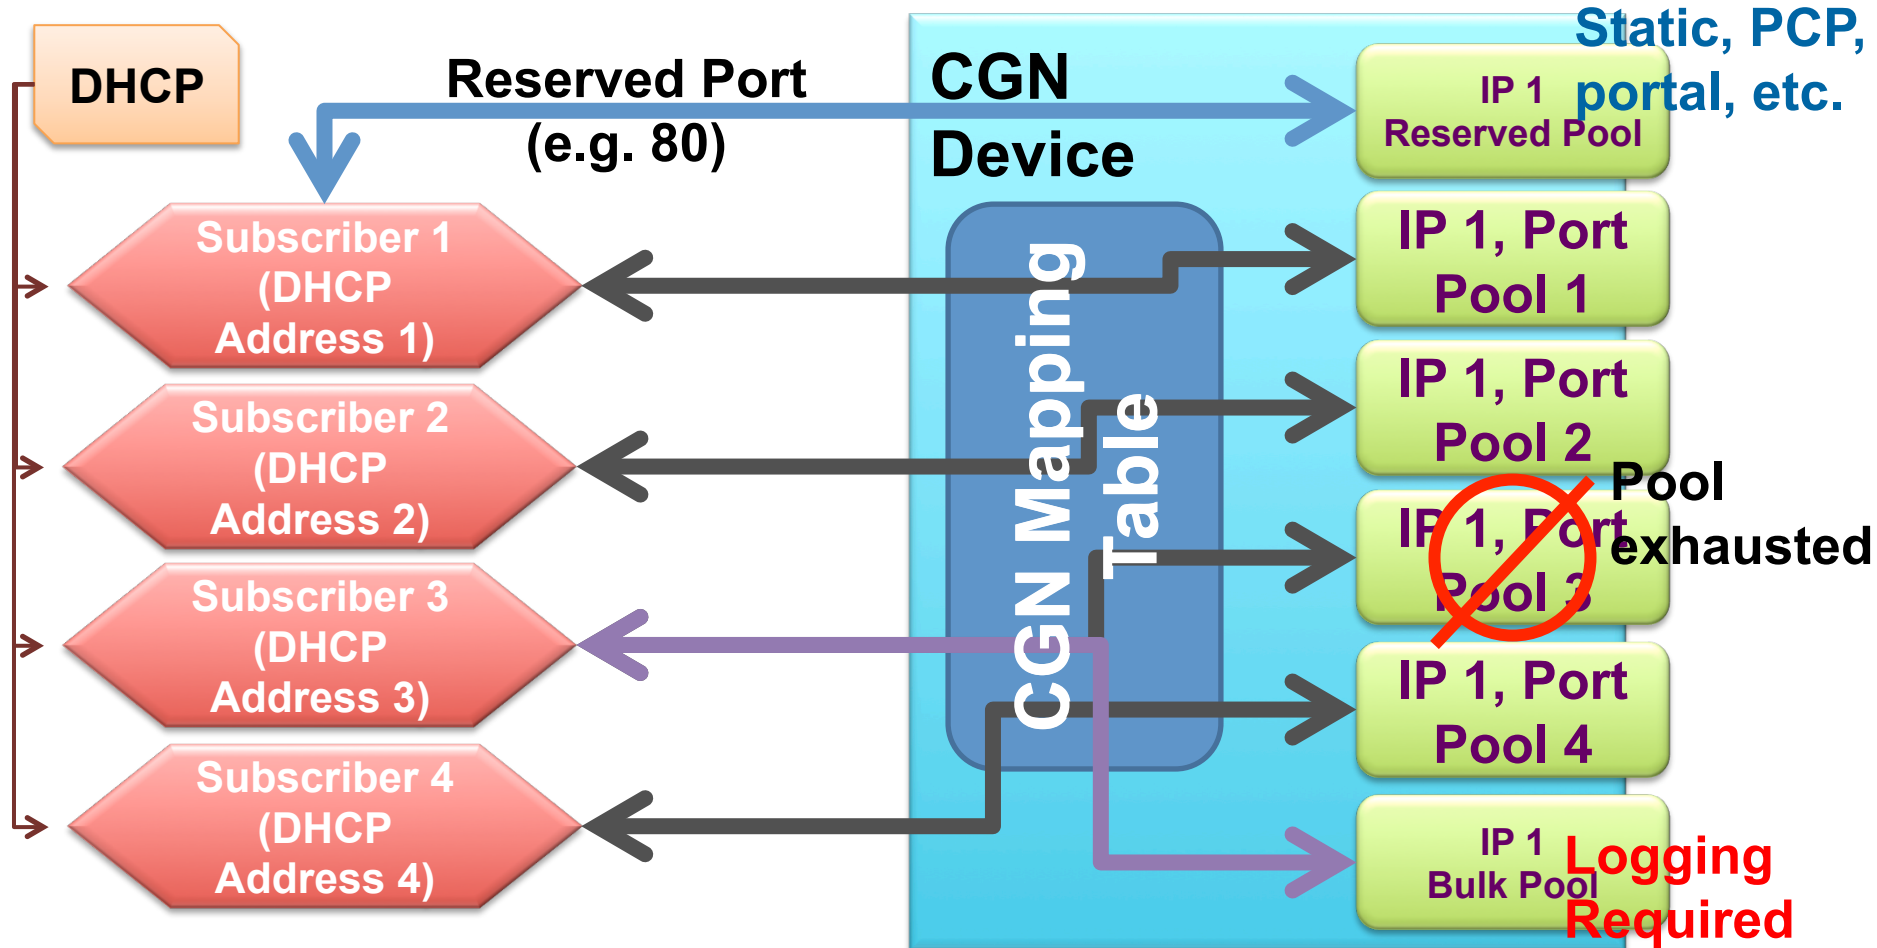
- [WAND \(NZ\) Research on CGNs](#)

	Outbound	Inbound
Max Mean Ports/ active subscriber	~700 (300 TCP/ 400 UDP)	~200 (60 TCP/140 UDP)
Max Median Ports/active subscriber	~70 (40 TCP/10 UDP)	~5
Peak TCP Ports/ active subscriber	99% <1000 90% <100	~15
Peak UDP Ports/ active subscriber	99% <2000 95% <1000 85% <100	~12

Deterministic CGN

- Uses an algorithm to allocate ports per subscriber
 - Abuse Response reverses algorithm to identify inside address
 - Eliminates the need for most logging (except for power users)
 - Reduces privacy impact on subscribers
- Logging only required for subs that exceed predefined port ranges
 - Uses bulk port logging for overflow log

Deterministic CGN Illustrated



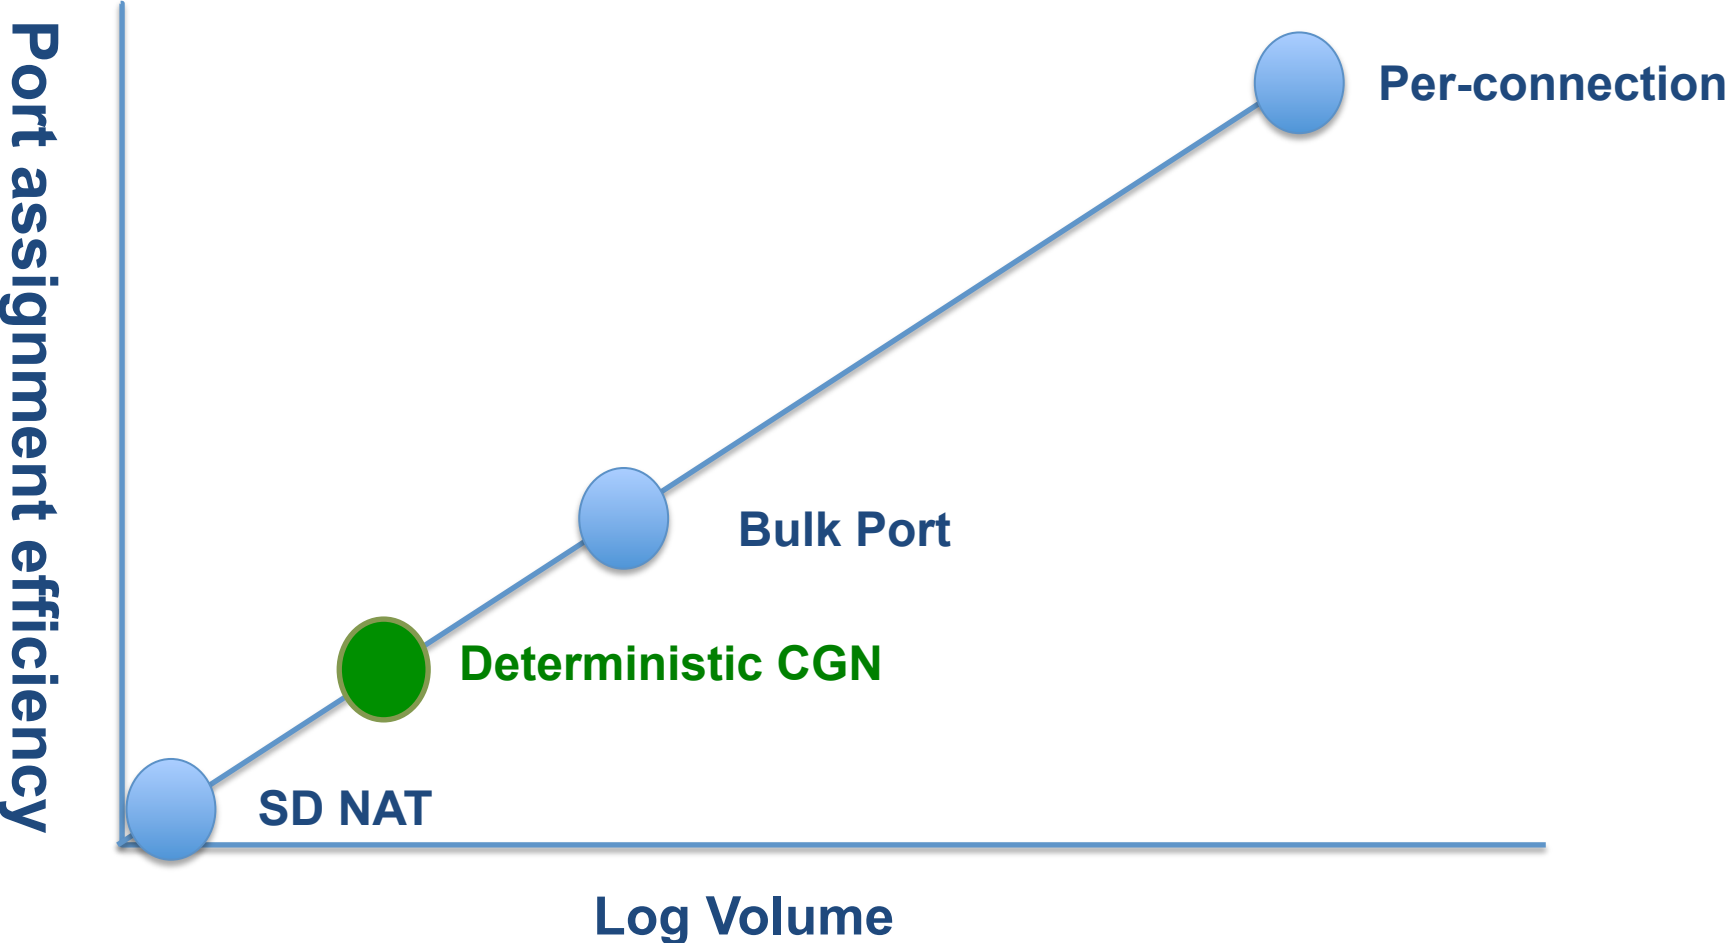
Example

- Assume 100 users; 33k connections/day/user; WAND power curve distribution
- Deterministic – 10:1 compression (~6k ports/subscriber)
- Bulk port assignment – 100 ports/block

	Standard	Bulk	Deterministic	SD-NAT
85 th percentile		85	0	0
90 th percentile		25	0	0
95 th percentile		50	0	0
99 th percentile		80	0	0
100 th percentile		100	40	0
Total	3,300,000	340	40	0

- Bulk - ~10,000x reduction
- Deterministic – ~100,000x reduction

Comparing log reduction approaches



Analysis

- bulk port logging, deterministic CGN, and SD-NAT significantly reduce log volumes
 - Bulk log messages approach number of active subs
 - Deterministic log messages approach 0
 - SD-NAT requires no logging
- This draft also compatible with other two approaches
 - Initial ports/user=0: bulk port
 - Overflow buffer=0: SD-NAT

Next Steps

- Update draft with mailing-list feedback
- WG draft?