Serbian Open eXchange

Ruting politike

Zoran Perović
iConsult@perke.NET
info@sox.rs

Serbian Open eXchange

Zašto tema ruting politike?

Zbog kvaliteta!!!

What Is BGP?

- **Border Gateway Protocol BGP-4**
- **The de-facto interdomain routing protocol**
- **BGP includes specifications:**
  - Which information gets advertised and how
- **BGP includes a routing protocol:**
  - Establishes and uses a routing table
- **Internal Gateway Protocol (I-BGP in the book)**

Why Is There Such Fuss about BGP?

- **BGP dictates routing at the AS level**
  - Absence of understanding: poor performance
- **BGP is complicated**
  - Designed to be flexible
  - Involves multiple fields
- **Understanding BGP behavior is not intuitive**
  - Implementation and business policies
- **The routing of the Internet relies on BGP**

Prefix Matching in More Detail

- **For a IP address of a packet, find longest match**
- **Example: Compare**
  - packet IP 128.32.101.1
  - With 128.32.0.0/16
  - IP : 01000000. 001000000. 01100101 .00000001
  - Mask : 11111111. 111111111. 00000000 .00000000
  - AND : 01000000. 001000000. 00000000 .00000000
  - Prefix : 01000000. 001000000. 00000000. 00000000
  - Equal? Yes

Serbian Open eXchange

L’art pour l’art

Basic AS relationships

- **Customer – Provider**
  - Customer pays Provider for service
  - The Customer is always right
- **Peer to Peer: mutual cooperation**
  - Ex. MCI and AT&T
- **Sibling-Sibling**
  - Ex. AT&T research and AT&T wireless

Basic AS Relationships

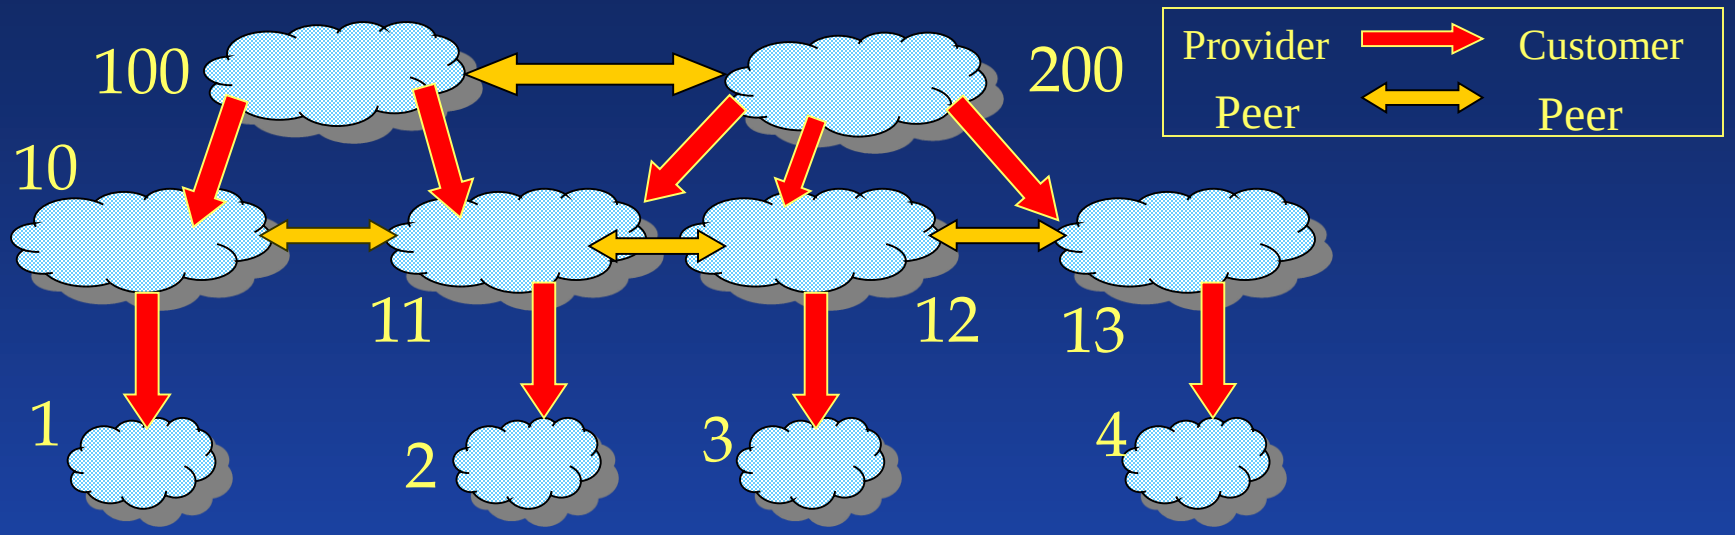

- **Customer – Provider:** customer pays and is always right
- **Peer to Peer:** Exchange traffic only between their customers
- **Sibling-Sibling:** Exchange traffic at will

How BGP Policy Restricts Routing

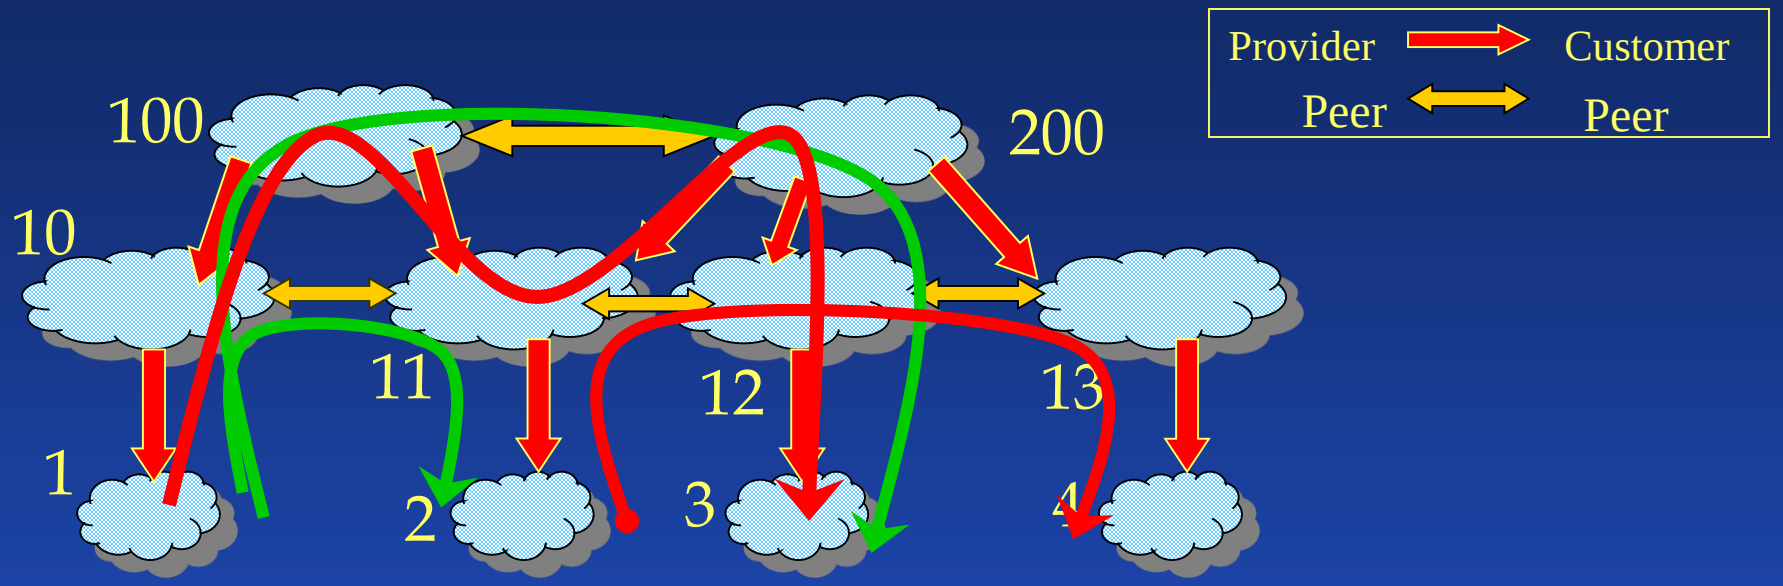

- **Routing rules:**
  - Provider accept everything
  - Peer only if it is for its customers
- **Path Properties:**
  - Up then down
  - No up-down-up, at most 1 peer-peer steps

Serbian Open eXchange

Kvalitet, a ne kvantitet

Serbian Open eXchange

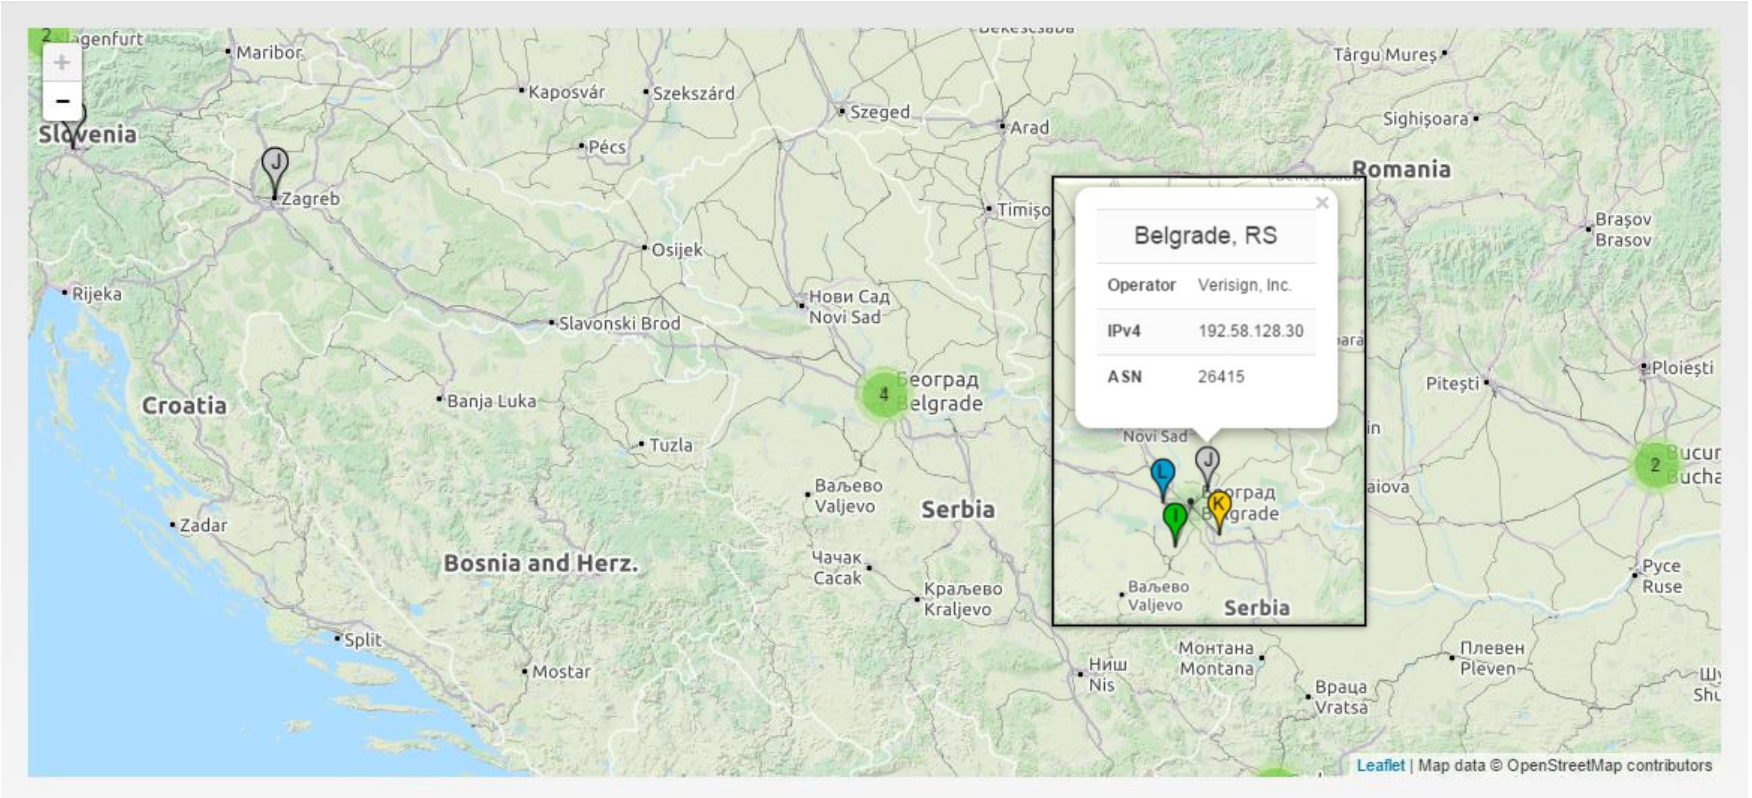

root DNS serveri – 4 od 13 u BGD-u, sa odzivom ispod 1 msec

Serbian Open eXchange

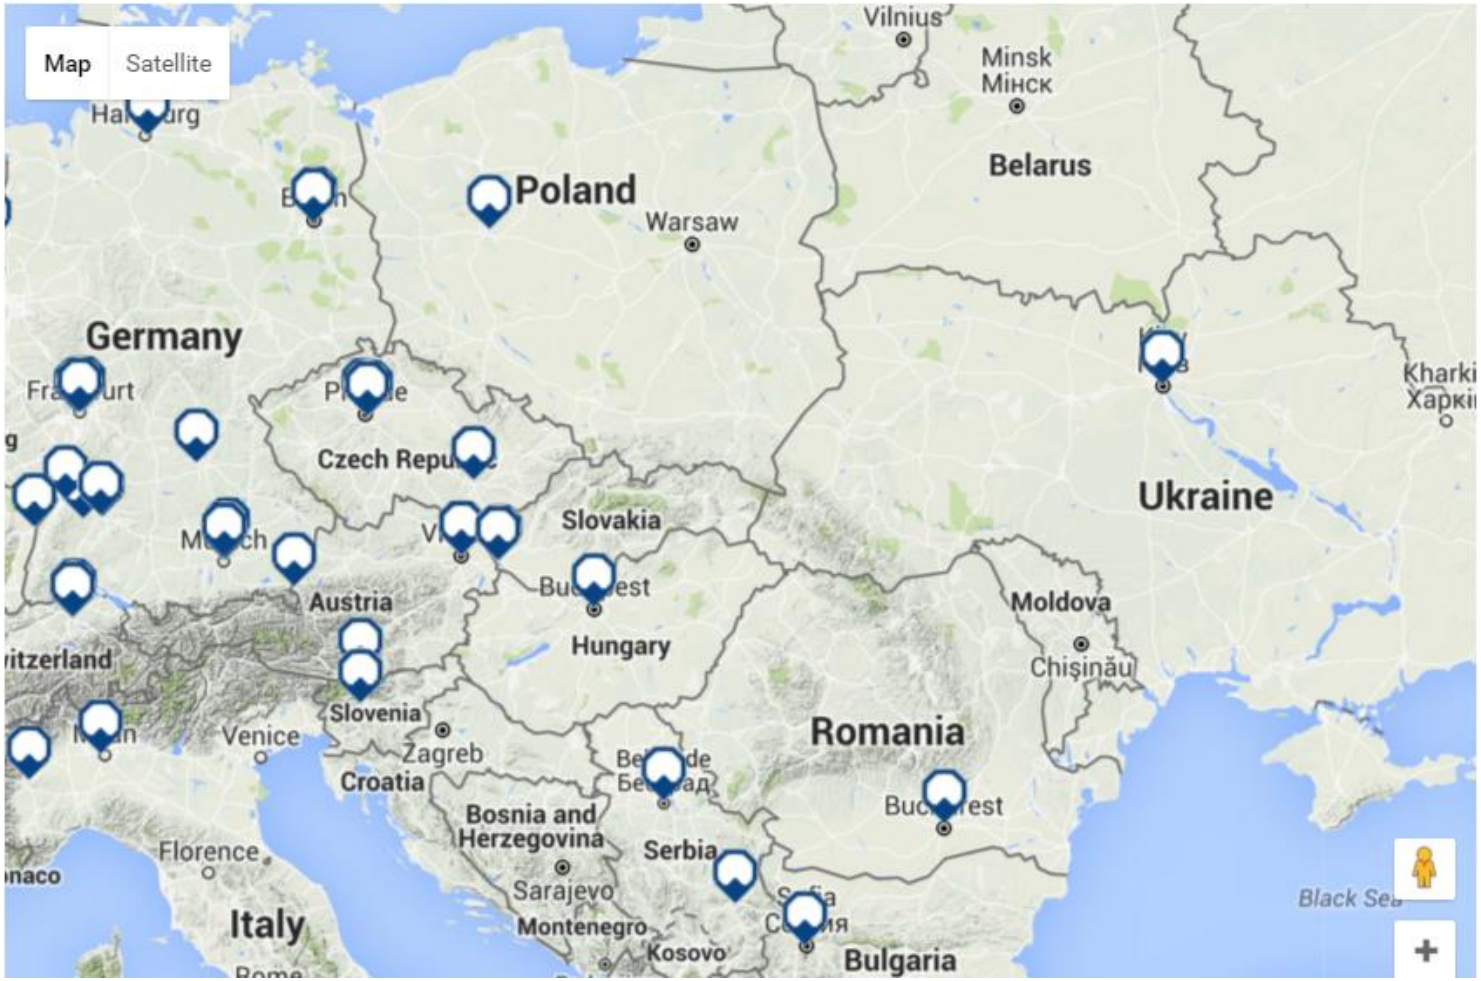

RIPE Atlas Anchor – 2 komada i nebrojeno sondi (probe)

Serbian Open eXchange

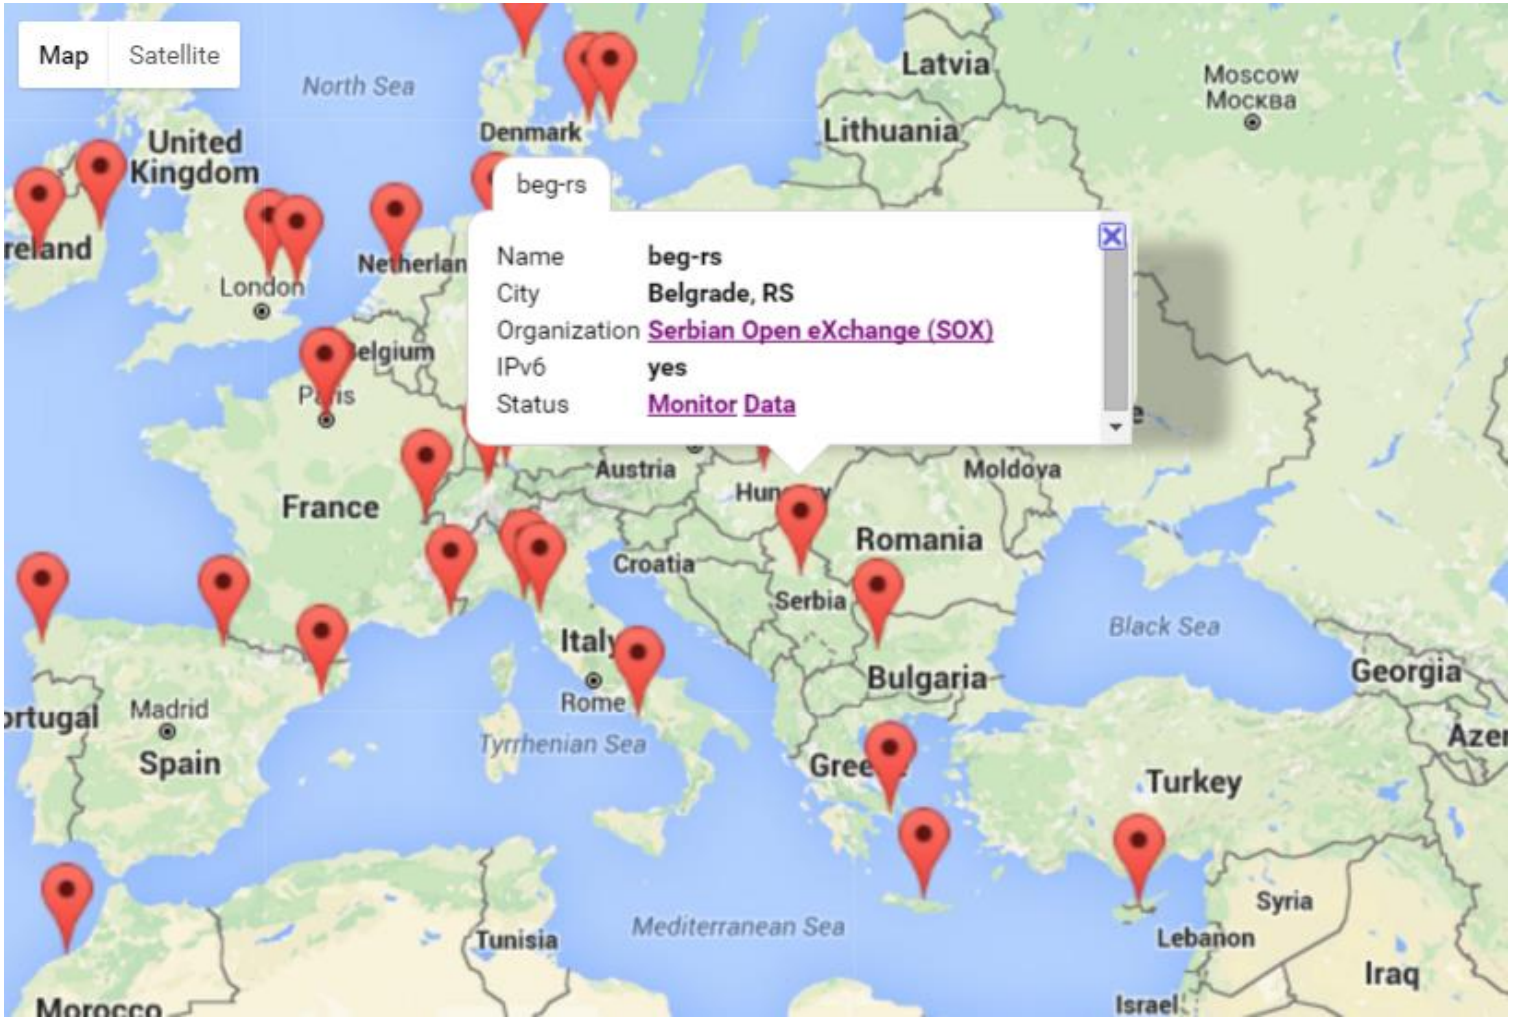

San Diego Supercomputing Centar – Archipelago projekt

SOX
Serbian Open eXchange

Serbian Open eXchange

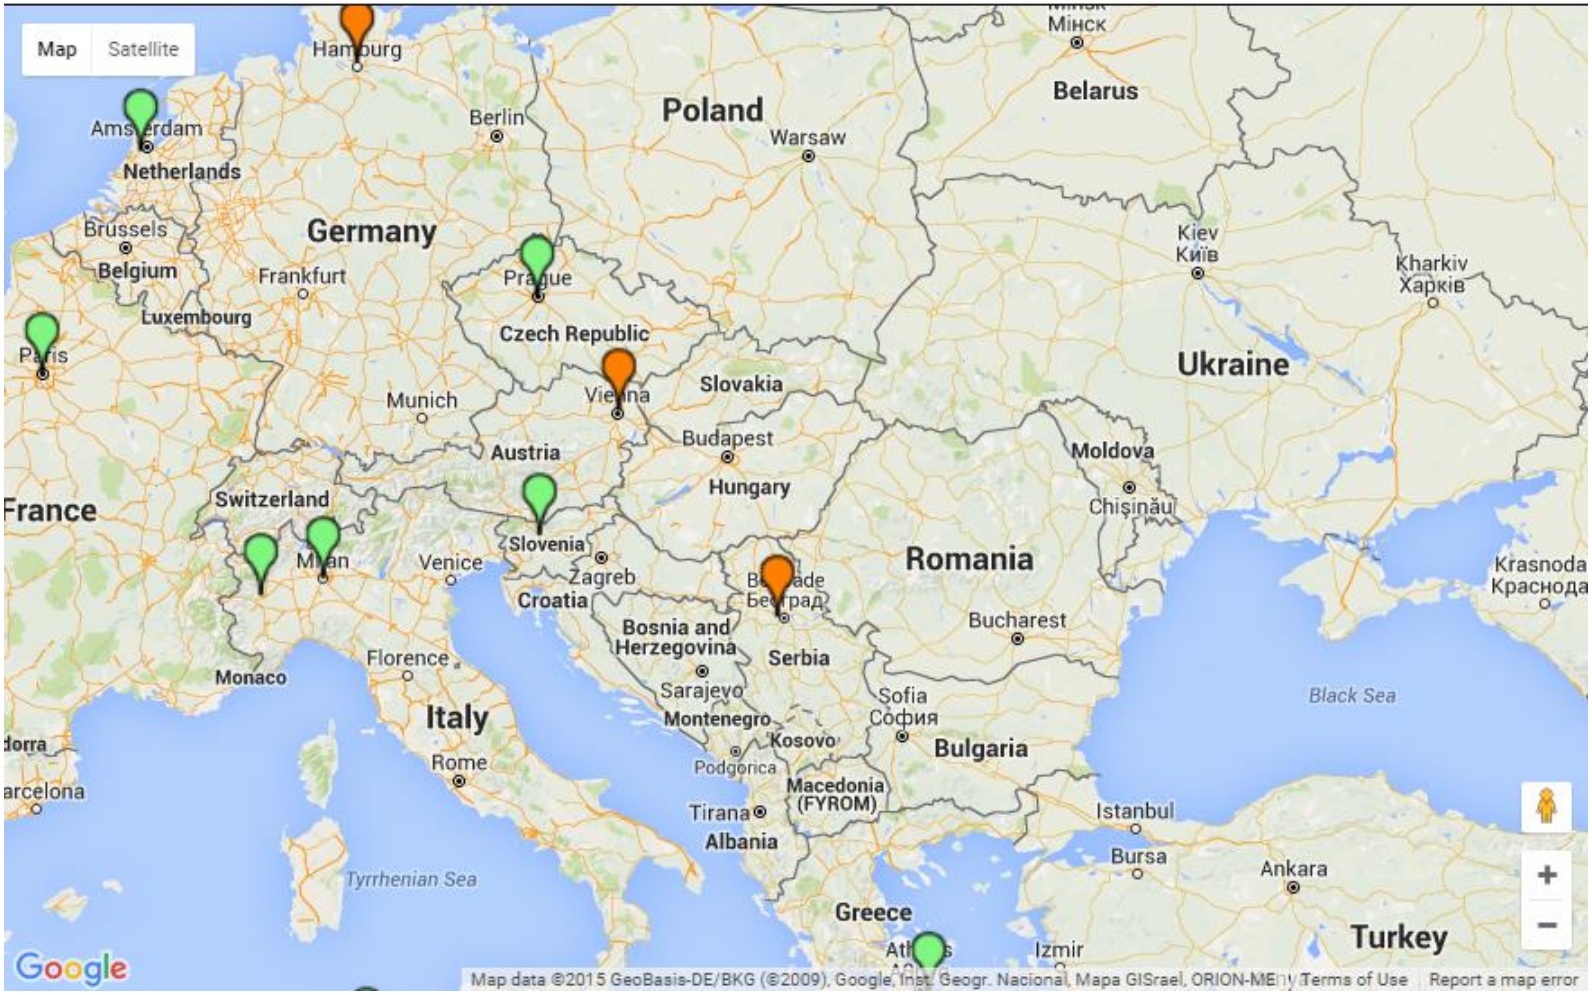

www.measurementlab.net

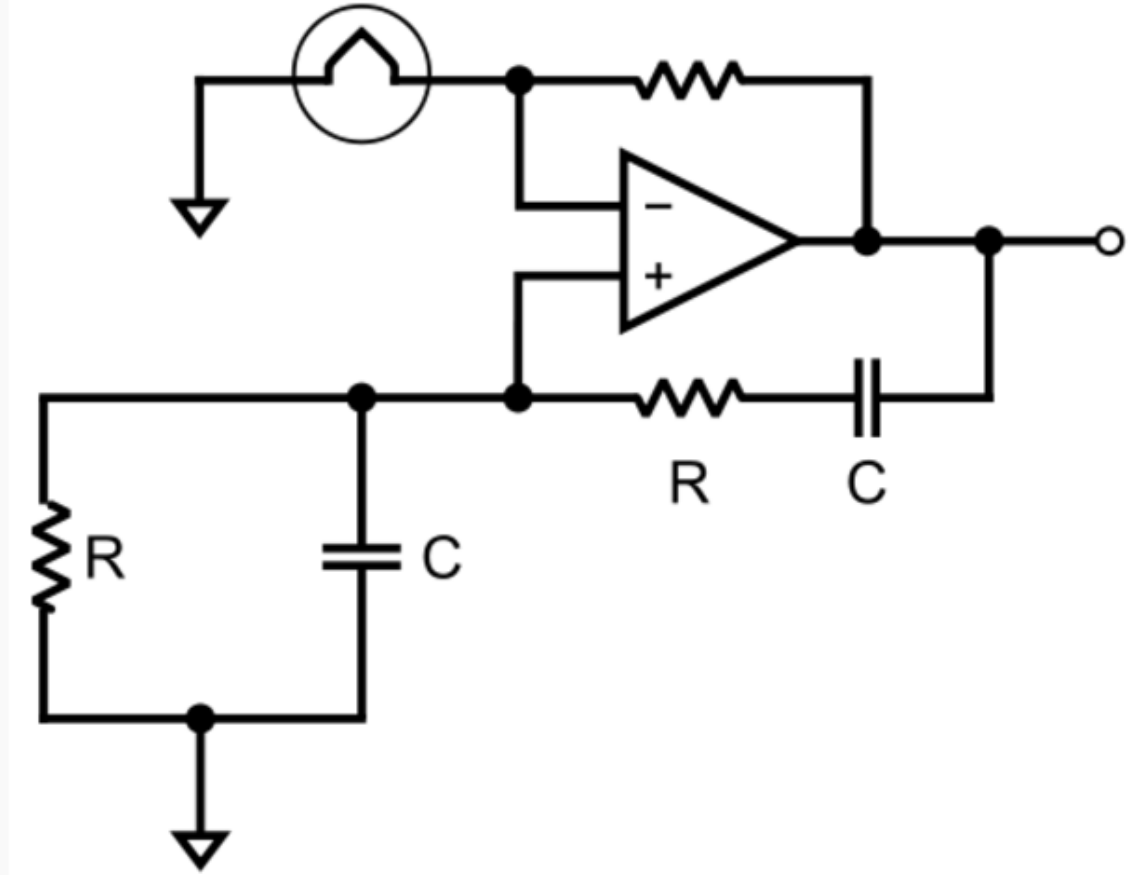

Vinov oscilator

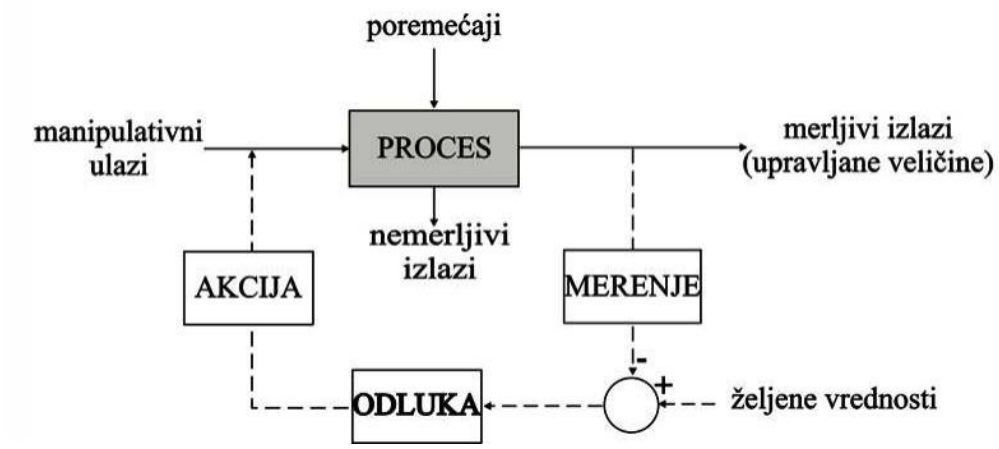
poremećaji
manipulativni ulazi
PROCES
merljivi izlazi (upravljane veličine)
nemerljivi izlazi
AKCIJA
MERENJE
ODLUKA
-
+
željene vrednosti

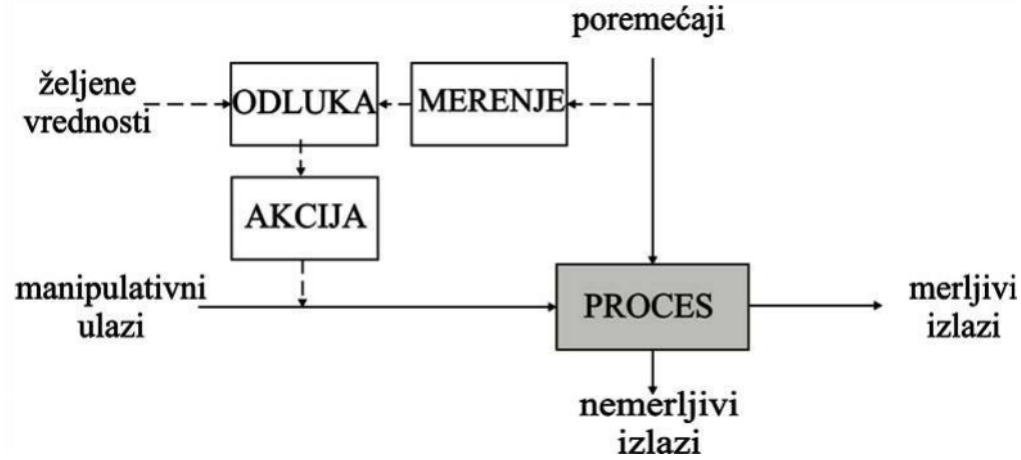
poremećaji
željene vrednosti
ODLUKA
MERENJE
AKCIJA
manipulativni ulazi
PROCES
merljivi izlazi
nemerljivi izlazi

Serbian Open eXchange

Waltzing on that gentle trade-off between internet routes and FIB space, an SDN story

David Barroso <dbarroso@spotify.com>

Paolo Lucente <paolo@pmacct.net>

RIPE SEE 4 meeting, Belgrade – Apr 2015

Serbian Open eXchange

pmacct

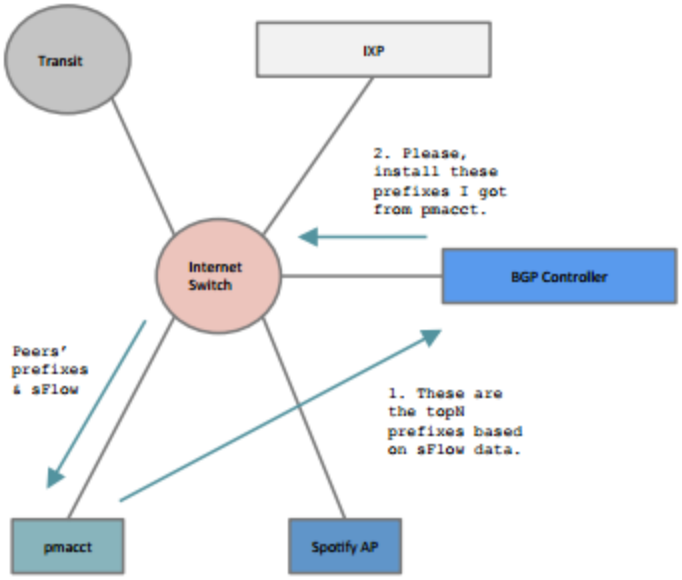

- pmacct aggregates sFlow data using the BGP information previously sent by the Internet Switch
- pmacct reports the TopN* prefixes to the BGP Controller
- The BGP controller instructs the Internet switch to install those TopN* prefixes

* N is a number close to the maximum number of entries that the FIB of the Internet Switch can support

Serbian Open eXchange

SDX: A Software Defined Internet Exchange

Arpit Gupta[†], Laurent Vanbever[⋆], Muhammad Shahbaz[†], Sean P. Donovan[†], Brandon Schlinker[‡]
Nick Feamster[†], Jennifer Rexford[⋆], Scott Shenker[◇], Russ Clark[†], Ethan Katz-Bassett[‡]
[†]Georgia Tech [⋆]Princeton University [◇]UC Berkeley [‡]Univ. of Southern California

Abstract

BGP severely constrains how networks can deliver traffic over the Internet. Today's networks can only forward traffic based on the destination IP prefix, by selecting among routes offered by their immediate neighbors. We believe Software Defined Networking (SDN) could revolutionize wide-area traffic delivery, by offering direct control over packet-processing rules that match on multiple header fields and perform a variety of actions. Internet exchange points (IXPs) are a compelling place to start, given their central role in interconnecting many networks and their growing importance in bringing popular content closer to end users.

To realize a Software Defined IXP (an "SDX"), we must create compelling applications, such as "application-specific peering"—where two networks peer only for (say) streaming video traffic. We

- **Routing only on destination IP prefix.** BGP selects and exports routes for destination prefixes. Networks cannot make more fine-grained decisions based on the type of application or the sender.
- **Influence only over direct neighbors.** A network selects among BGP routes learned from its direct neighbors, and exports selected routes to these neighbors. Networks have little control over end-to-end paths.
- **Indirect expression of policy.** Networks rely on indirect, obscure mechanisms (*e.g.*, "local preference", "AS Path Prepending") to influence path selection. Networks cannot directly express preferred inbound and outbound paths.

These problems are well-known, yet incremental deployment of alternative solutions is a perennial problem in a global Internet with more than 50,000 independently operated networks and a huge

Serbian Open eXchange

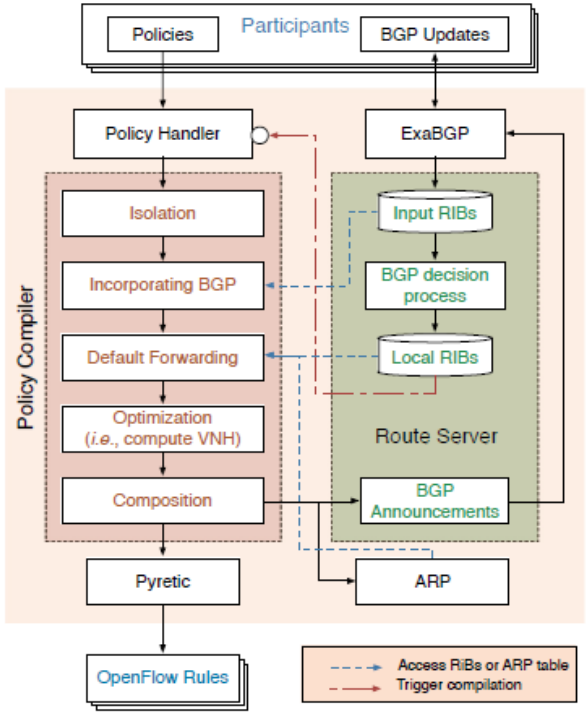

Figure 3: *The SDX controller implementation, which has two pipelines: a policy compiler and a route server.*

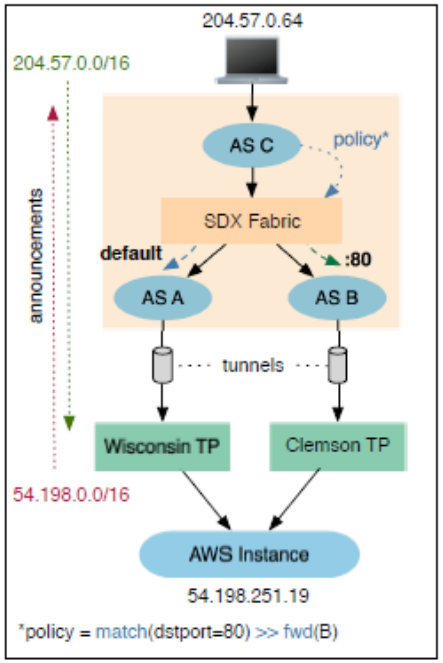

(a) *Application-Specific Peering.*

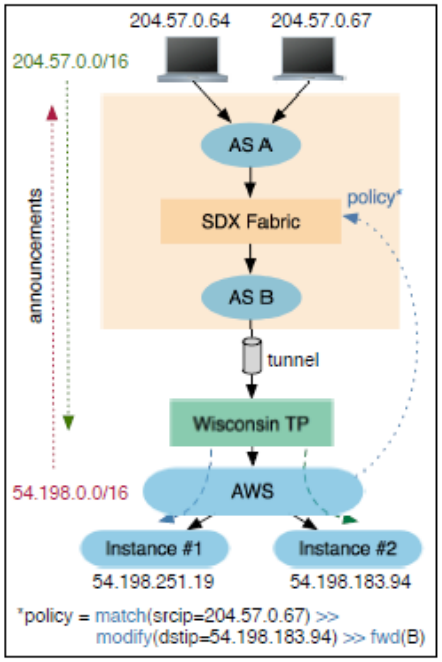

(b) *Wide-Area Load Balance.*

Figure 4: *Setup for deployment experiments.*

Serbian Open eXchange

Ako moderator već nije prozvao sledećeg izlagača, spreman sam da odgovorim i na par pitanja

Zoran Perović
iConsult@perke.NET
info@sox.rs