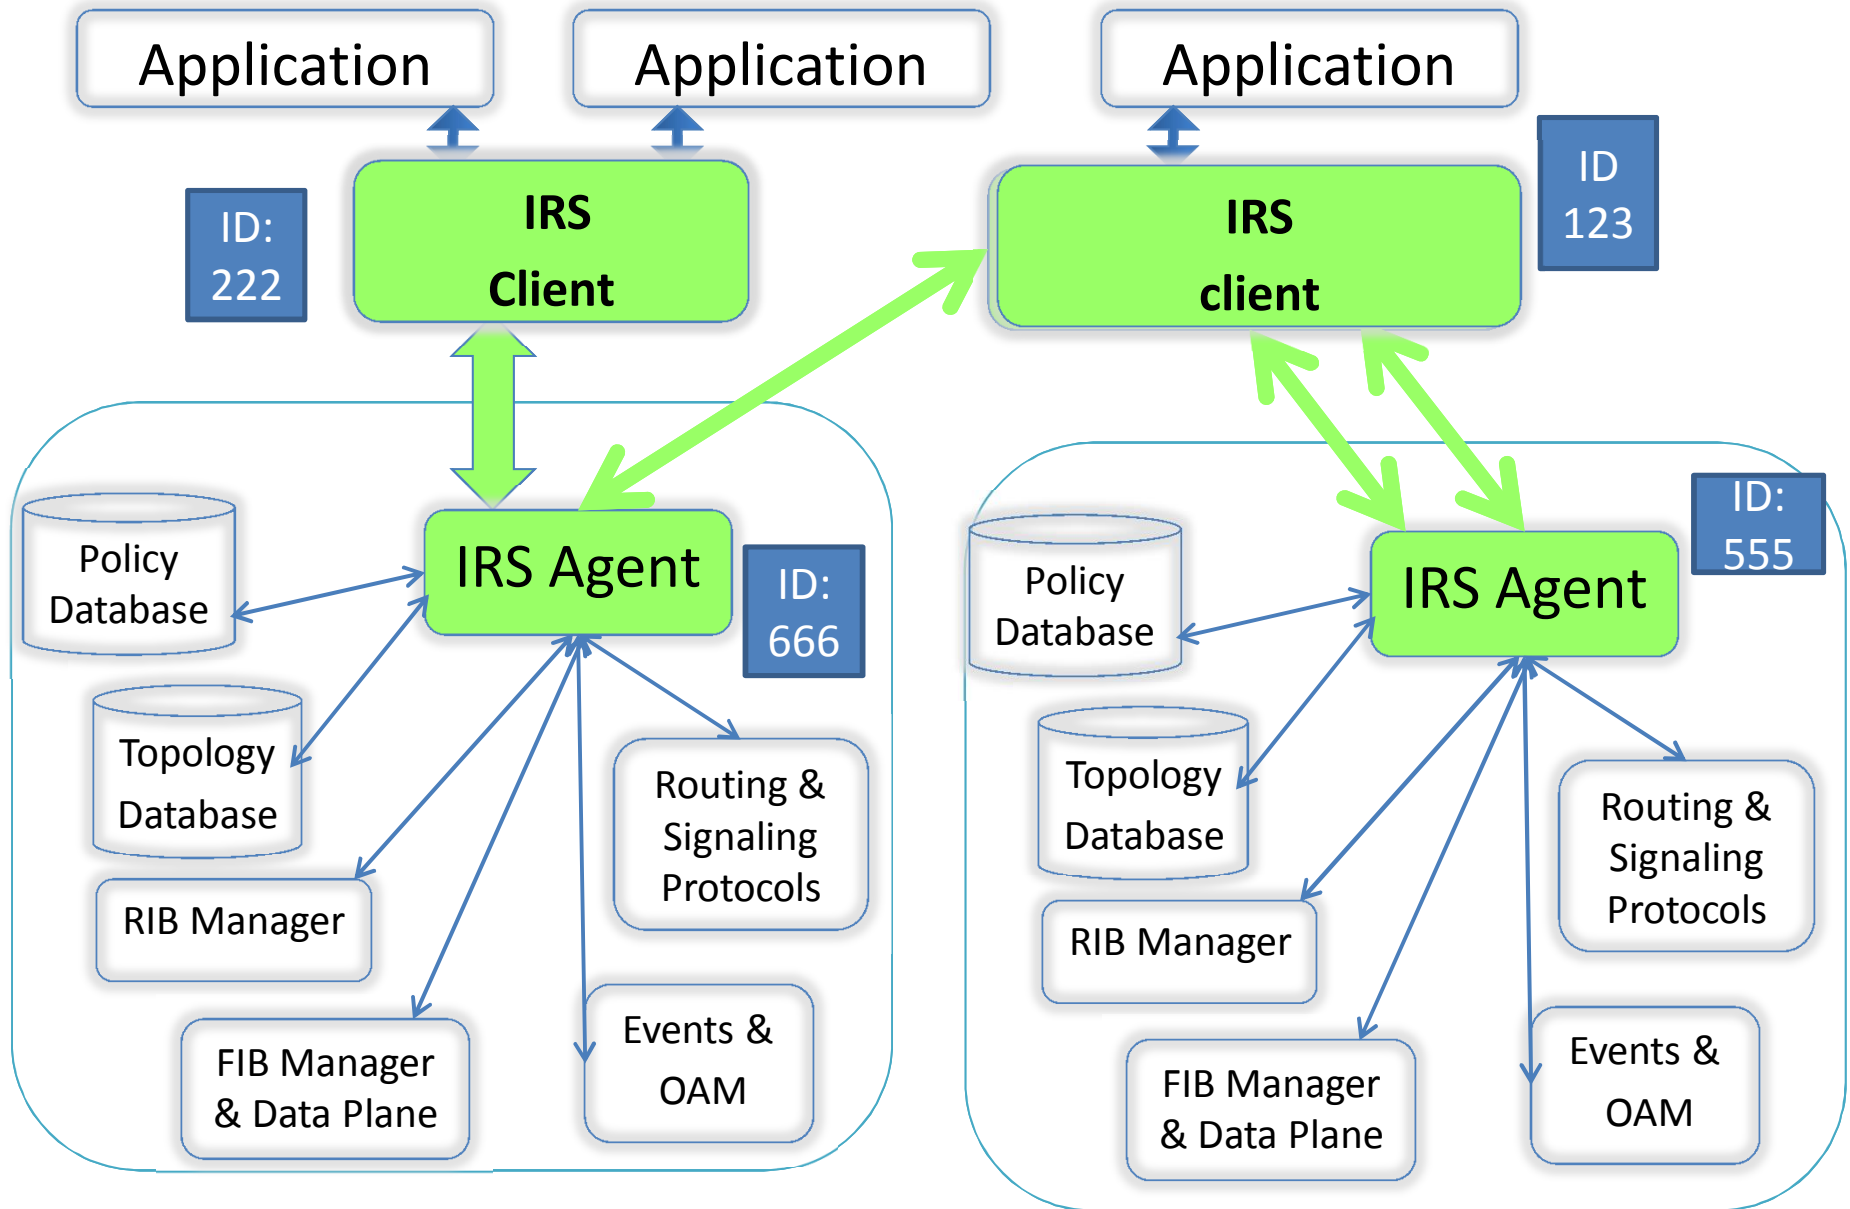


i2rs Policy – the basics

Sue Hares

I2rs interim April 22, 2013

IRS Policy Framework





Policy Framework

101

Policy Definitions

- “ Identity
 - . Not tied to a single channel
 - . One per commissioner
 - . One per agent
- “ Role
 - . Each commission has a security role
- “ Read Scope - what I can read
- “ Write Scope - what I can write
- “ Resources
 - . what agent can consume
 - . Example: # of installs, # of events, # operations
- “ Policy – explicit and implicit
 - . Explicit: what you configure
 - . Implicit: What’s implied in protocols or “doing the right thing” in configuration

Policy Actions

- “ Connectivity
 - . No need for active connection
- “ State
 - . Tied to Actions such as get this topology;
- “ Priority
 - . Commissioner gives 3 tasks:
 - “ pull routes,
 - “ status on interface 2,
 - “ turn on interface 3
 - . What’s the order
- “ Precedence Decisions
 1. Assume configured a route 192.165.2/24
 2. Multiple people use IR to move traffic for 192.165.2/24 short term
 - . Who gets to install
 - . what happens when they get done
 - . What happens on a reboot

Key Questions

- “ What i2rs clients can talk to an i2rs agent?
- “ What happens if there are two i2rs clients?
- “ What happens if local configuration overlaps with i2rs state?
- “ What happens when i2rs clients remove state
- “ What happens if node is rebooted?

A Simple Example will answer

Without I2RS clients


BGP picks Exit 2

Read Scope: Interface
Write scope: FIB

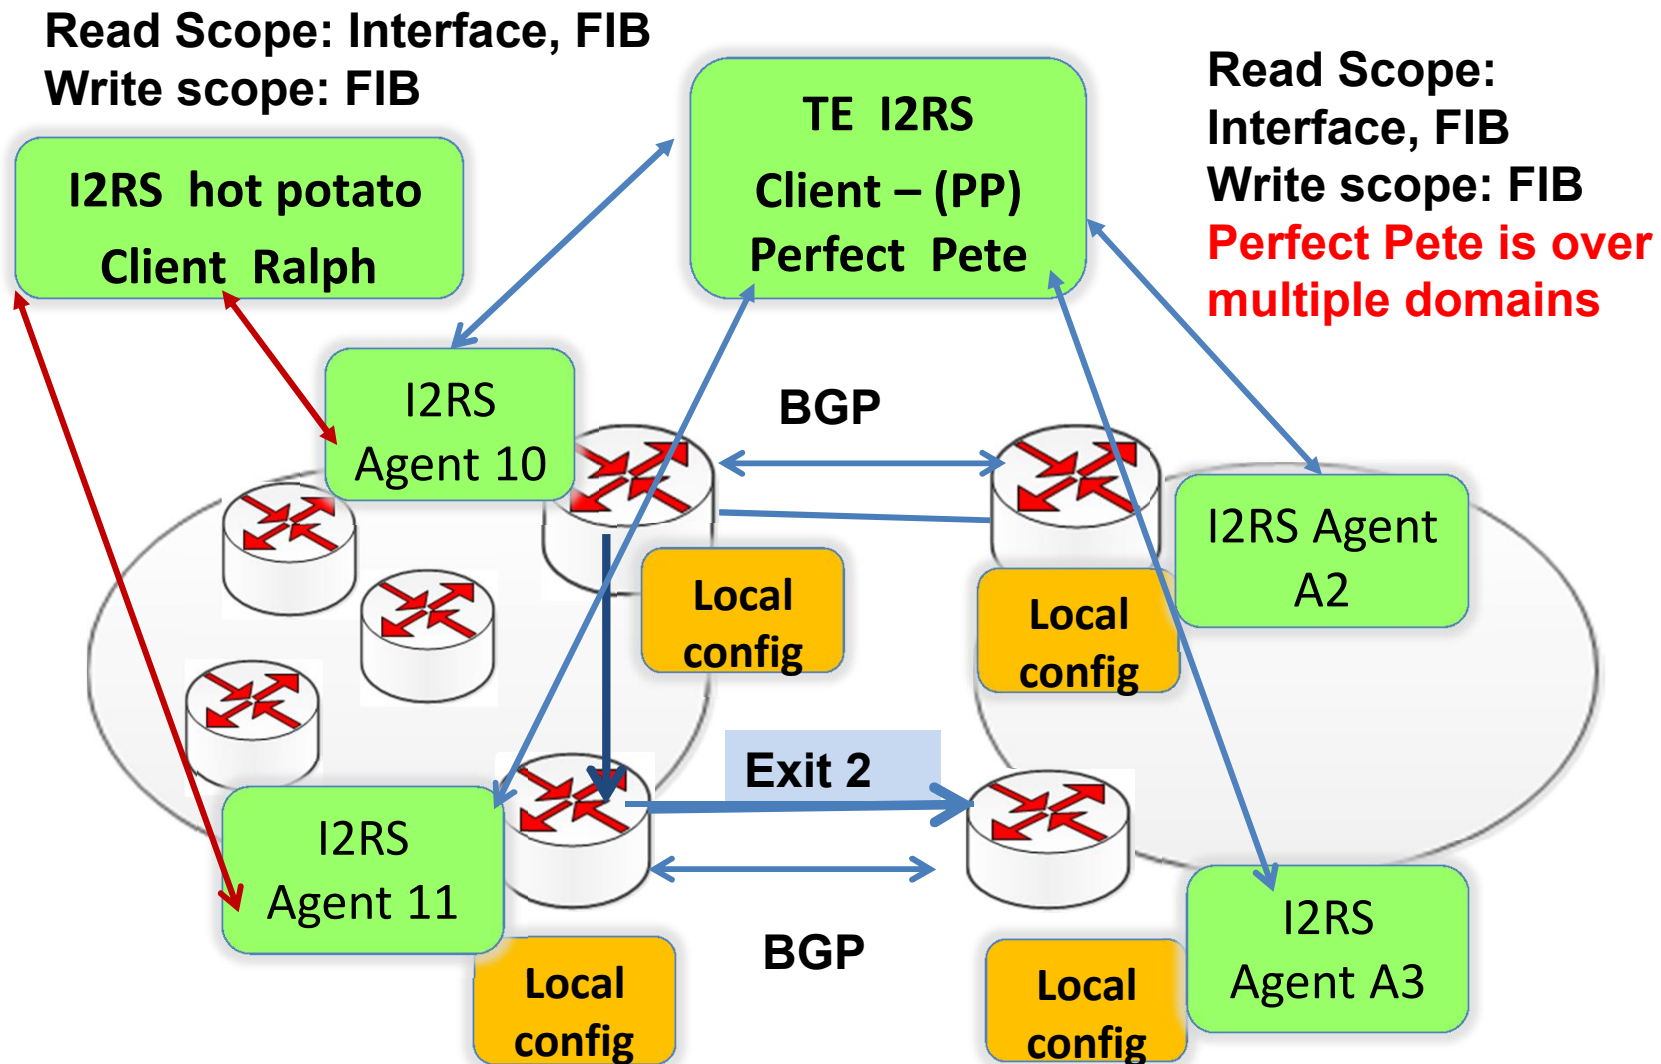
I2RS hot potato
Client

TE I2RS
Client

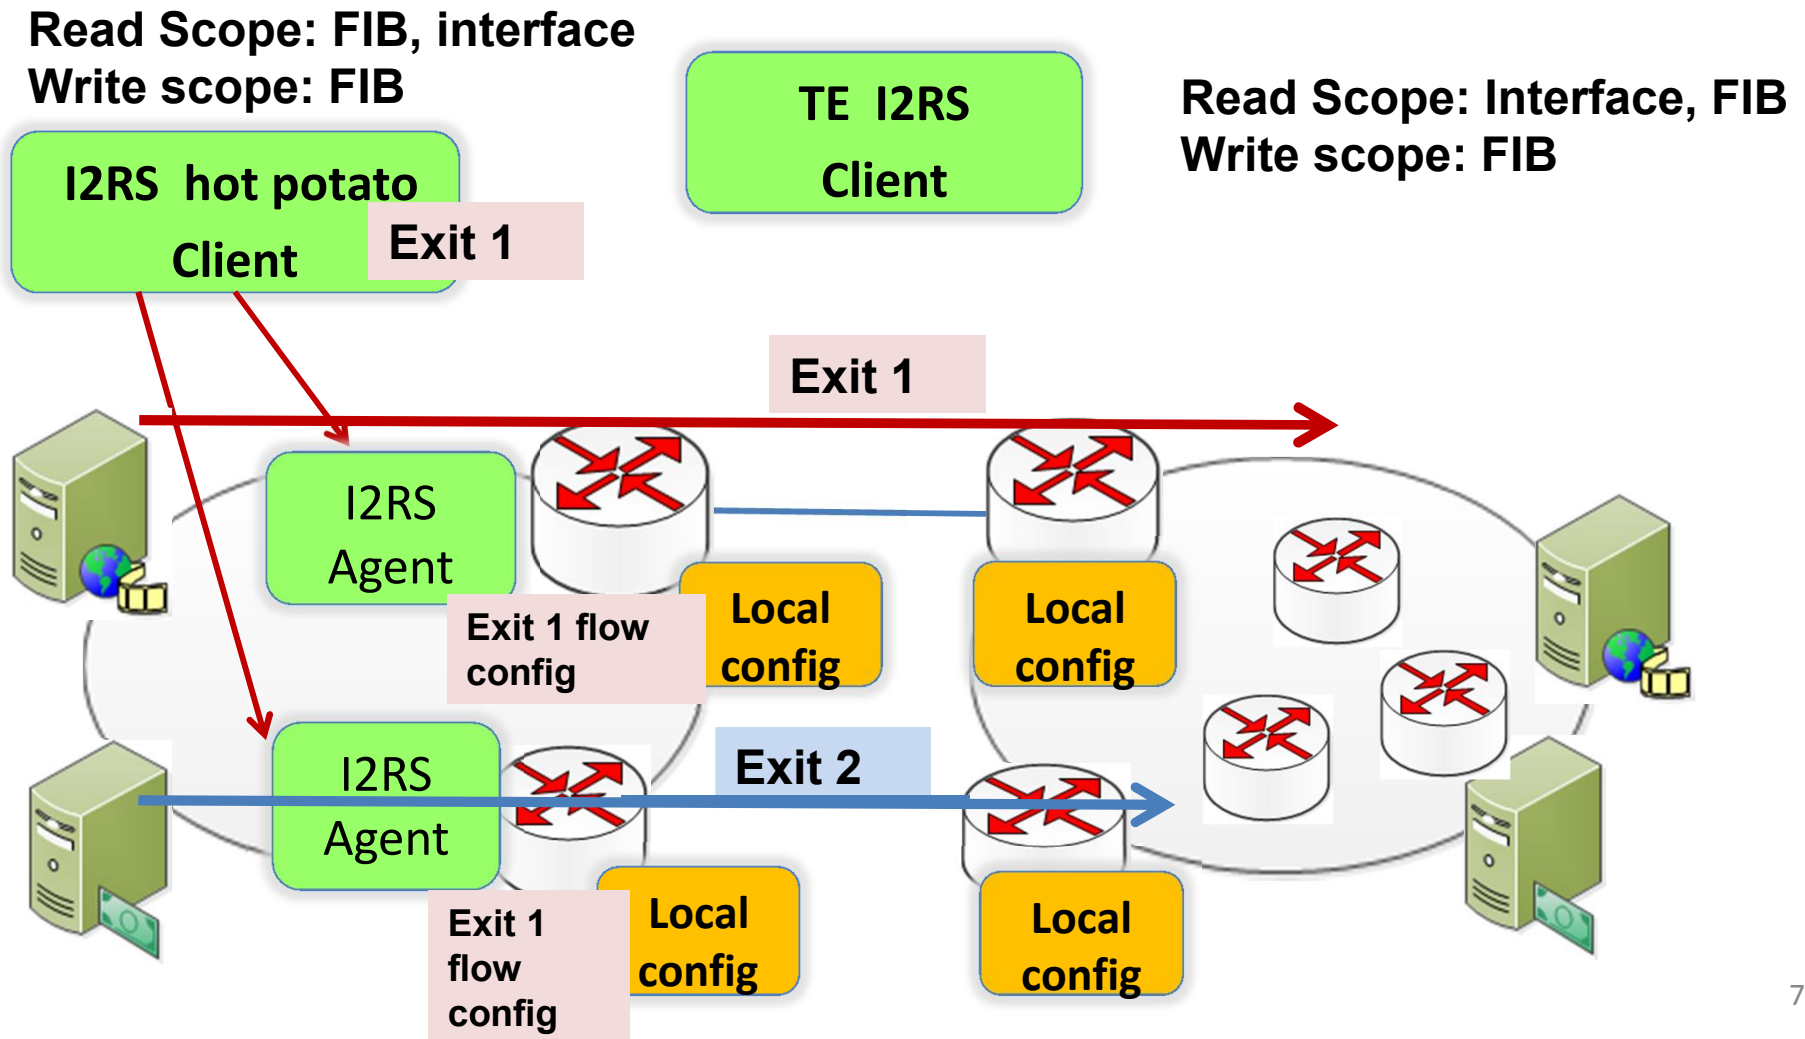
Read Scope: Interface
Write scope: FIB



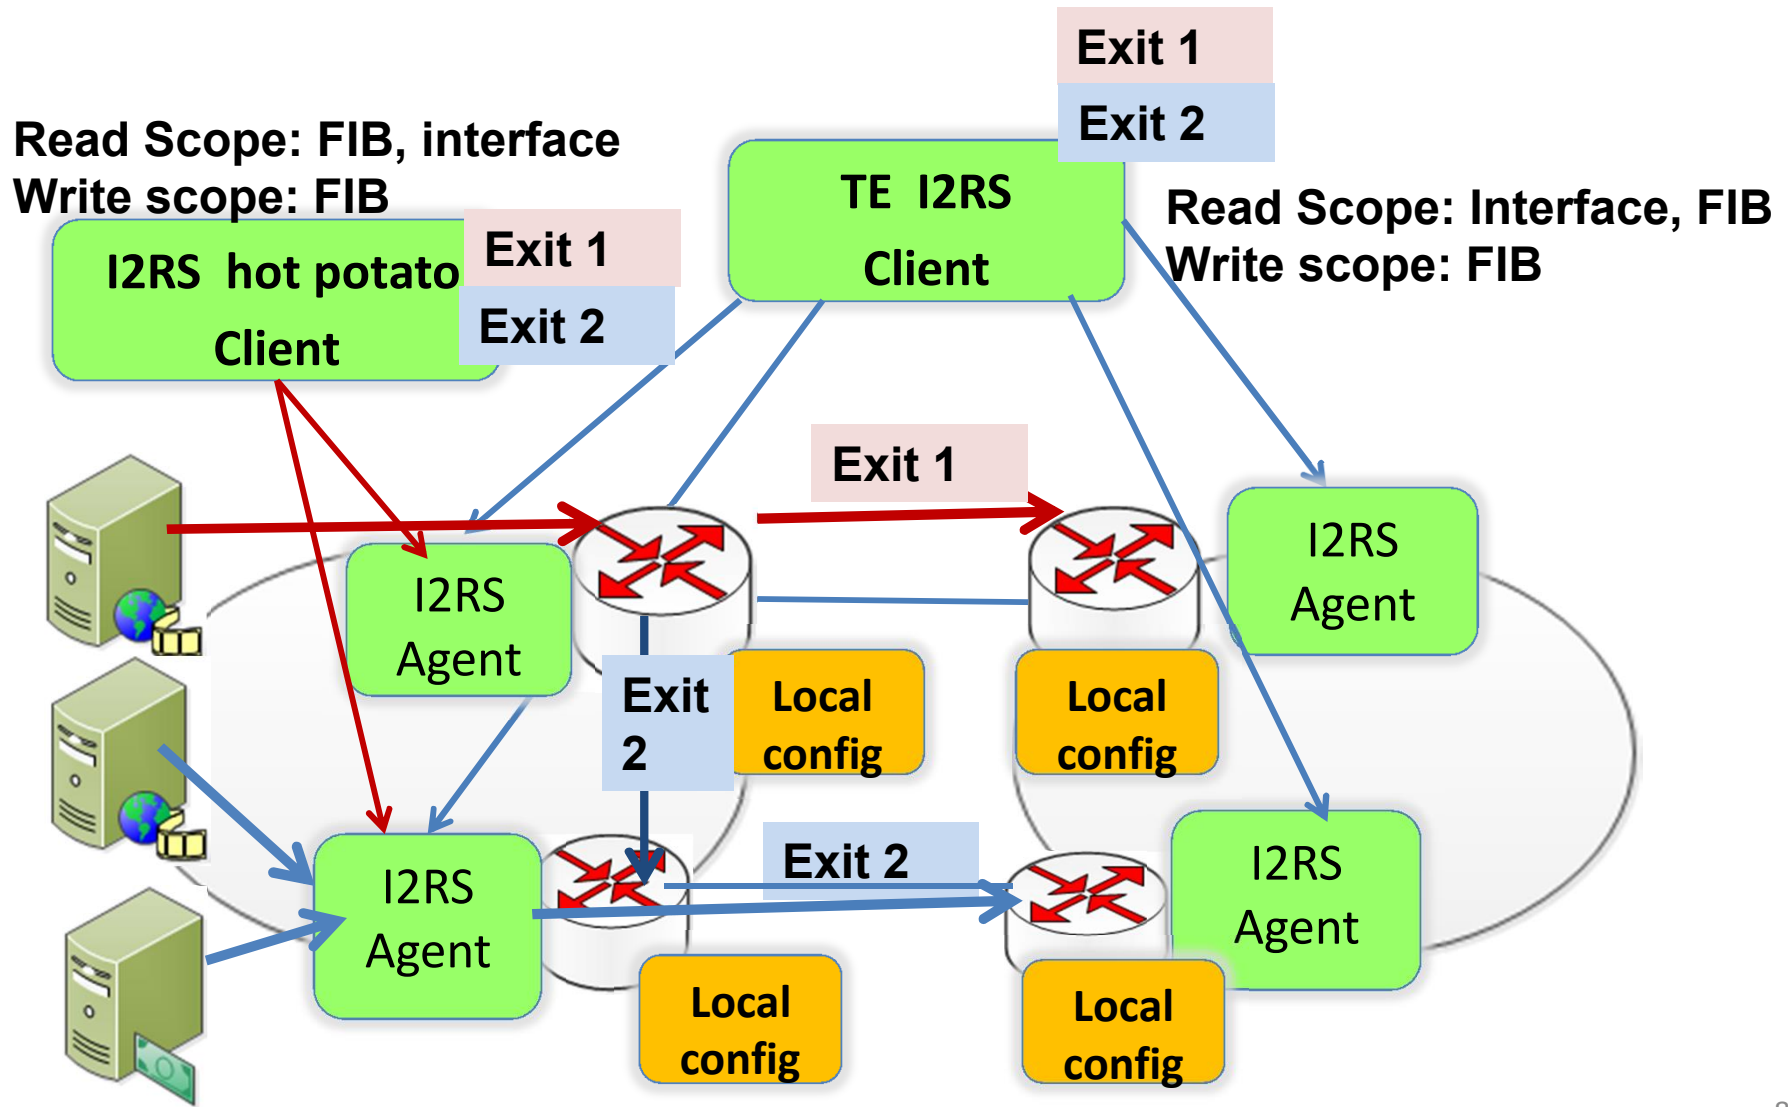
Identity Yourself



Hot potato selects Exit 1 and Exit 2 based on distance



TE Load Balance Exit 1 & Exit 2



Hot potato + TE client want

Read Scope: FIB, interface
Write scope: FIB

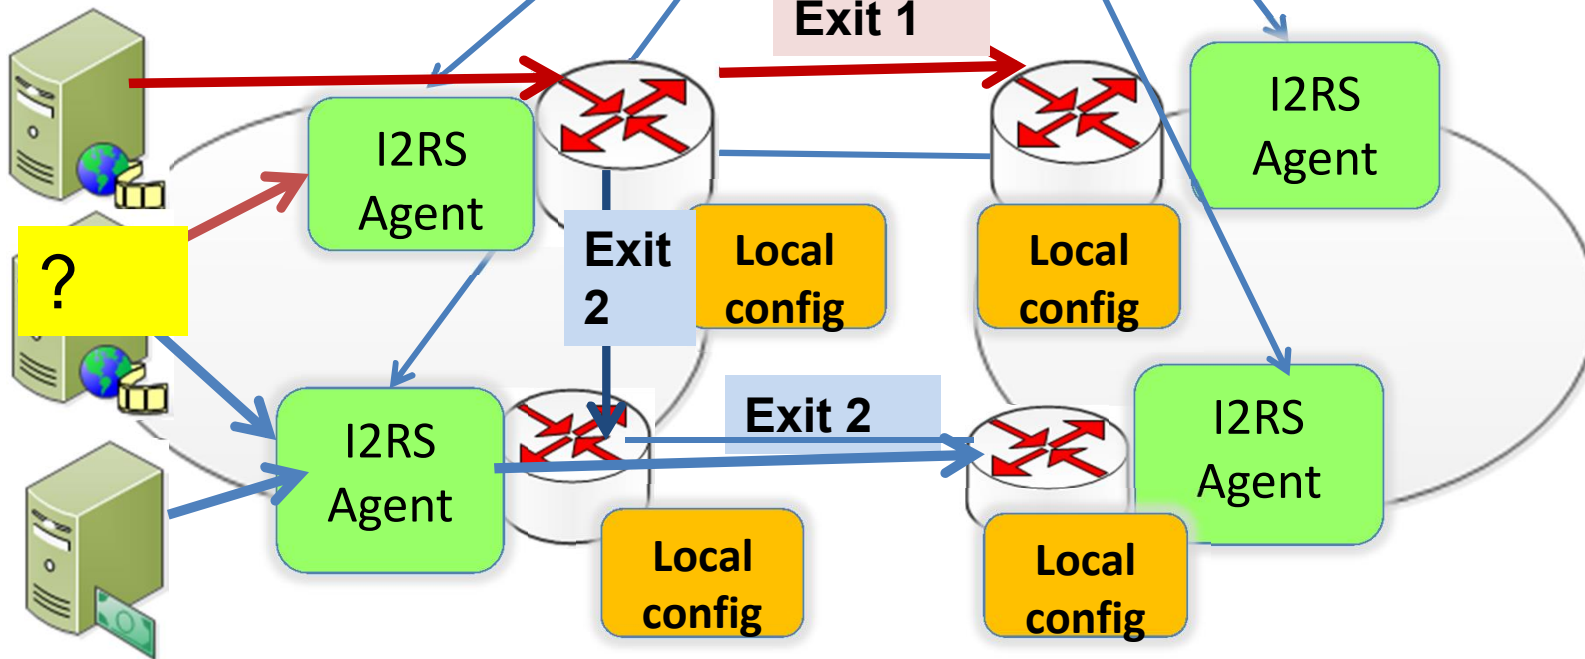
I2RS hot potato Client

Exit 1

Exit 2

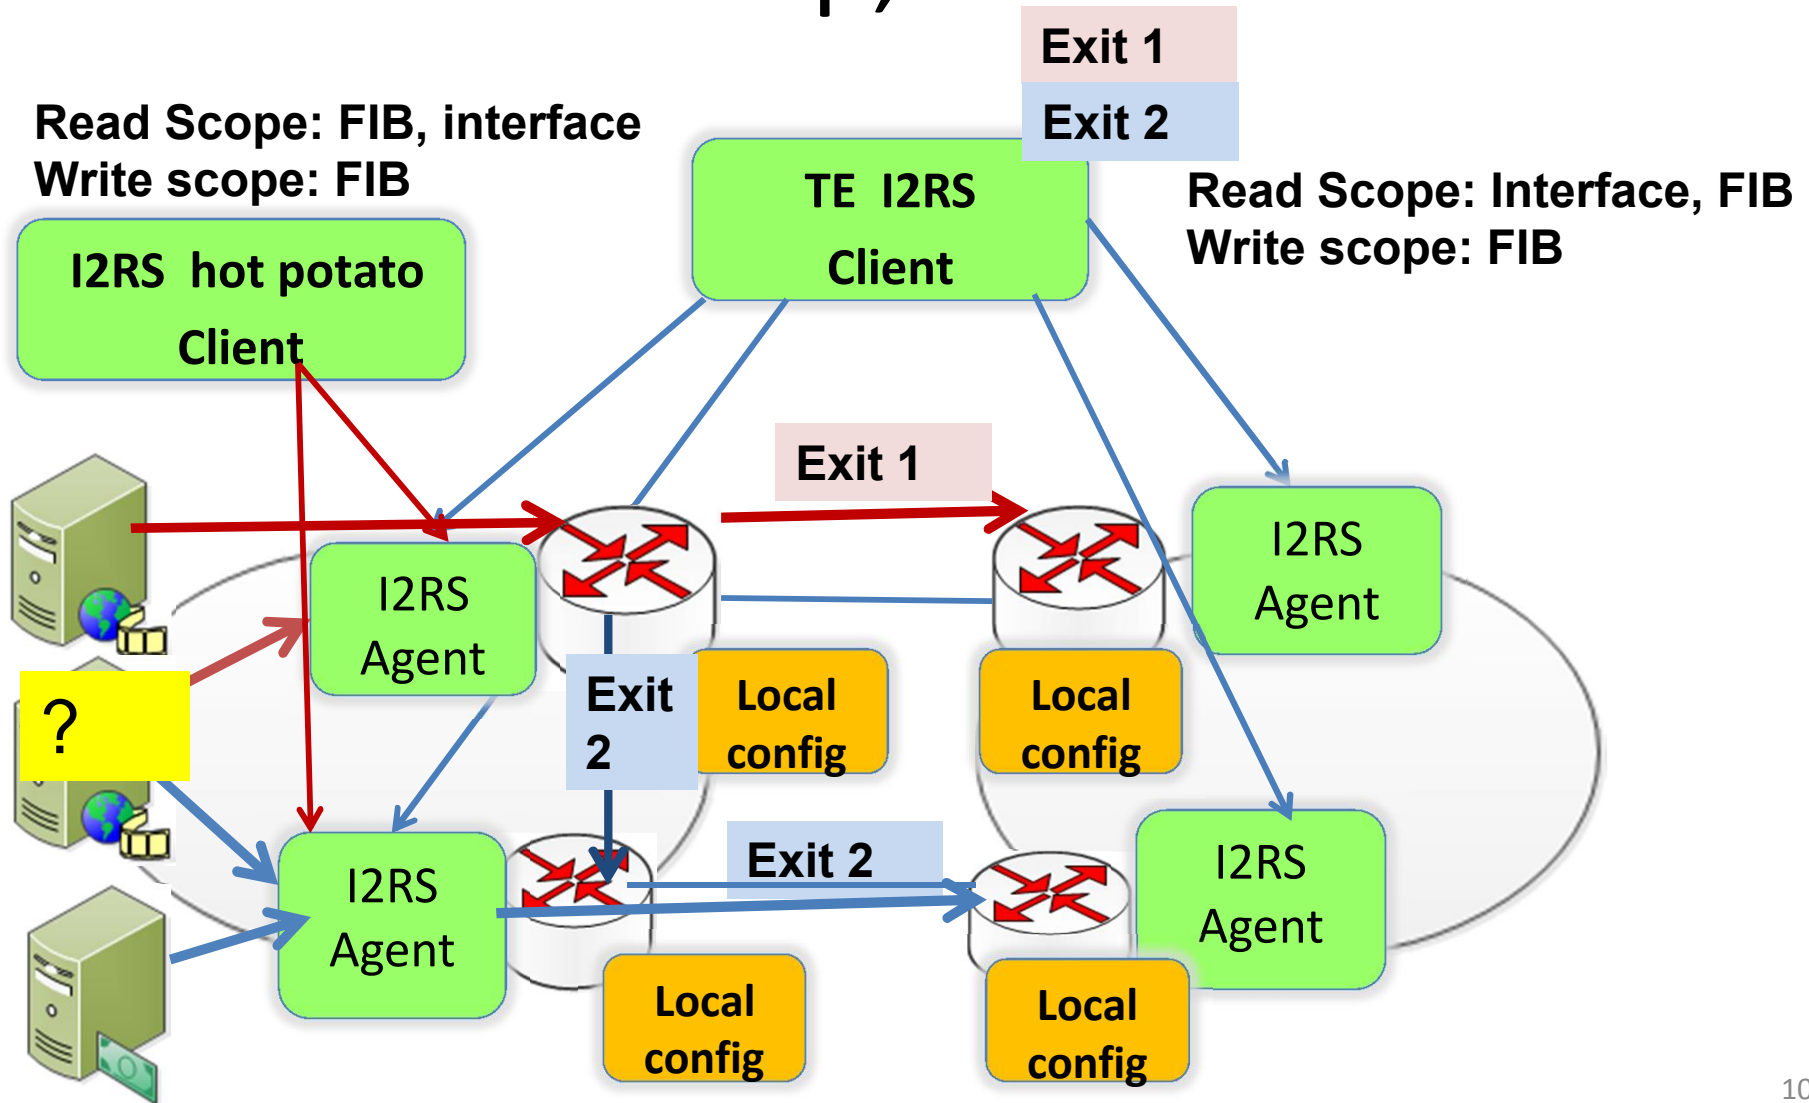
TE I2RS Client

Read Scope: Interface, FIB
Write scope: FIB

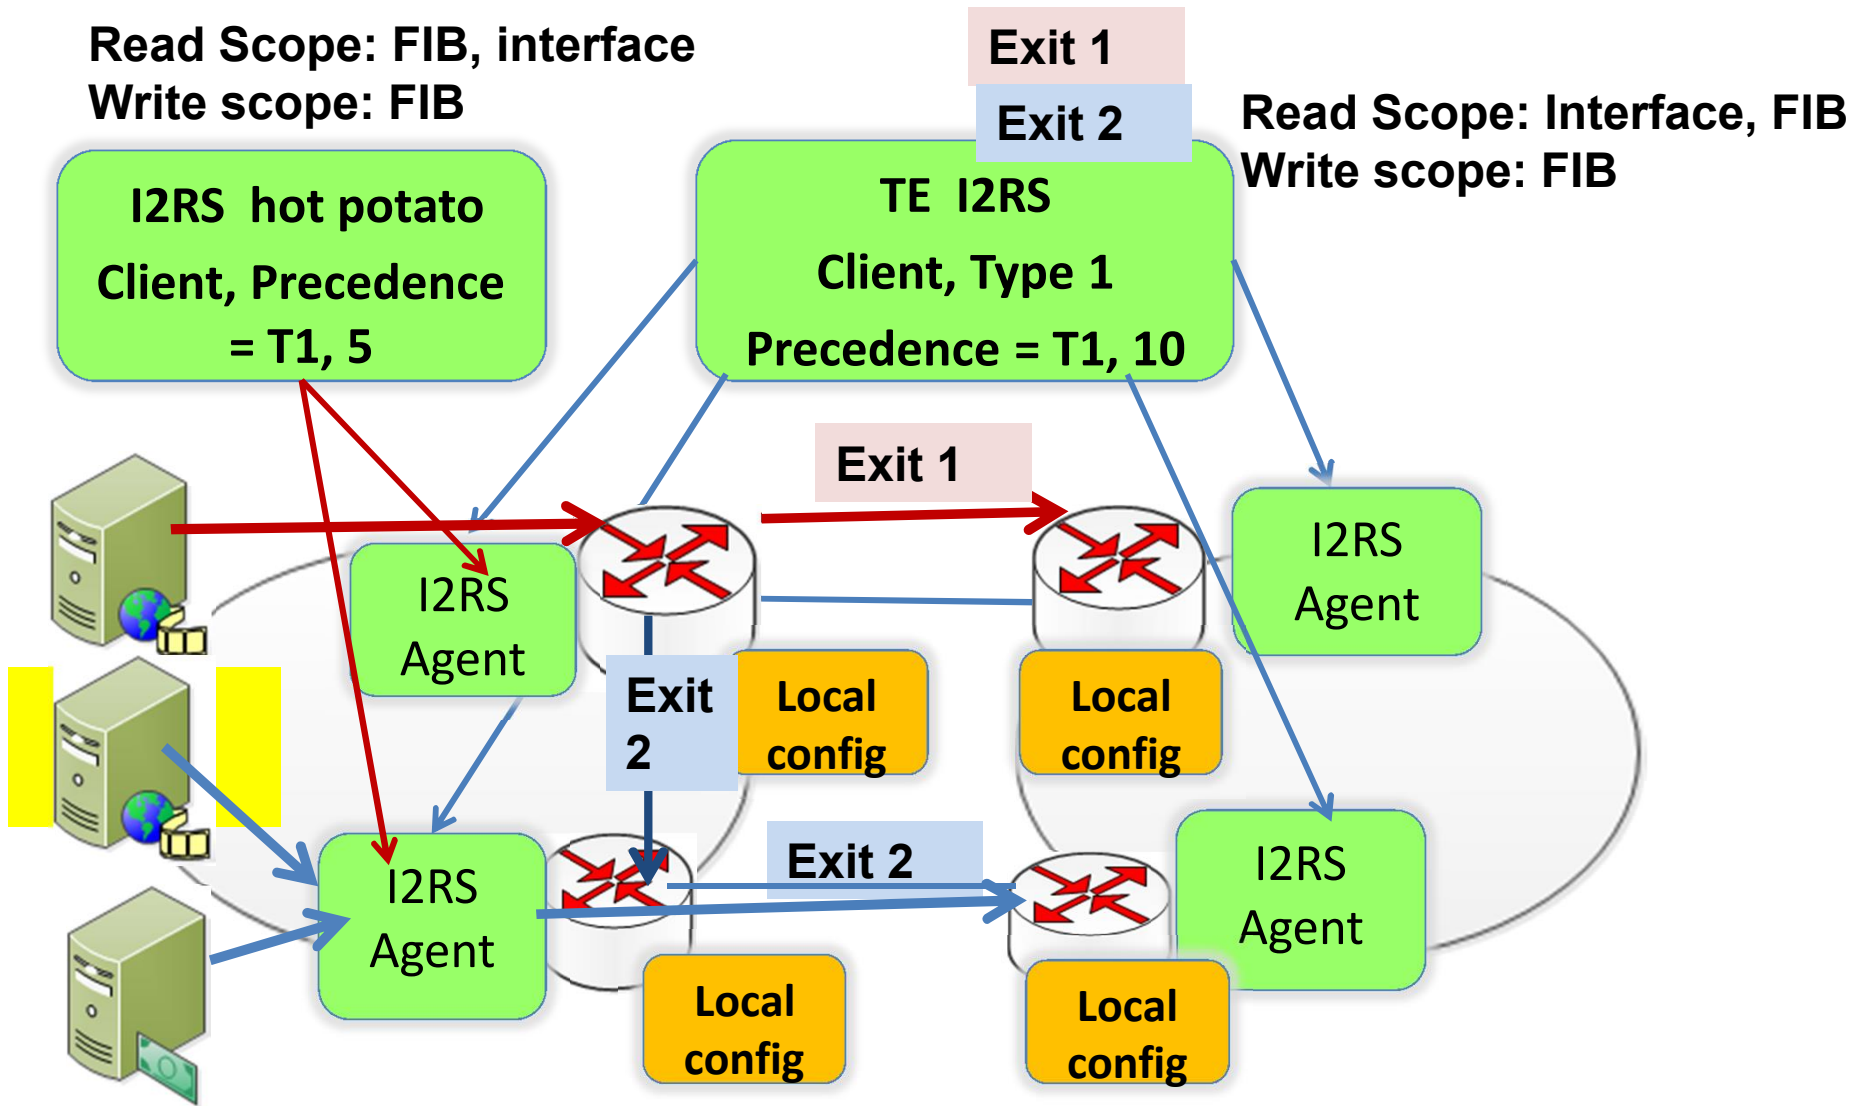


Hot potato + TE client

80% overlap, 20% different

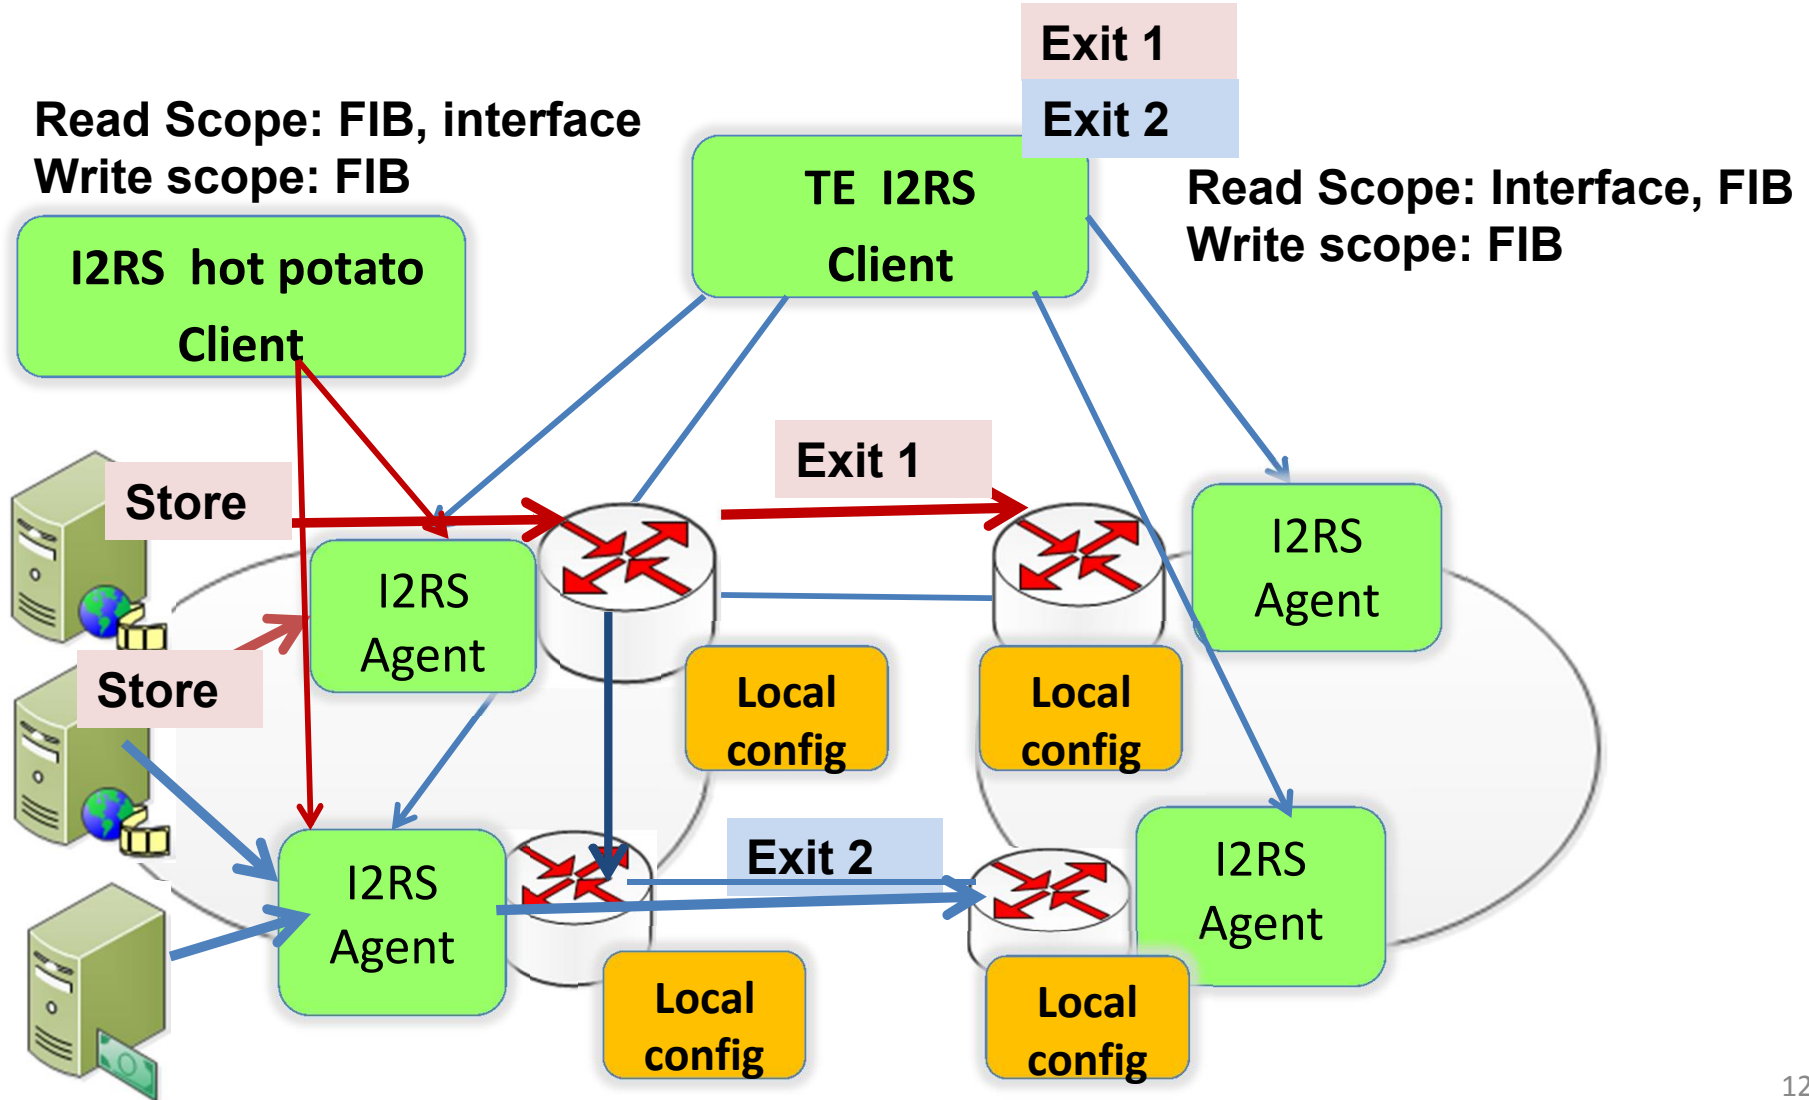


Precedence Resolves issue

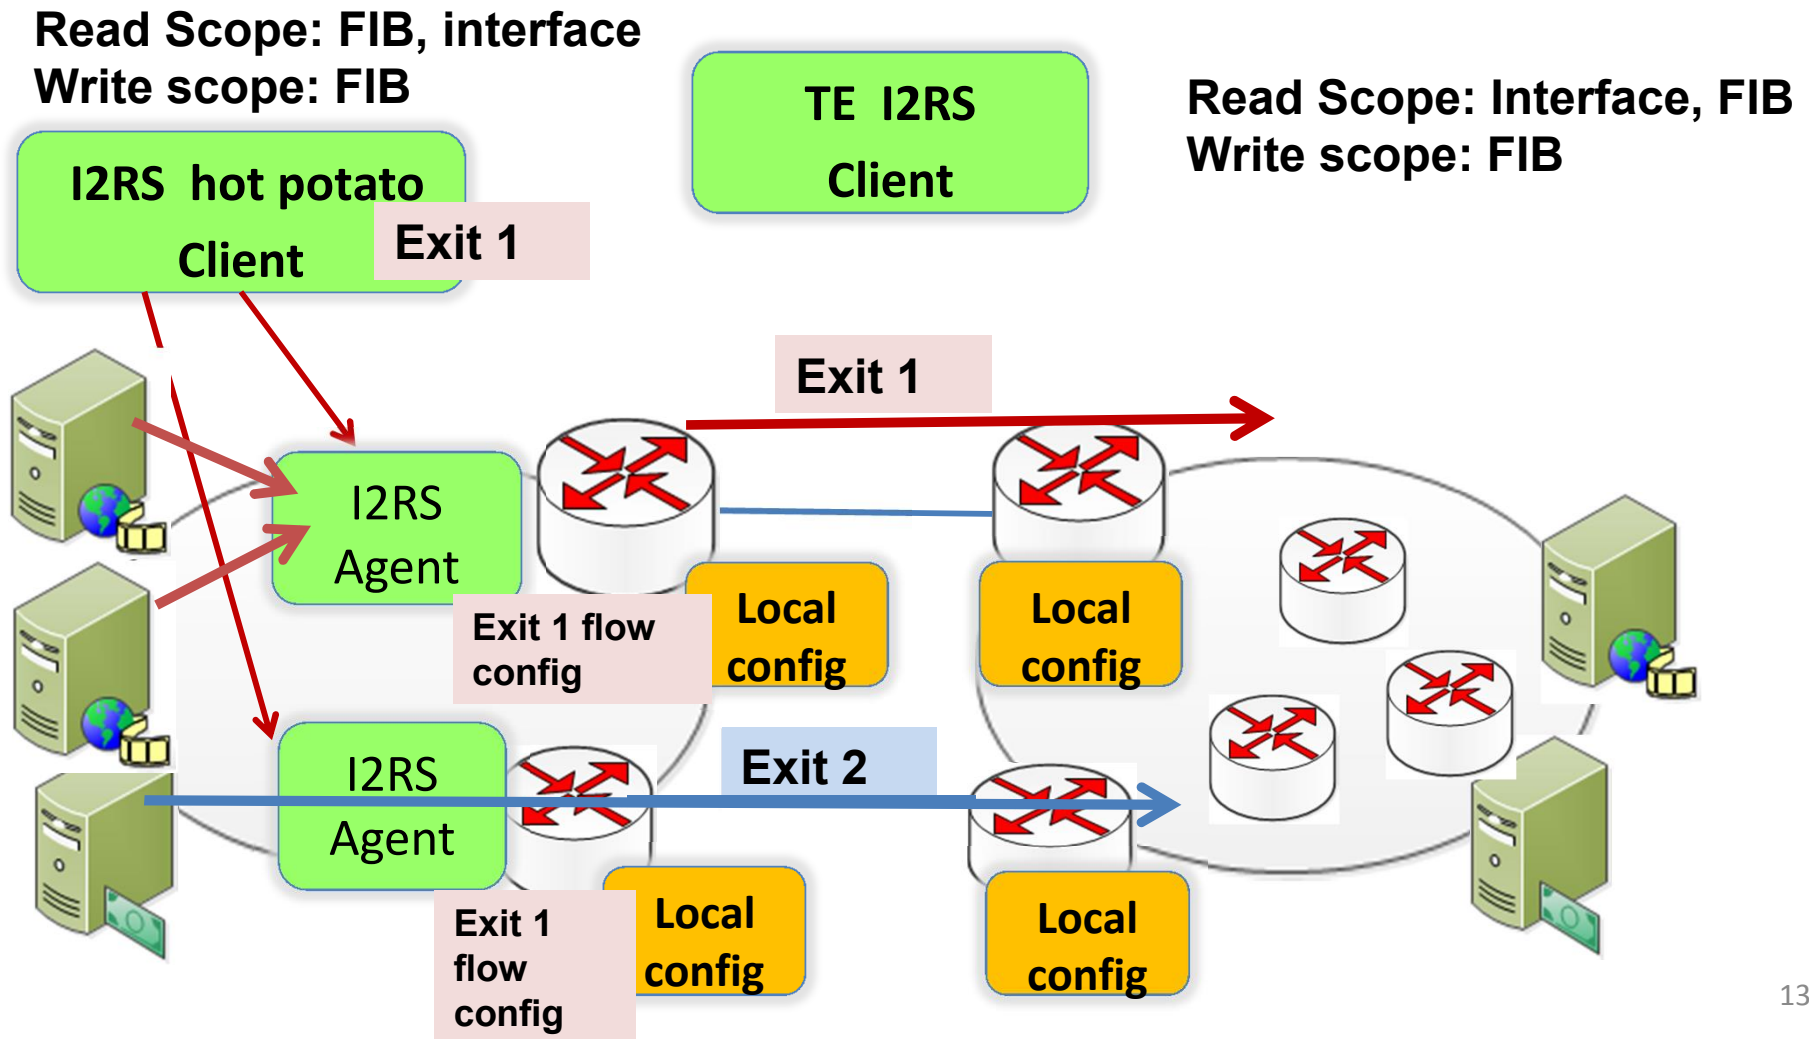


Hot potato + TE client

Store if not best



Return to Hot potato when TE clients stops



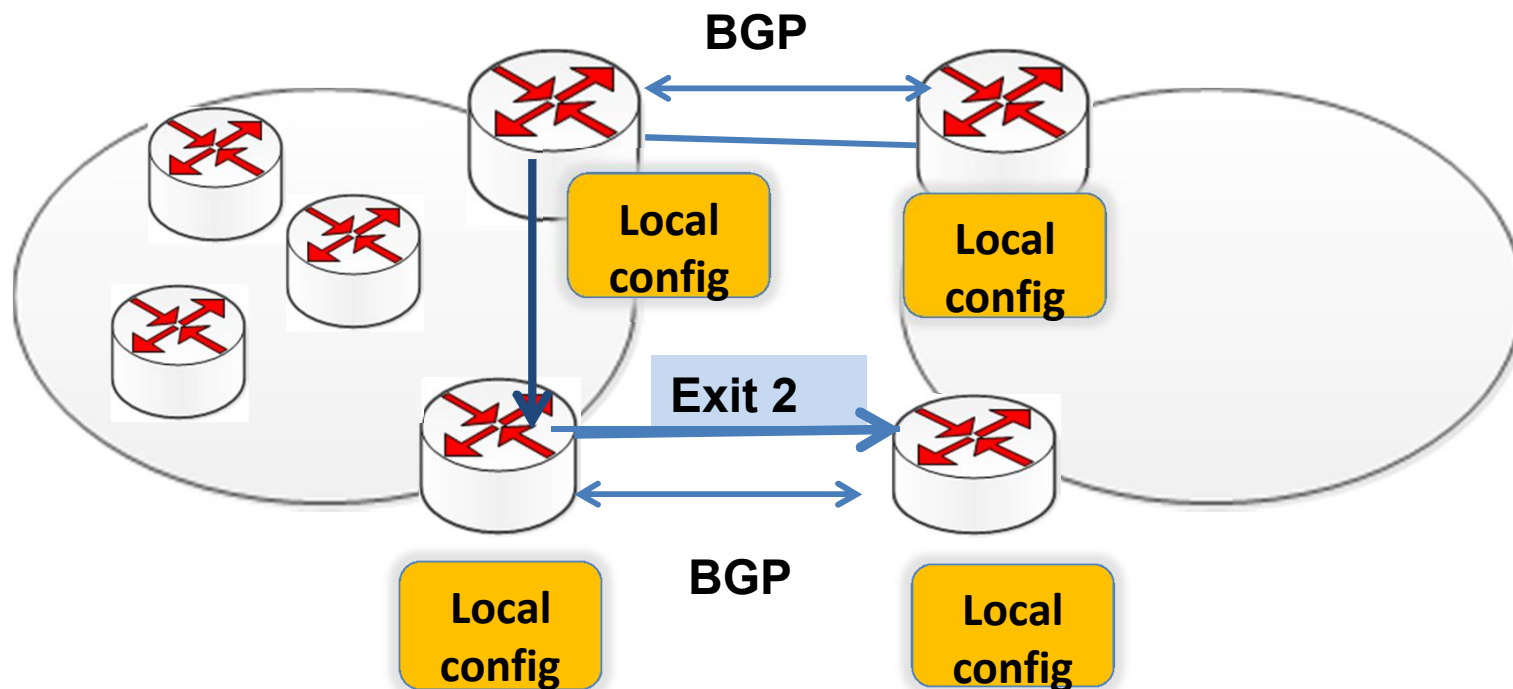
Upon Reboot – return to BGP

Read Scope: FIB, interface
Write scope: FIB

I2RS hot potato
Client

TE I2RS
Client

Read Scope: Interface, FIB
Write scope: FIB



Key Questions Answered (1)

“ What i2rs clients can talk to an i2rs agent?

- . Identity and security
- . NfV causes Perfect Pete to exist (see DT talk)

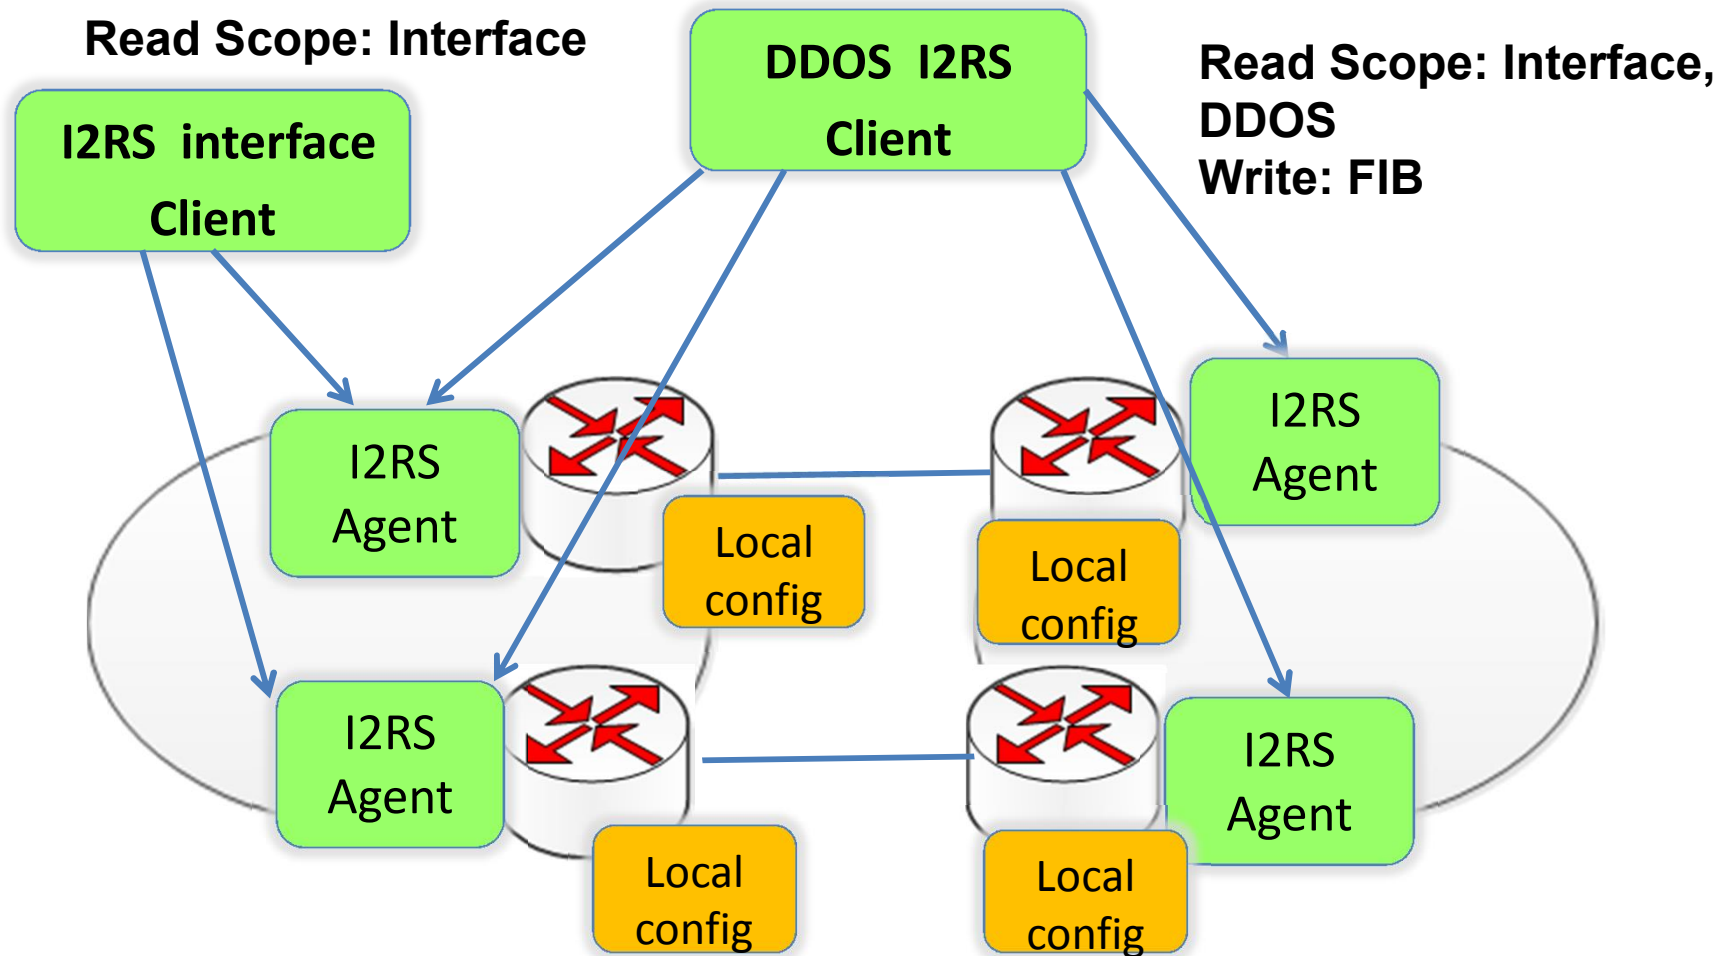
“ What happens if there are two i2rs clients?

- . If both try to change – precedence
- . (?) Should precedence be all per client or overlap?

Key Questions answered (2)

- “ What happens if local configuration overlaps with i2rs state?
 - . BGP lost in above examples to more specific hot-potato, TE load balancing, and hot-potato + TE load balancing
- “ What happens when i2rs clients remove state
 - . Depends on store-if-not-best flag
- “ What happens if reboot
 - . Goes away or persists – it’s the designers choice

Is this Multi-headed ?



Types of Commits for i2rs config

- “ Reserve the TE for this time – so I can distribute my video to Cache boxes at edge
- “ Start and Stop time

- “ Simple Local Config/i2rs config gotcha’s
 - . Policy is both implicit and Explicit
 - . Implicit: BGP keepalive is 30 sec (by default)
 - . Explicit: Configure this BGP peer

How to match this to Andy's talk

- “ This is What, Andy's is “how”
- “ Dynamic Configuration – that's how the i2rs writes the policy for rib forwarding
 - . Should we do conditional store?
Draft-kwatsen-conditional-enablement-00
 - . Should we have confliglets: static, conditional, triggered?
 - . Precedence based on access groups
Draft-rfernando-i2rs-yang-mods-00

Will Simply Policy for other use cases help?

If So.. I'll write-up over lunch and send out

Other use cases

Choose 1 or 2

Q&A

draft-white-i2rs-use-case-00

- “ Fine-grained tuning of traffic flow(s) in an IP network
- “ **Optimized Exit Control (What you saw)**
 - . Current dynamic routing protocols (BGP) do not provide the granularity to fine-tune exit paths in a network
- “ **Reacting to DDoS attacks**
 - . Redirect some traffic through “traffic scrubbing” points in the network.
- “ **Dynamically optimize traffic flows in a hub & spoke network**
 - . Instead of forwarding all traffic spoke-to-spoke traffic through hub-site; the hub site could dynamically program a spoke site to directly forward traffic to other spoke sites.
- “ **Inside DataCenter Routing**
 - . Quickly modify routing based on topology changes and shifts in traffic patterns
- “ **Between DataCenter Routing**
 - . “Bandwidth on Demand” across the WAN to move or replicate resources from one DC to a second DC.

draft-amante-i2rs-topology-use-cases-00

- “ Framework for Topology, Policy & Orchestration
‘Manager’ Functions
- “ Use Cases
 - . Capacity Planning & **Traffic Engineering**
 - “ Operate at different time scales, but critical requirement to incorporate information from multiple data sources: statistics & inventory data warehouses
 - . **VPN Services Provisioning**
 - . **Rapid IP & ASN Renumbering**
 - . Path Computation Element (PCE)
 - . ALTO Server

draft-keyupdate-i2rs-bgp-usecases-01

- “ Mass BGP Protocol & Policy Configuration Changes; Analysis and Troubleshooting of BGP Routing State across an entire network
- “ Side note: BGP configuration is overwhelming majority of configuration on routers and modified the most frequently, by Service Providers
- “ BGP Protocol Configuration
 - . Dynamically change AFI/SAFI; **ASN migration scenarios; RT (RD) changes**, etc.
- “ BGP Policy Configuration: Route filtering, **Route summarization**
- “ Internal BGP Error Handling (?)
- “ BGP Route Manipulation
 - . Customize Best Path Selection
 - . **Flowspec (react to DDoS attacks) – similar to draft-white-irs-use-case-00**
 - . **Optimized Exit Control, a.k.a.: TE – similar to draft-white-irs-use-case-00**
 - . **Change RT values on RR’s – similar to draft-white-irs-use-case-00**
- “ BGP Events
 - . **Notify applications when changes occur to “important routes”**
 - . Troubleshooting Filtered BGP Routes
 - . BGP Protocol Statistics – monitor & change ‘max-prefix limit’

draft-medved-irs-topology-requirements-00

- “ Reqmt’s for ‘Topology Manager’ function of draft-amante-irs-topology-use-cases
- “ Topology Manager (TM) constructs virtualized views of global network topology for consumption by Clients
- “ General:
 - . Define standards-based data models with common vocabulary to describe various network components
- “ Data Model:
 - . Layer-2 & higher Data Model Reqmt’s
 - . MUST capture Visible & ‘Invisible’ Network Components
 - . Hierarchical representation, composition and summarization of network components into real or virtual/abstract depictions of network topologies
- “ Northbound (Client) API
 - . Efficient, flow-control-capable protocol for large data transfers between TM & Clients
 - . **MUST support publish/subscribe capability**
 - . MUST support ‘non-Routers’ as Clients, (up to now Clients needed to run a dynamic Routing Protocol to learn of network topology or events).
- “ Southbound (Network & Device) API – better reqmt’s in draft-rfernando-irs-framework-requirement-00 ... 😊

draft-rfernando-i2rs-framework-requirement-00

- “ IRS Framework Terminology and Requirements
 - . In-depth requirements for protocol and (service) data models used between Clients & Servers in IRS
- “ One, p2p transport connection between a client and a server
- “ In order, reliable data delivery in both directions
 - . Critical when client is adding, changing or deleting state on a server.
- “ **Publish-Subscribe for Asynchronous Notifications of Events** that occur on the Server
- “ Security Requirements: server needs to validate Identify of client, before allowing client read-only or read-write access to server state.
- “ Application Programmability Reqmt’s:
 - . Apps should focus on functionality, they should not have to focus on mechanisms wrt communicating with Servers.
 - . Apps should be re-usable across different environments
 - . App “templates” (design patterns?) should be available in a common repository for re-use by other app developers ...