DEMYSTIFYING IPV6 OVER MPLS

TACKLING THE CHALLENGE OF CONNECTING IPV6 ISLANDS

KAM AGAHIAN CCIE#25341 (SP – R&S)

SENIOR MANGER OF NETWORK ENGINEERING, QUALCOMM

HOW THIS SESSION IS STRUCTURED?

- We will cover
  - Basics of IPv6 and MPLS (101 with some oversimplification)
  - Need to IPv6 in modern networks
  - Current migration solutions
  - **Outside the box thinking**
  - Real world scenario

THE "OUT OF THE BOX" THING?

- **Our big issue as an industry; a bit off topic**
- Let's look at some common themes:

- The curios case of vulnerability scanning
  - New CISO or Security manager: Security strategy vs vulnerability scanning
  - ICS (Industrial Control Systems), SCADA/PLC scanning (Future NANOG?)
    - IoT scanning and even more scanning!

- The curios case of IPv6 dual stacking and tunneling – Show of hands!

THIS SESSION…

- Adds yet another tool to your toolbox
  - Maybe an unusual one

LET'S FORGET ABOUT THE IP WORLD

- How many of you have travelled to more than 40 countries?
- How many of you have ever had issues overseas with non-local ID's?
  - Driving on a non-local ID….

LET'S FORGET ABOUT THE IP WORLD

- The imaginary trip:
  - Your document is in English
  - No one understands English even a bit
  - You are passing through

LET'S FORGET ABOUT THE IP WORLD

- Can you think of a (semi-) common language?

  - ❑ Perhaps numbers? attached to your actual ID.

HOW ABOUT THE IP WORLD?

- Applying the same concept:
  - The layer 2 and layer 3 headers are sometimes non-sense to the middle network
  - Can we still use numbers? And forget about the headers…

WHAT TO CALL IT?

- Let's pick a name for it…
- Multi Protocol Label Switching (MPLS)
  - Provides that common language for passing through foreign territories
  - "Staples" numbers (aka labels) to each packet passing through
  - The labels provide a means understandable throughout the new territory
  - Even you can change your path throughout the foreign territory without speaking their language

LET'S TAKE A CLOSER LOOK : HTTP

- Header of a HTTP packet in IPv6 going from 2006::6 to 2001::200

```
> Frame 65: 82 bytes on wire (656 bits), 82 bytes captured (656 bits) on interface 0
> Ethernet II, Src: ca:03:0b:f8:00:08 (ca:03:0b:f8:00:08), Dst: ca:02:00:08:00:1c (ca:02:00:08:00:1c)
> MultiProtocol Label Switching Header, Label: 17, Exp: 6, S: 0, TTL: 253
> MultiProtocol Label Switching Header, Label: 25, Exp: 6, S: 1, TTL: 255
> Internet Protocol Version 6, Src: 2006::6, Dst: 2001::200
> Transmission Control Protocol, Src Port: 80, Dst Port: 60121, Seq: 1, Ack: 11, Len: 0
```

ENOUGH FOR NOW…

- But, "I owe you", a lot more on the topic.

NOW BACK TO NANOG 76

- **Where We Are At:** Companies already have an IPv4 network including LAN and WAN
- **The Challenge**: Now they are facing a growing need to connect IPv6 islands
- **The Option:** MPLS can help you without adding IPv6 to your IPv4 stable network (at least for years)

HANG ON…FIRST A BIT ON IPV6

- IPv6 Primer – Who doesn't know?!
  - 128-bit address space instead of 32-bit in IPv6
  - 16-17 billion times larger than the IPv4 address space
  - IPv6 Internet is fully functional; ~69,000 IPv6 address globally
  - Most networking features are available across multiple vendors
    - BGP, ISIS, OSPF, RIP, EIGRP, QoS, Multicast (BSR, SSM) and more

IP VERSION 6 PRIMER

- So why don't we all do IPv6 everywhere and call it a day?
  - Availability and stability concerns….bugs?
  - Security concerns
  - Fairly steep learning curve when rolled out globally
  - Systems and applications engineering are (still) lagging behind networking
    - AAA - Just one example
    - Hard coded applications and the list goes on
  - **Complicated migration strategies (see next slide)**

IPV6 MIGRATION STRATEGIES

- Dual stack
  - Run both IPv4 and IPv6 on the same device
  - Still very invasive: stability, learning and troubleshooting
  - Umm, it's not really a "migration" strategy

- Tunneling
  - Many techniques have come and gone
  - Vendor support issues
  - Our experience with vendors

WHERE DO WE SEE THE IPV6 DEMAND?

- A bit on Qualcomm
  - Founded in 1985
  - Leader in 5G, IoT and various SoC's
  - 40+ major acquisitions, 10s of locations worldwide with up to 100Gbps WAN connections
  - Security through segmentation
- Hundreds of labs supporting thousands of hardware design and Q/A engineers
- IPv6 is an integral part
- Not much on the corporate-side of the house YET

ALL IN ALL…

- The demand is real…
  - 1500+ open positions on LinkedIn…

DOWN TO THE BUSINESS

IPv4 packet
Typical Layer 2
Typical Layer 3 (IP v4 src/dst)
ICMP

```
> Frame 14: 114 bytes on wire (912 bits), 114 bytes captured (912 bits) on interface 0
> Ethernet II, Src: ca:01:2f:14:00:08 (ca:01:2f:14:00:08), Dst: ca:00:10:94:00:08 (ca:00:10:94:00:08)
> Internet Protocol Version 4, Src: 6.6.6.6, Dst: 12.0.0.1
> Internet Control Message Protocol
```

THE HAPPY IPV4 WORLD

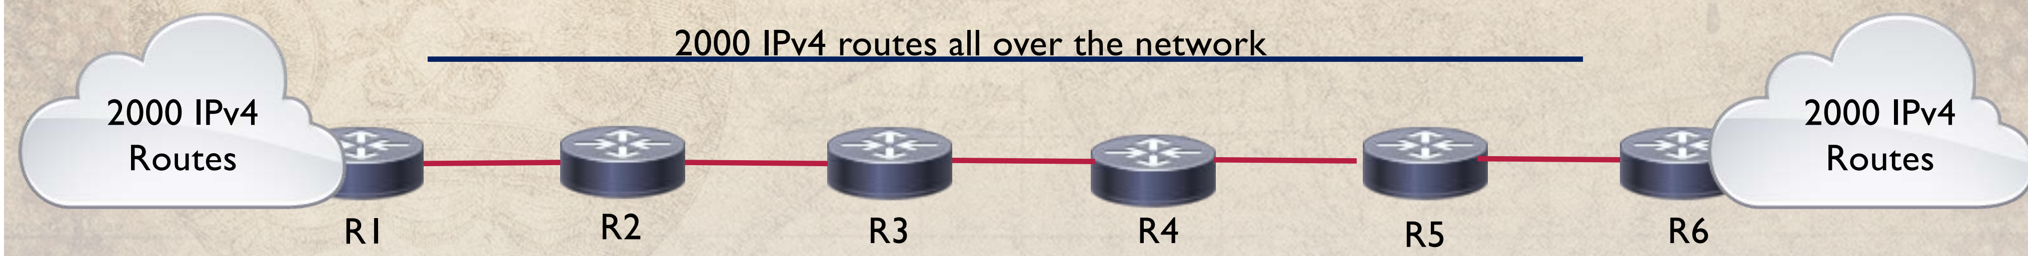

OSPFv2 or ISIS Single Area/Level

```
> Frame 14: 114 bytes on wire (912 bits), 114 bytes captured (912 bits) on interface 0
> Ethernet II, Src: ca:01:2f:14:00:08 (ca:01:2f:14:00:08), Dst: ca:00:10:94:00:08 (ca:00:10:94:00:08)
> Internet Protocol Version 4, Src: 6.6.6.6, Dst: 12.0.0.1
> Internet Control Message Protocol
```

THE HAPPY IPV4 WORLD

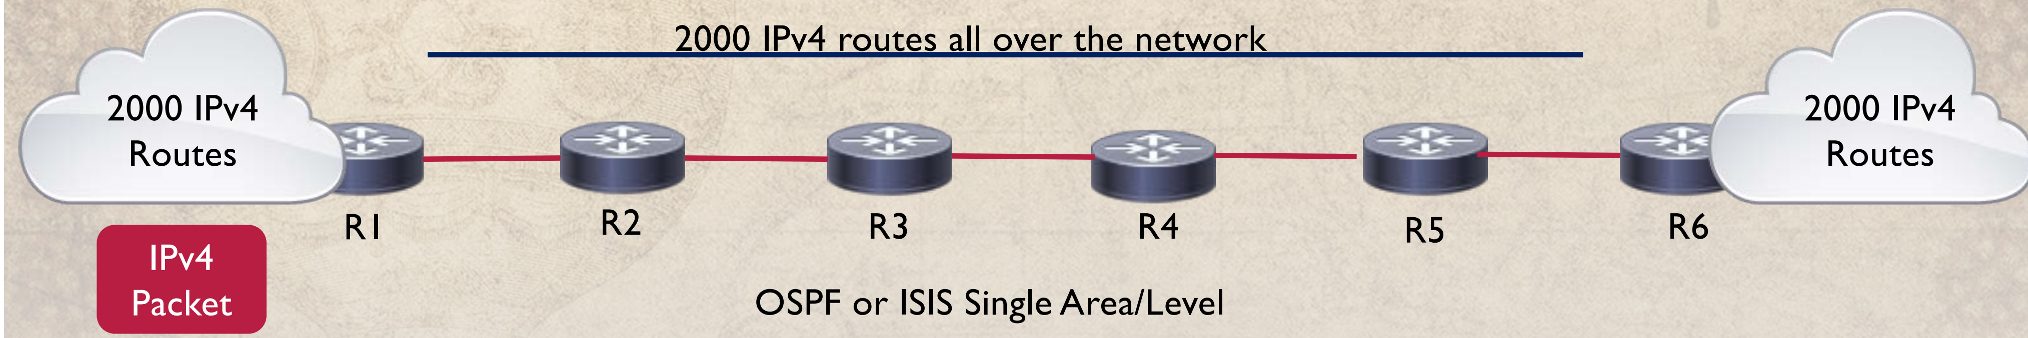

THE HAPPY IPV4 WORLD

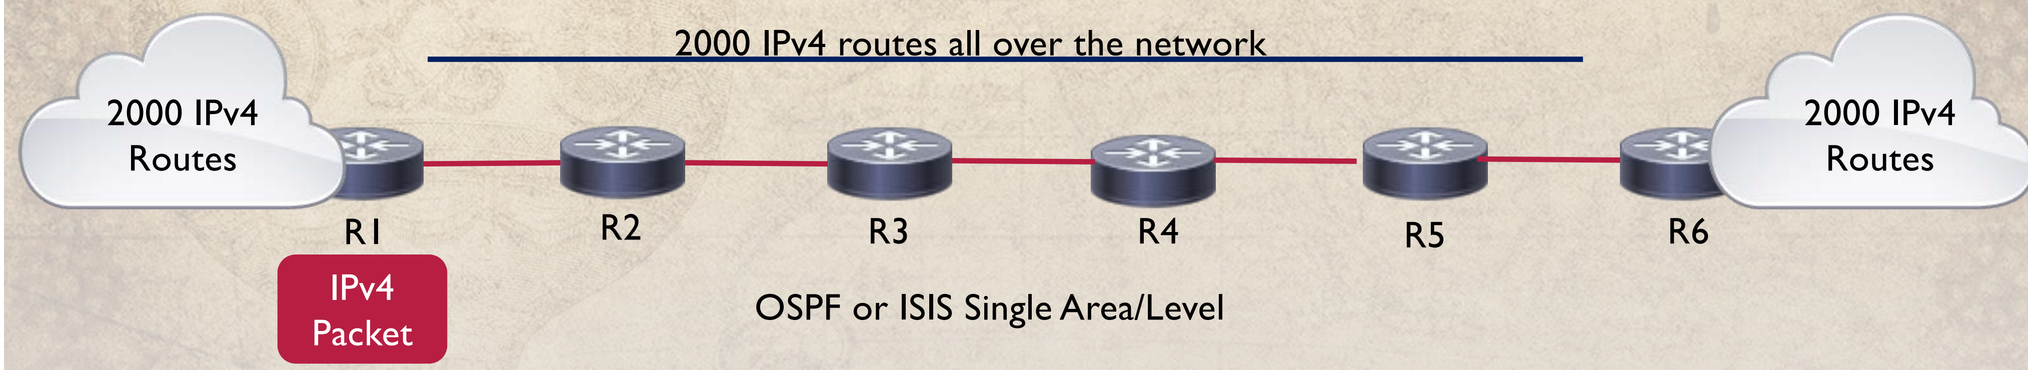

THE HAPPY IPV4 WORLD

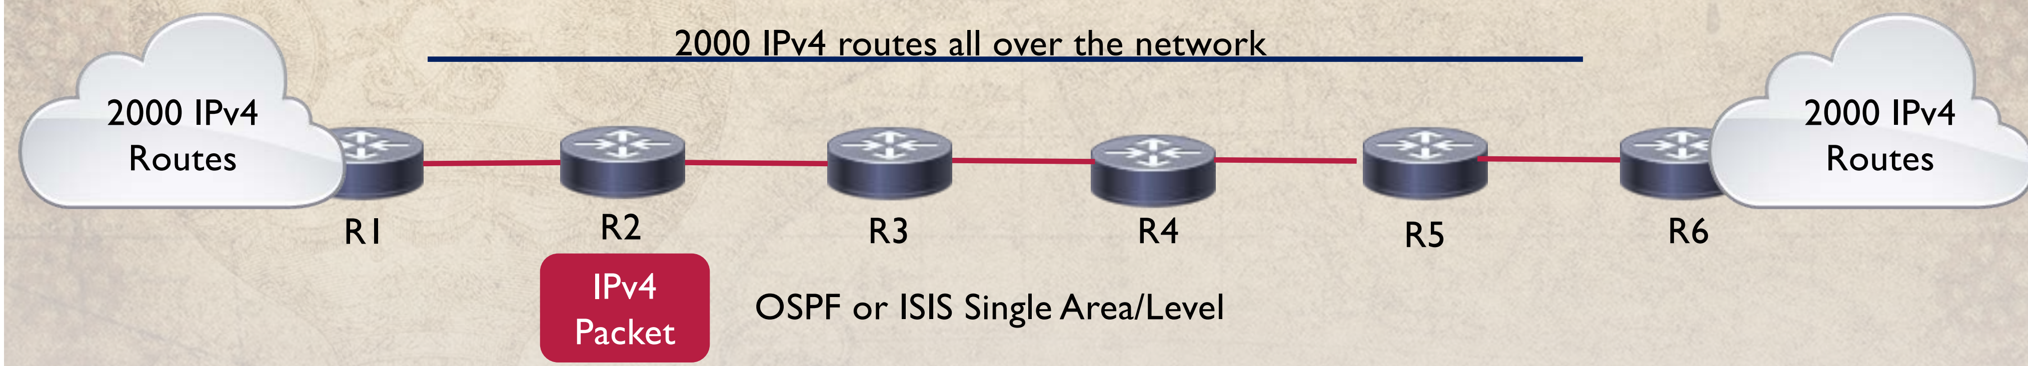

THE HAPPY IPV4 WORLD

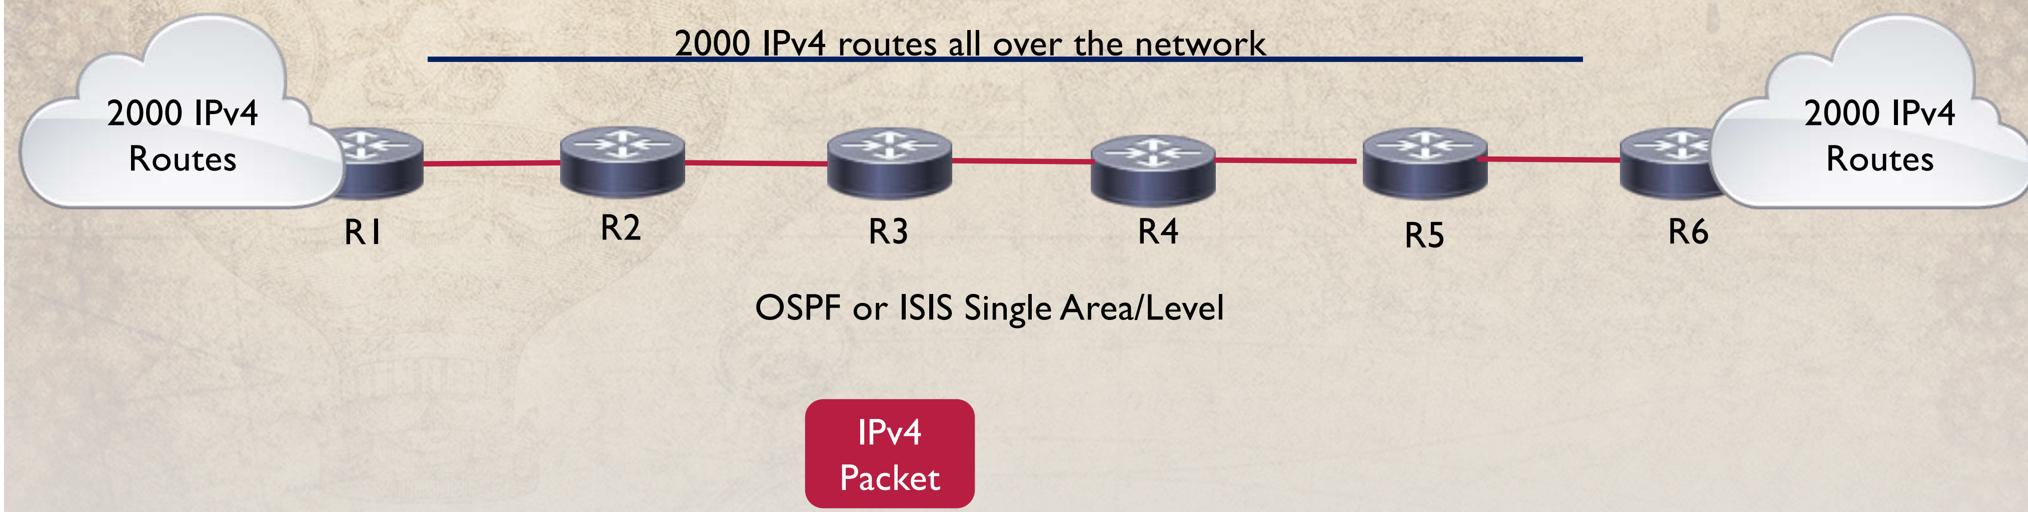

THE HAPPY IPV4 WORLD

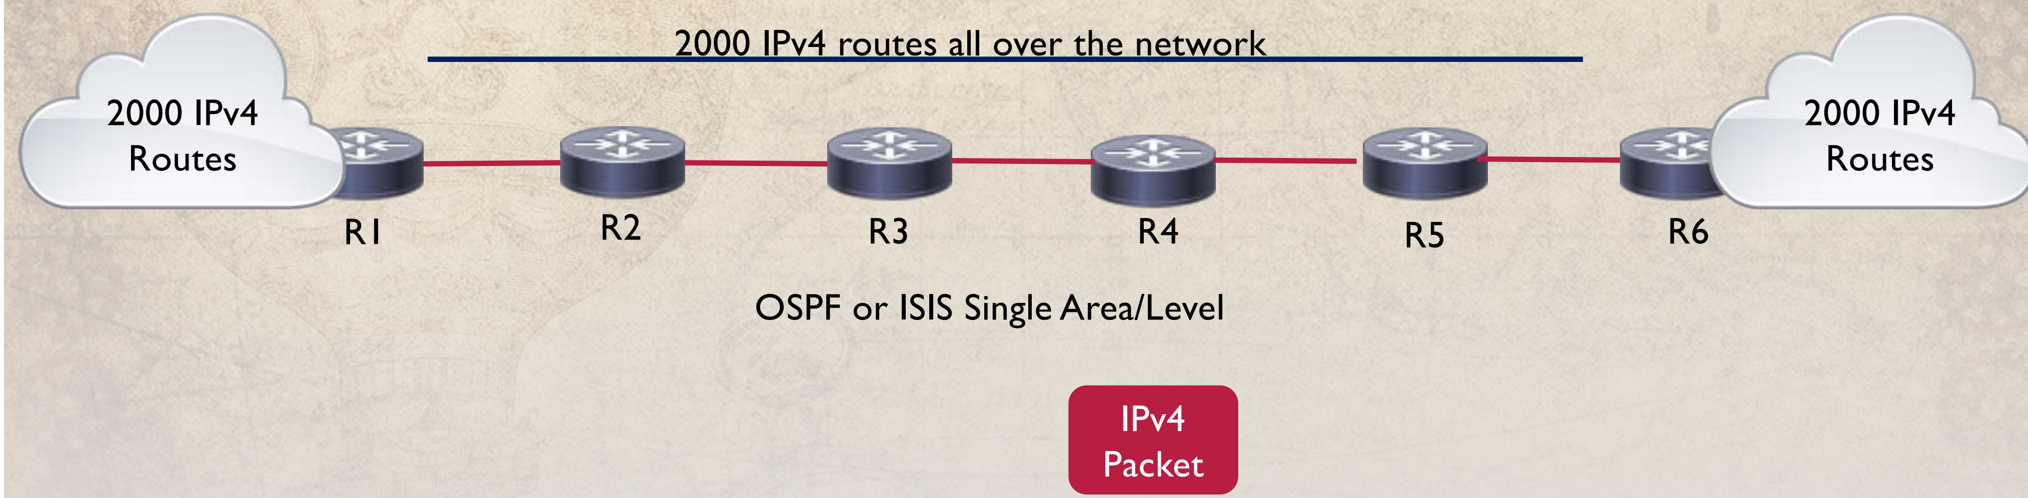

THE HAPPY IPV4 WORLD

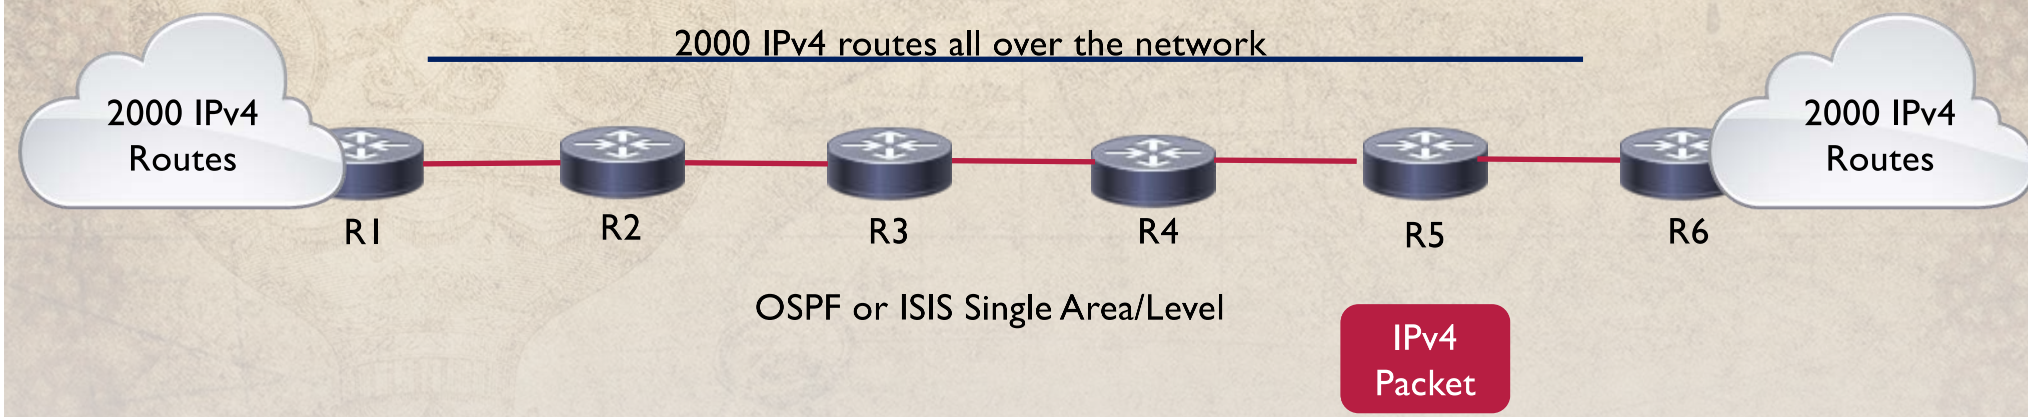

THE HAPPY IPV4 WORLD

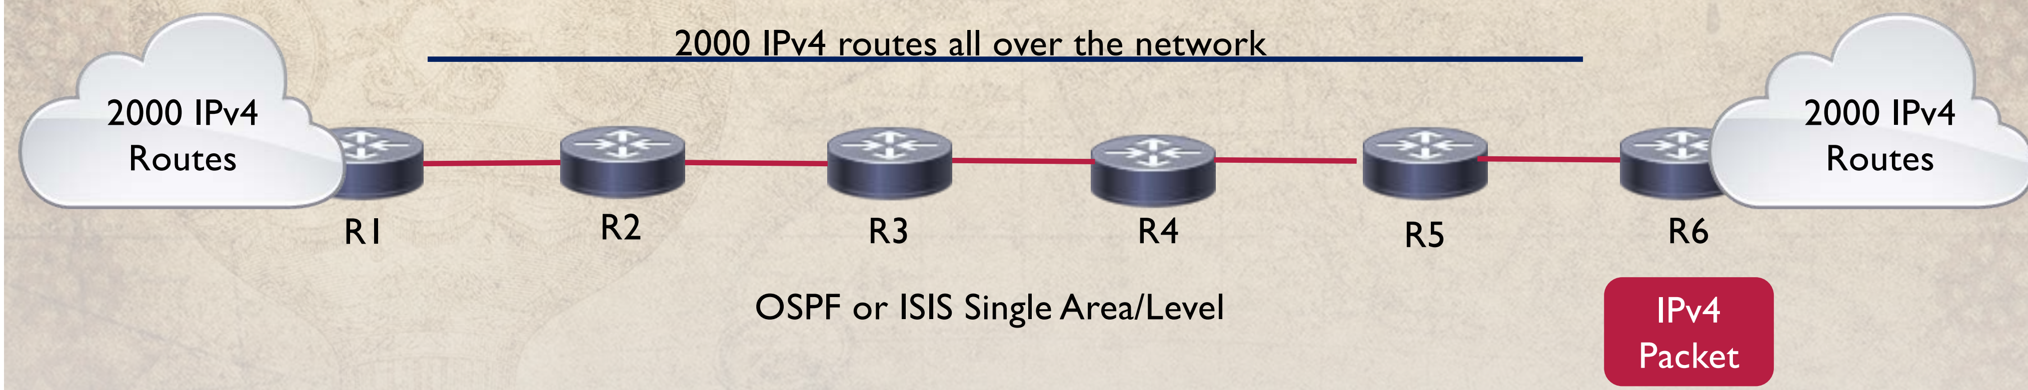

THE HAPPY IPV4 WORLD

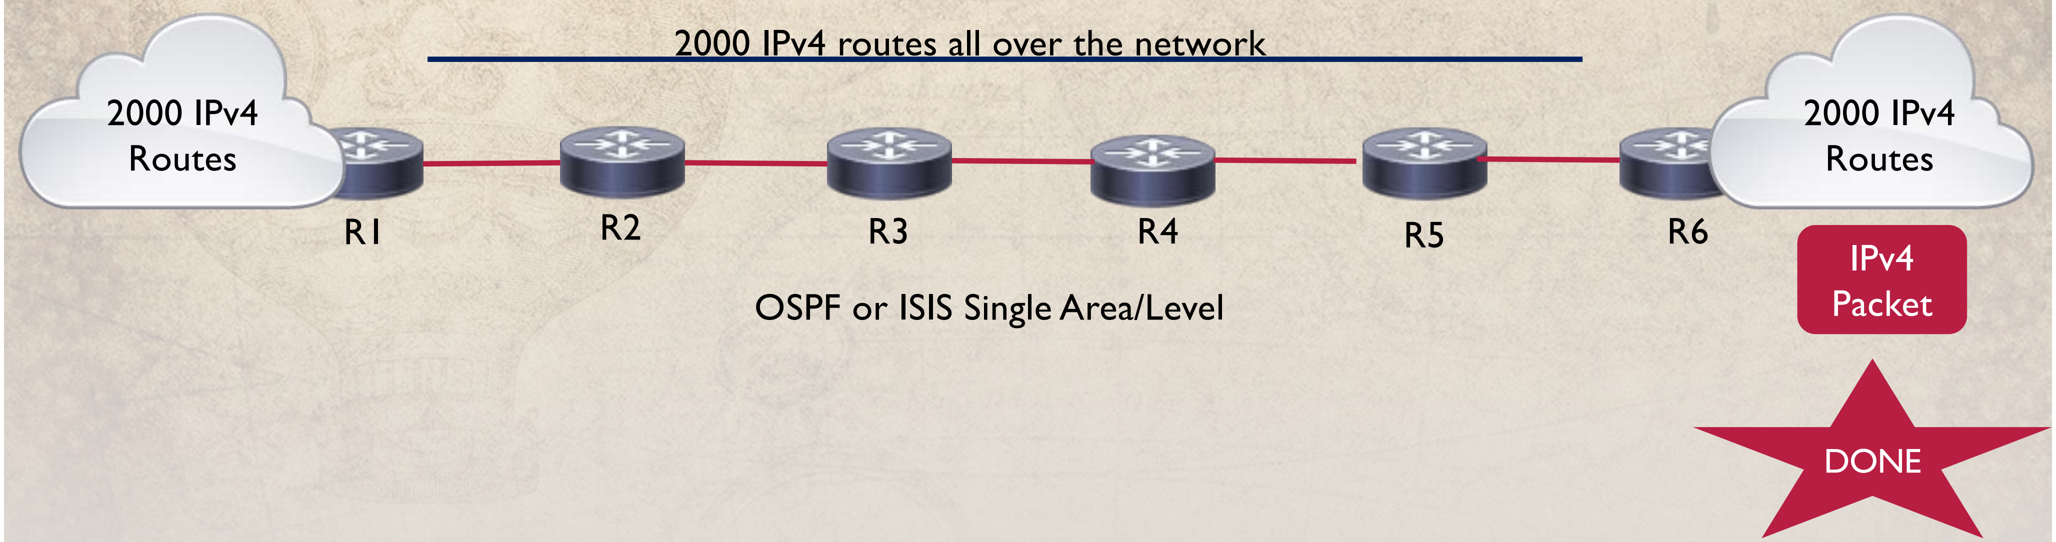

THE SAME STORY BUT AN IPV6 PACKET

IPv6 packet
Typical Layer 2
Typical Layer 3 (IP v6 src/dst)
ICMPv6

```
> Frame 394: 114 bytes on wire (912 bits), 114 bytes captured (912 bits) on interface 0
> Ethernet II, Src: ca:01:2f:14:00:08 (ca:01:2f:14:00:08), Dst: ca:00:10:94:00:08 (ca:00:10:94:00:08)
> Internet Protocol Version 6, Src: 2002::2, Dst: 2002::1
> Internet Control Message Protocol v6
```

THE IPV4 NETWORK WITH IPV6 ISLANDS

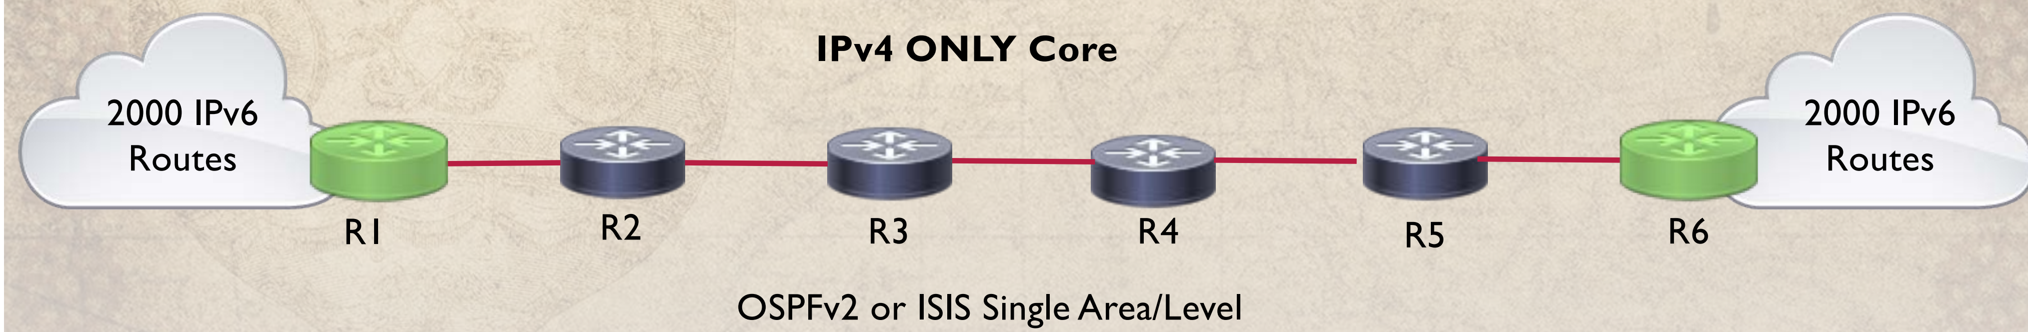

THE IPV4 NETWORK WITH IPV6 ISLANDS

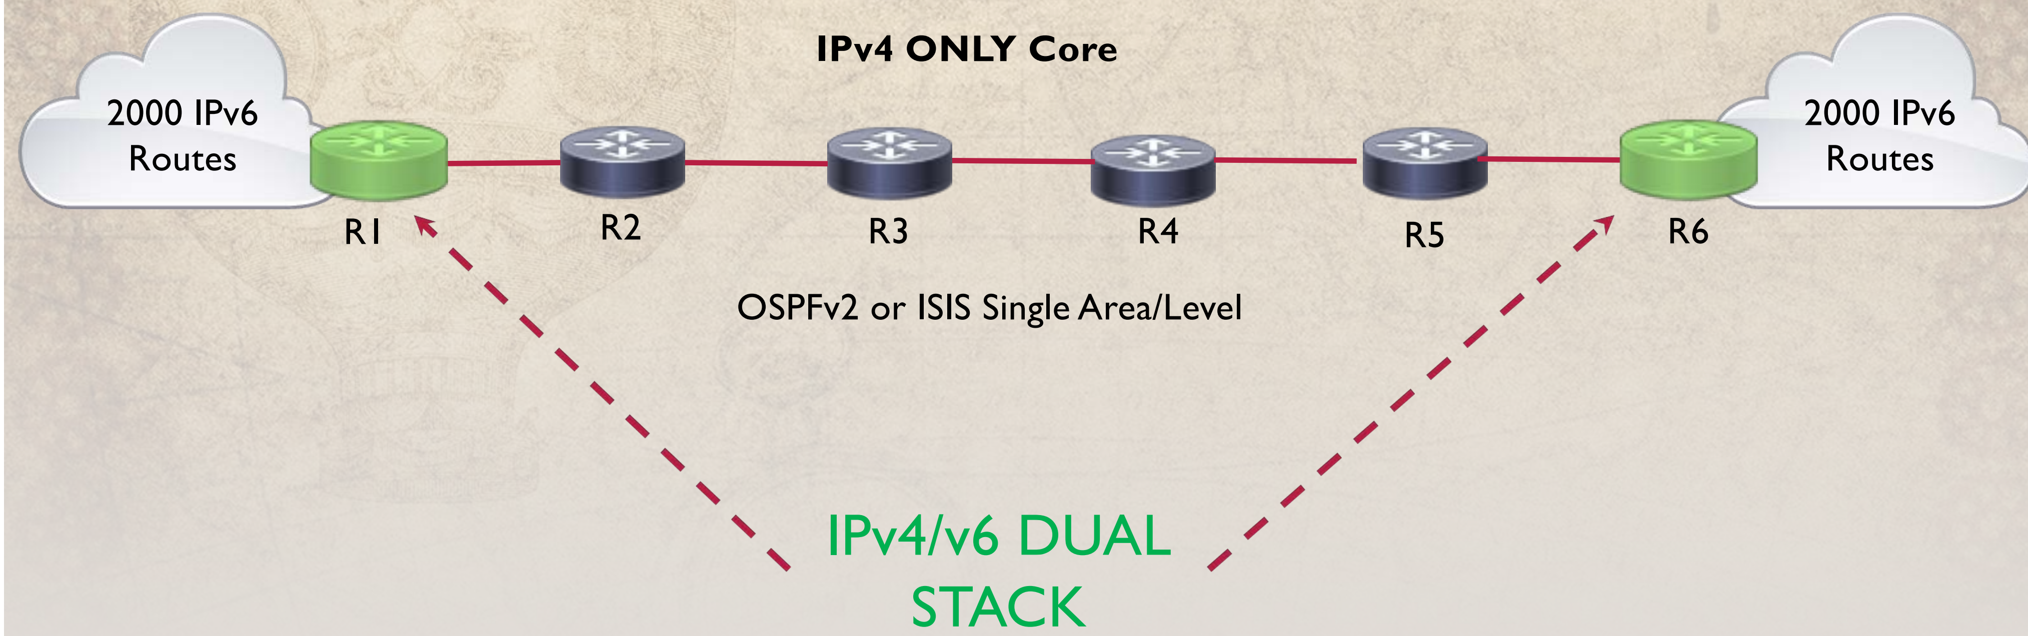

THE IPV4 NETWORK WITH IPV6 ISLANDS

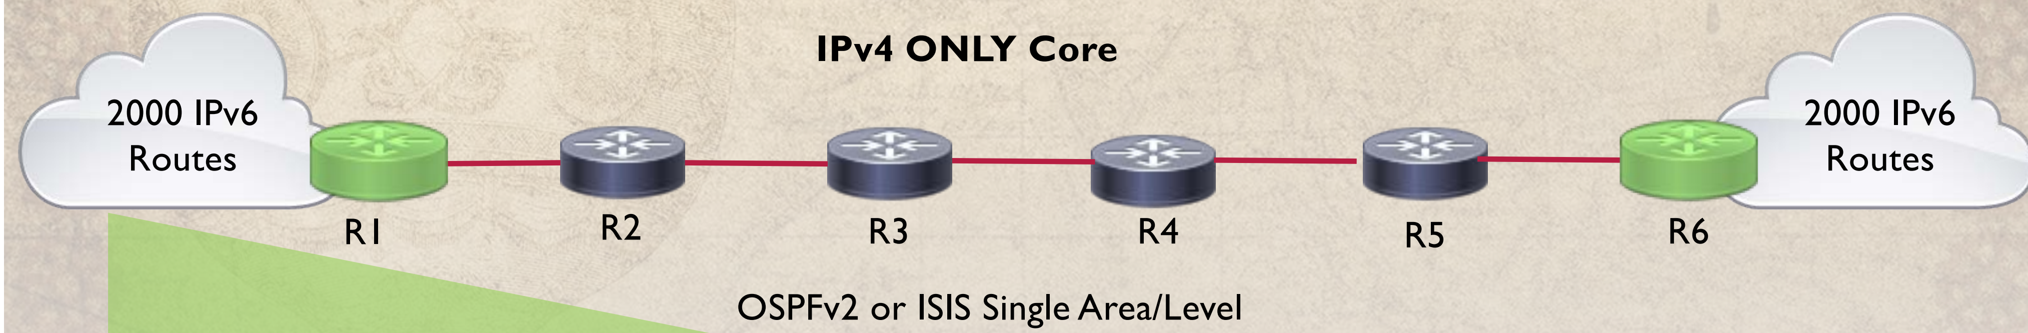


```
> Frame 394: 114 bytes on wire (912 bits), 114 bytes captured (912 bits) on interface 0
> Ethernet II, Src: ca:01:2f:14:00:08 (ca:01:2f:14:00:08), Dst: ca:00:10:94:00:08 (ca:00:10:94:00:08)
> Internet Protocol Version 6, Src: 2002::2, Dst: 2002::1
> Internet Control Message Protocol v6
```

THE IPV4 NETWORK WITH IPV6 ISLANDS

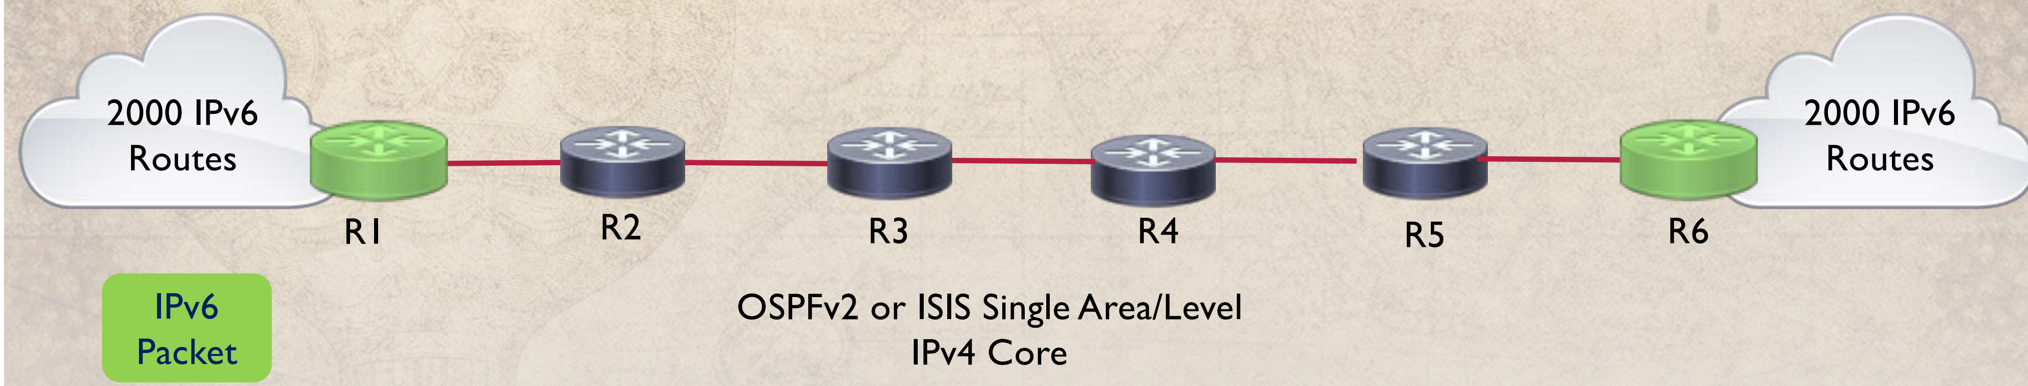

THE IPV4 NETWORK WITH IPV6 ISLANDS

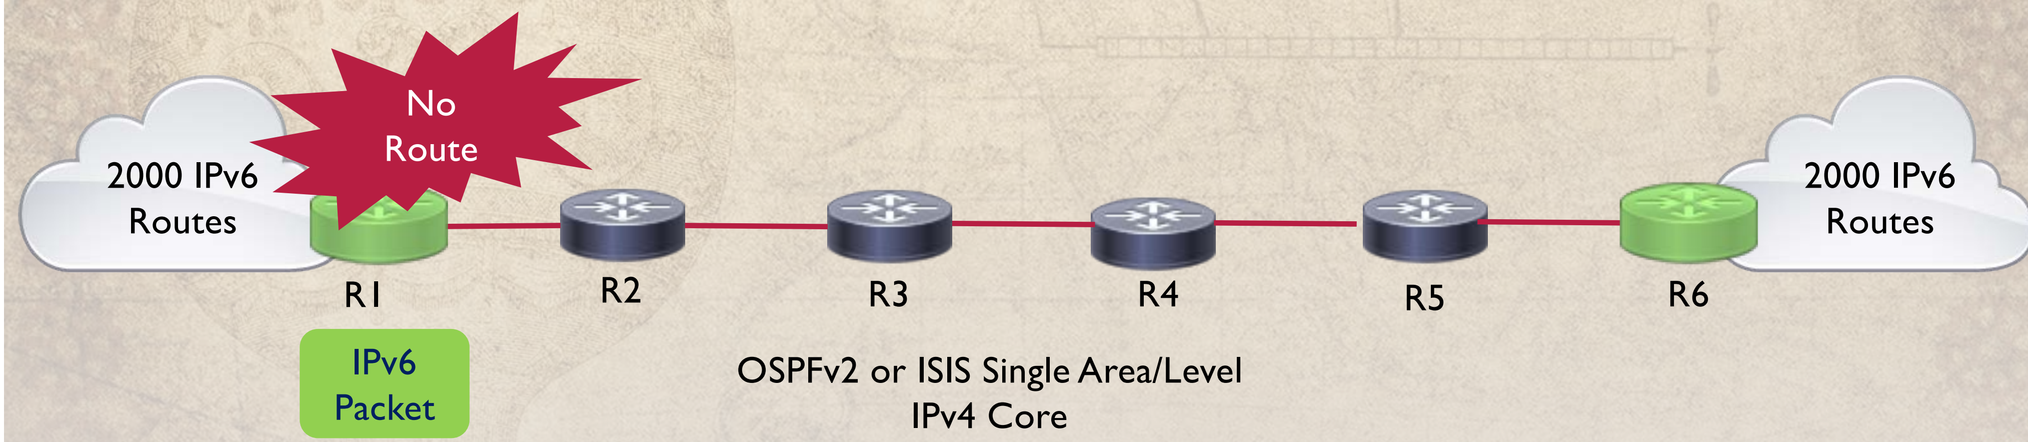

THE IPV4 NETWORK WITH IPV6 ISLANDS

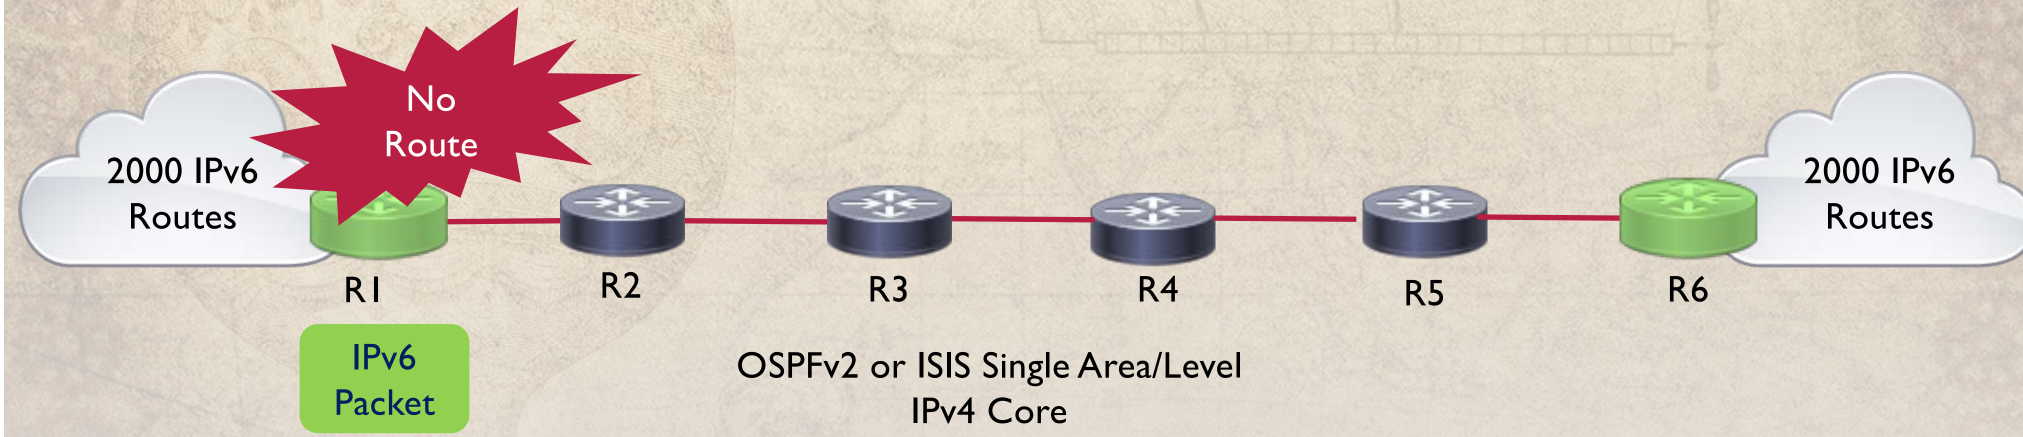

IDEAS?

THE IPV4 NETWORK WITH IPV6 ISLANDS

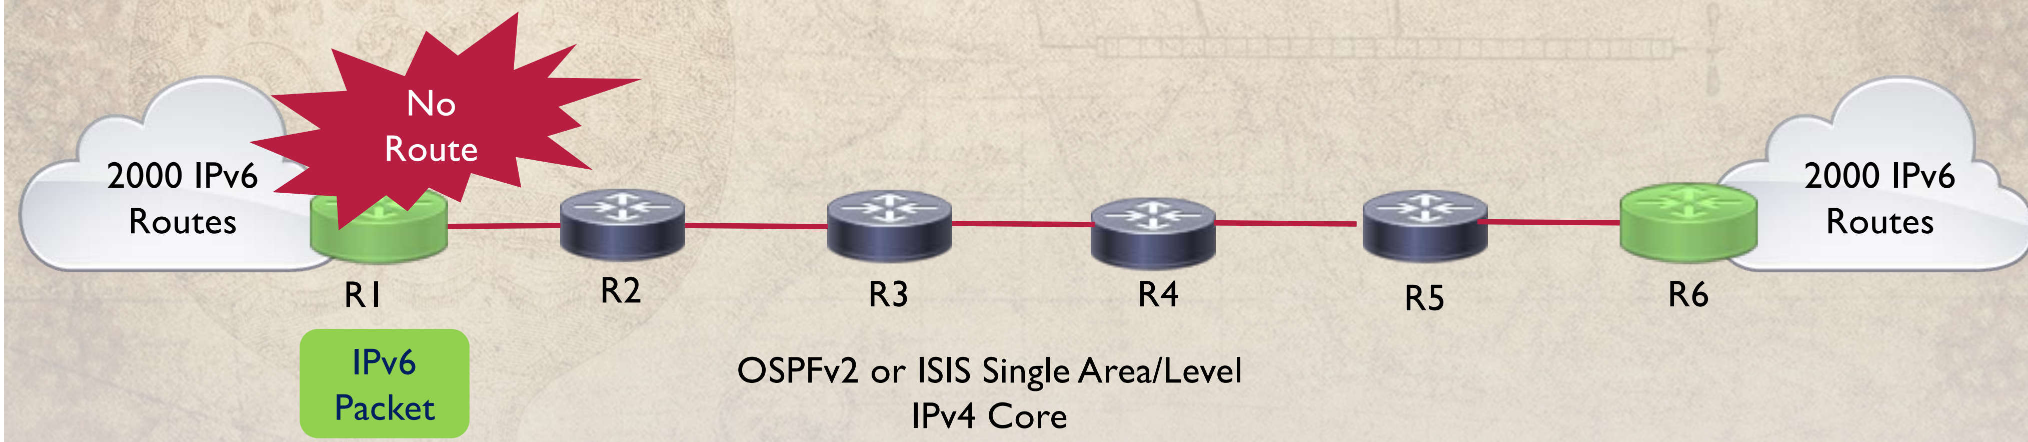

A TUNNEL MAYBE?

LET'S TUNNEL

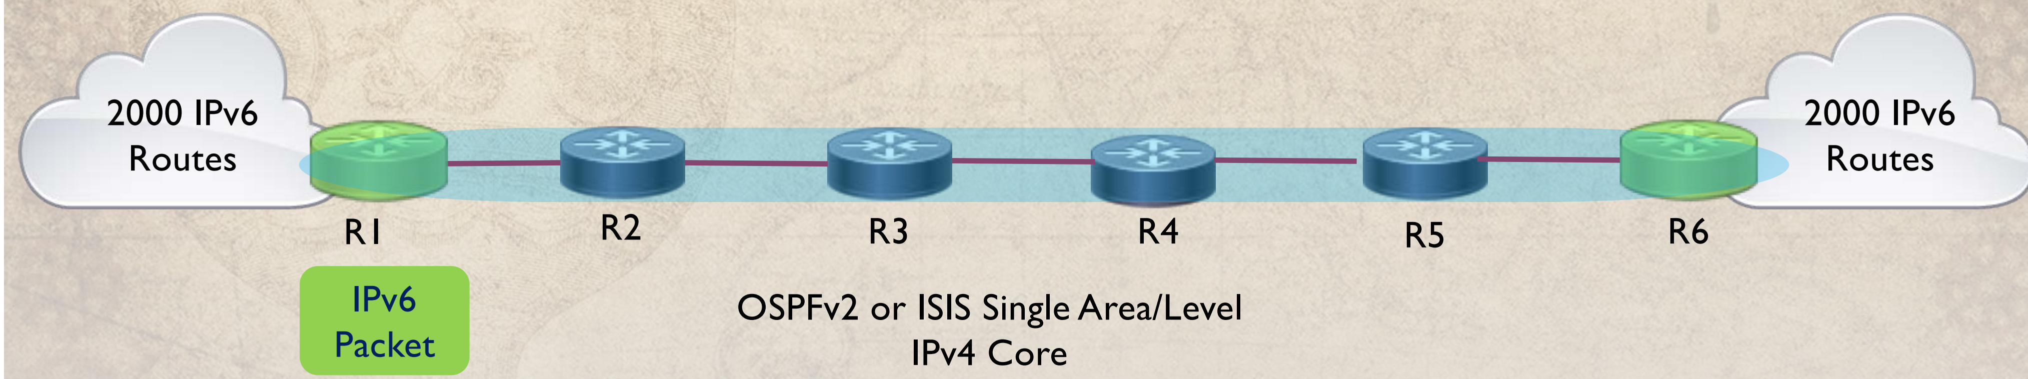

Cross vendor issues

Still evolving but…

LET'S TUNNEL

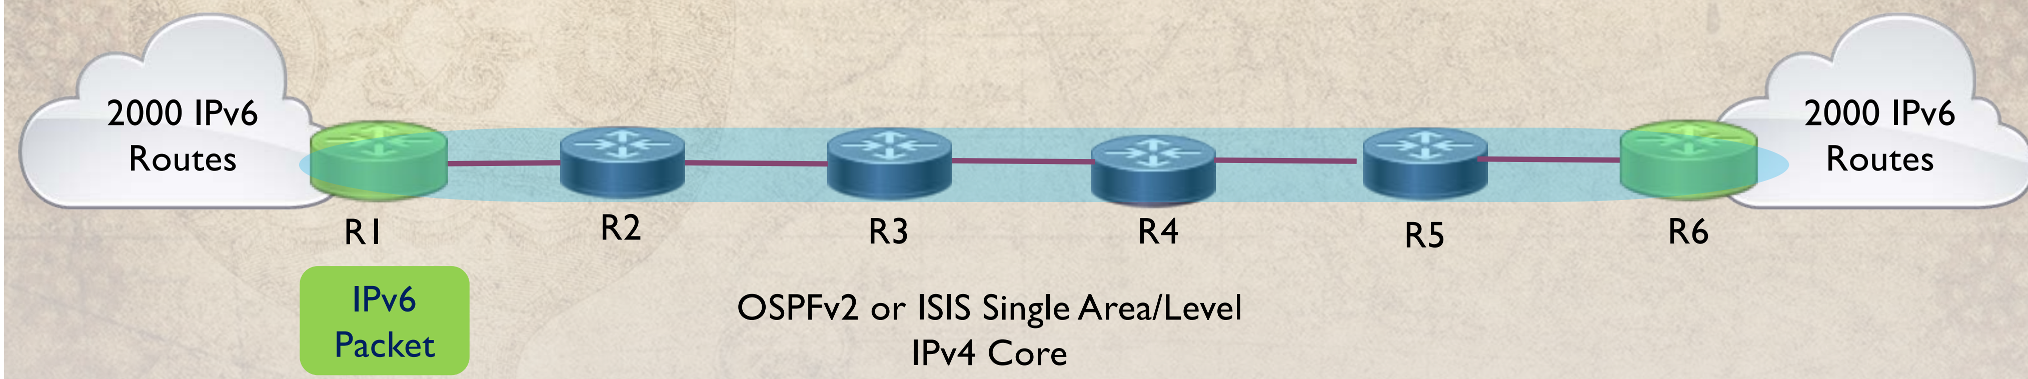

Scalability issues as you grow

FOR NOW…

- Your temporary proposal to the CIO, surgical change.
- Just get them off your back
- Day 1 ideas develop with or without you

In a year or two we will come up with better ideas. For now this will fix it for the two labs or sites. Off my back.

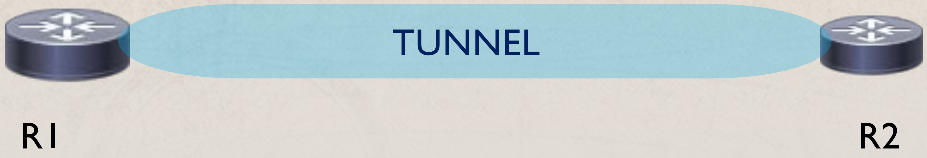

BUT THEY USUALLY DO NOT STOP…

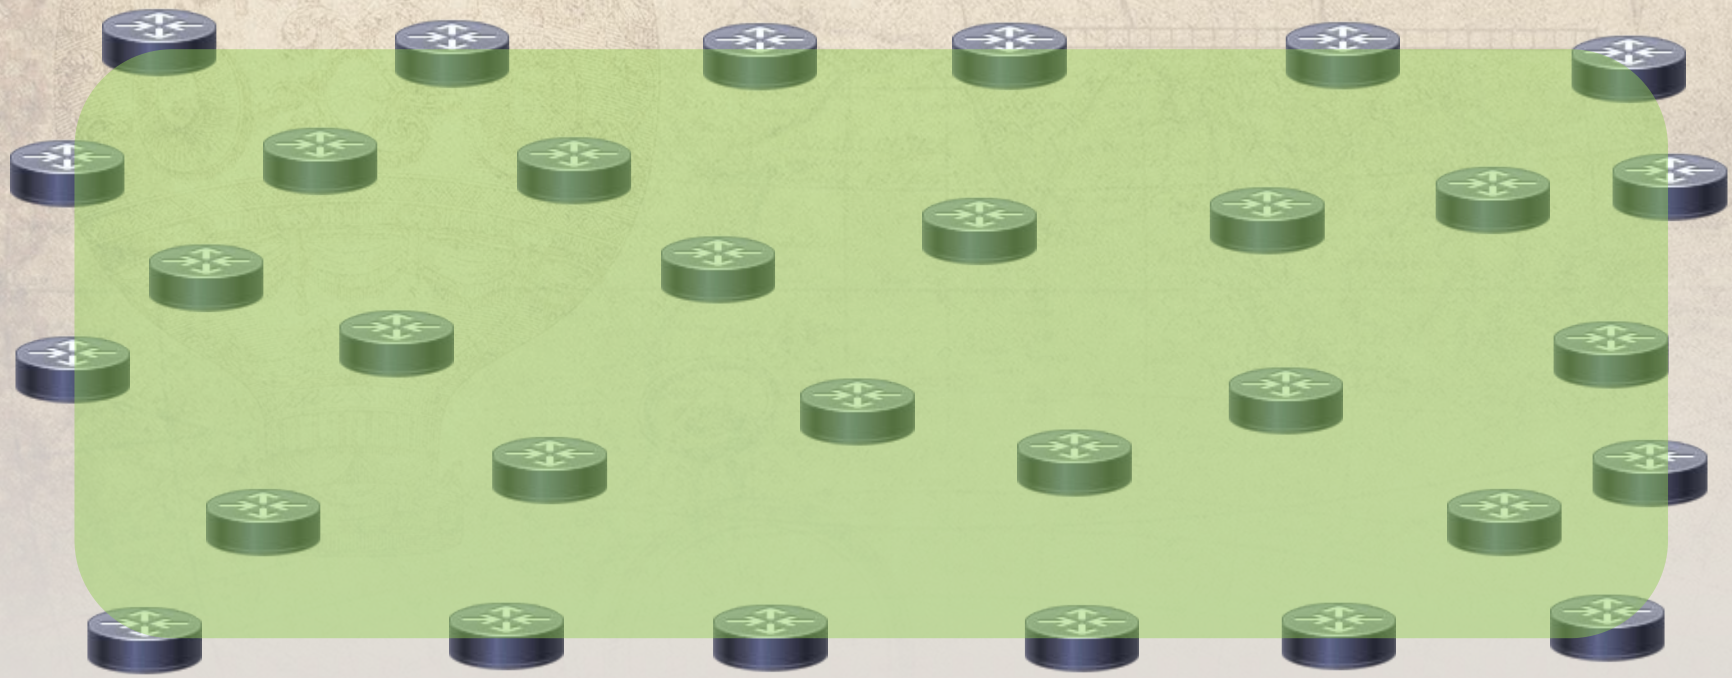

THEY ACTUALLY "GROW" *OR COOK*

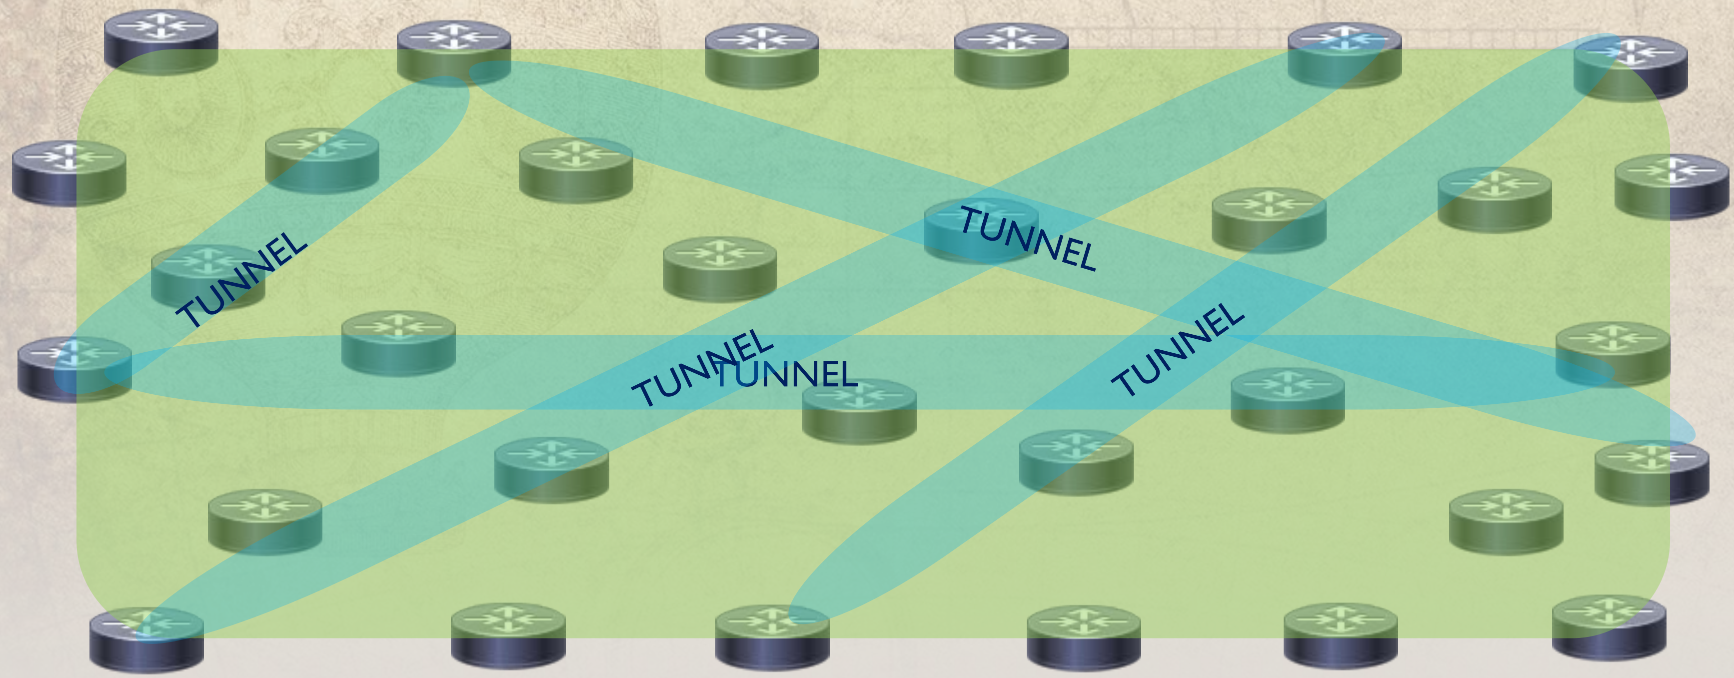

TUNNELS ACTUALLY "COOK" *NOT GROW*

BUT WE ARE SMART AND…

- Won't let that happen
- So after fixing it temporary for 6 months….
- Now we have a plan so we're upgrading everything to dual stack

GUESS WE'RE DONE…

GREAT BUT…

- For most features now you need to activate the IPv6 equivalent of them
  - Which means:
    - You have 50 routers in the core (pretty small scale)
    - Now you need the IPv6 version of them on 50 new routers
      - Automation is your friend!
    - Does your vendor support them for your code? Multicast maybe?

OH YEAH ALL DAY LONG…

- Just got off the phone with both vendors and they have the IPv6 equivalent of everything

NOW…

"Welcome to the first day of the rest of your life" – Negan.

CAN’T BE **THAT** BAD…

- To what extent?
- New features sometimes mean new codes (scrub it?)
- IPv6 features in many vendors are not “beaten up”
- Does your hardware support them?
- New codes means more maintenance windows
- New features (or even hardware) means new bugs
- Stability in the core?

CAN'T BE **THAT** BAD…

- New bugs perhaps new security issues in the core?
- How about core automation and monitoring?
- The learning curve…times the number of vendors!

A LITTLE BEYOND OUR SCOPE

- Routing instability
- Systems readiness (usually beyond your control)
- Unpredicted behavior

SHOW ME THE REQUIREMENTS

- Okay, it's complicated. But show me the requirements.

THE BEST OF BOTH WORLD?

- Want to have to end IPv6 reachability
- Something point to multi-point without having the "Spaghetti effect"
- Do not want to introduce a lot of new features to the core (remember Negan?)
- Do not want downtime in the core (code upgrades etc.)
- Want something that's proven and tested across multiple vendors for years at scale

THE BEST OF BOTH WORLD?

- Want something that the largest operators in the world beat up in “in the core”
- Would be nice to have some steering control over my traffic later
- Keep the v6 in my enclaves for now
- Do not want to add anything major to monitoring, Ops and automation burden
- Perhaps not a steep learning curve

THE BEST OF BOTH WORLD?

IDEALLY :

- Quick expansion and scalable
  - Might end up with 1000's of IPv6 routes
  - Can I keep them out of the core?
- Some bonus points too maybe?
  - Future VLAN/L2 extension
  - Support VRF's as well (a V-R-F or a Vurf!)
  - Very low convergence time (maybe 50-100ms?)

WHAT DID THE TUNNEL DO FOR YOU?

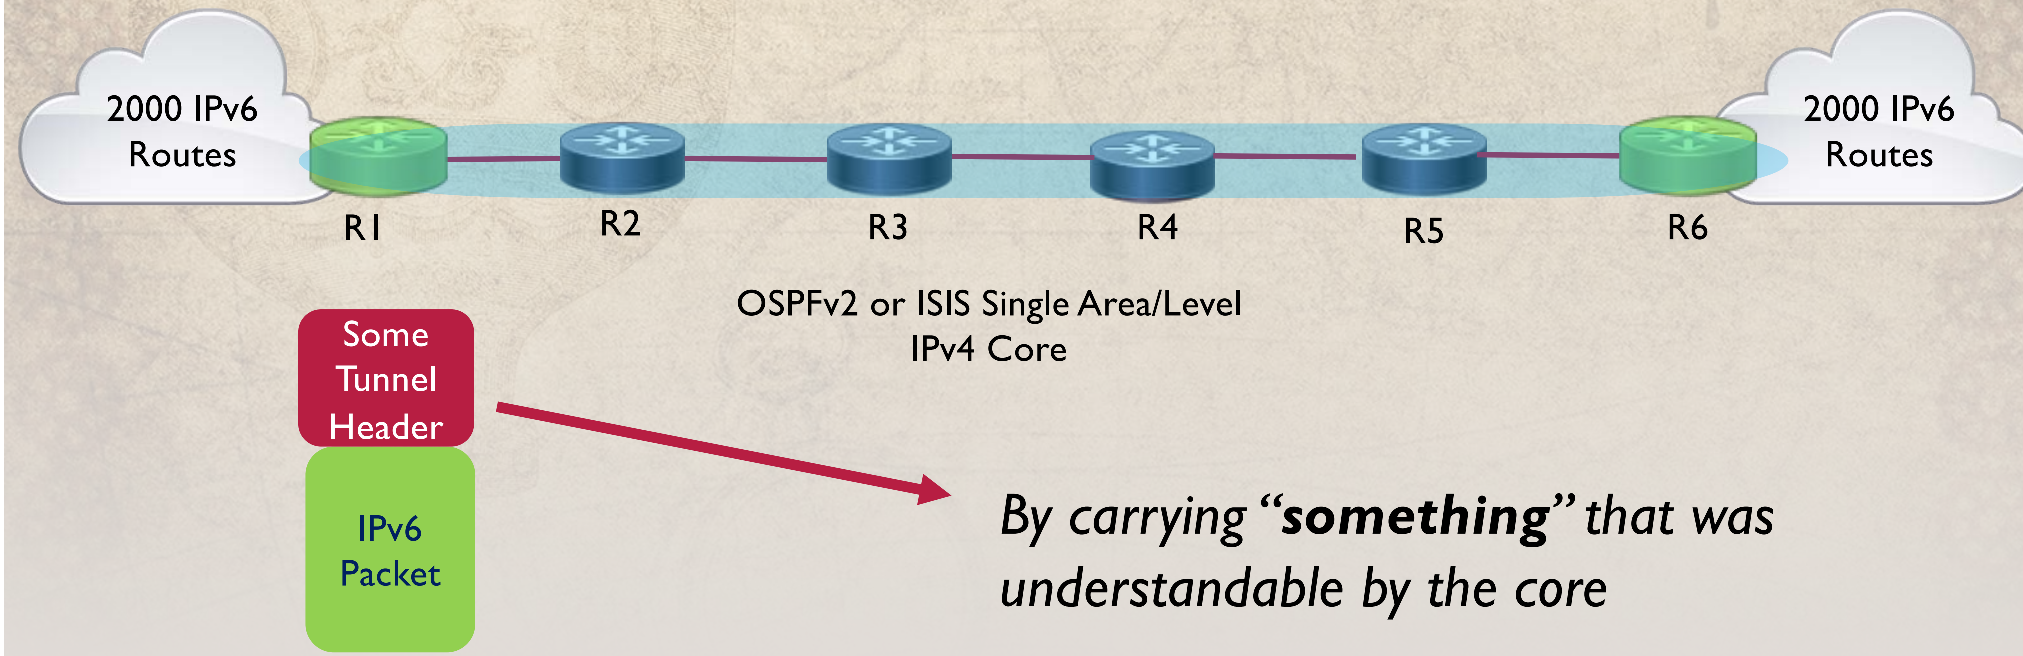

*By carrying "**something**" that was understandable by the core*

REMEMBER OUR ANALOGY?

- Maybe that **something** can be a number (a label?)

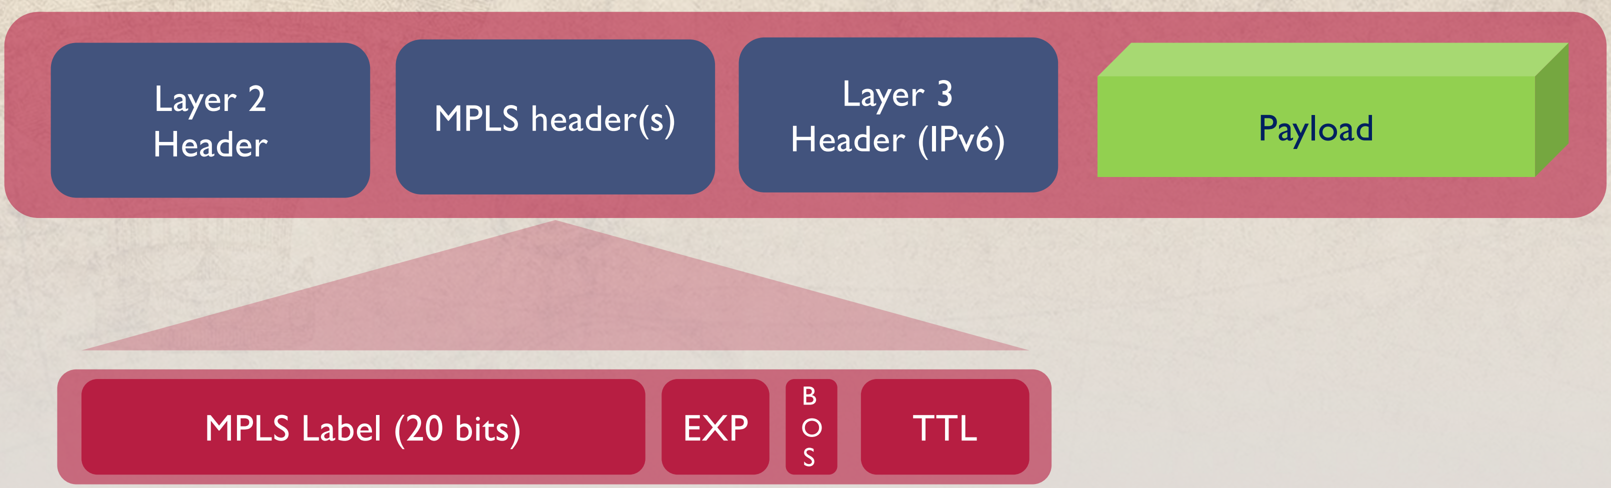

LET'S TAKE A CLOSER LOOK

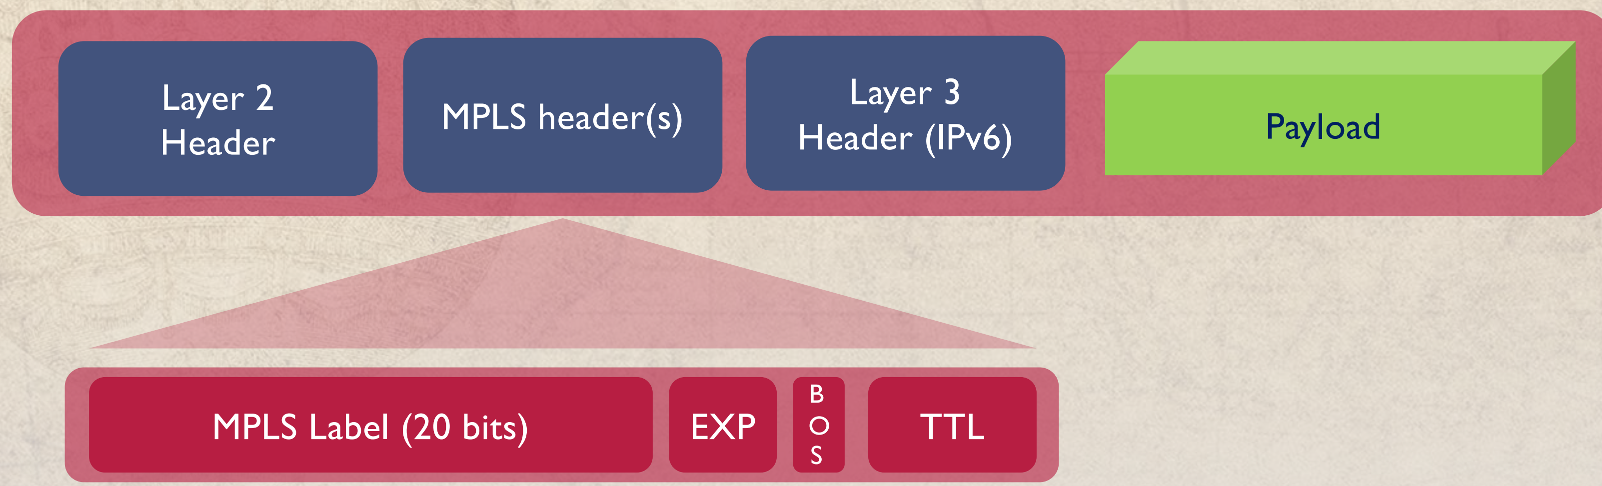


```
> Frame 65: 82 bytes on wire (656 bits), 82 bytes captured (656 bits) on interface 0
> Ethernet II, Src: ca:03:0b:f8:00:08 (ca:03:0b:f8:00:08), Dst: ca:02:00:08:00:1c (ca:02:00:08:00:1c)
> MultiProtocol Label Switching Header, Label: 17, Exp: 6, S: 0, TTL: 253
> MultiProtocol Label Switching Header, Label: 25, Exp: 6, S: 1, TTL: 255
> Internet Protocol Version 6, Src: 2006::6, Dst: 2001::200
> Transmission Control Protocol, Src Port: 80, Dst Port: 60121, Seq: 1, Ack: 11, Len: 0
```

LET'S TAKE A CLOSER LOOK

- Header of a HTTP packet in IPv6
- Going from 2006::6 (www server) to 2001::200 (client)

```
> Frame 65: 82 bytes on wire (656 bits), 82 bytes captured (656 bits) on interface 0
> Ethernet II, Src: ca:03:0b:f8:00:08 (ca:03:0b:f8:00:08), Dst: ca:02:00:08:00:1c (ca:02:00:08:00:1c)
> MultiProtocol Label Switching Header, Label: 17, Exp: 6, S: 0, TTL: 253
> MultiProtocol Label Switching Header, Label: 25, Exp: 6, S: 1, TTL: 255
> Internet Protocol Version 6, Src: 2006::6, Dst: 2001::200
> Transmission Control Protocol, Src Port: 80, Dst Port: 60121, Seq: 1, Ack: 11, Len: 0
```

TOP LABEL

BOTTOM LABEL

A BIT OF MPLS INTERNALS

- Each router assigns labels to its reachable destinations (for the most part)
- And tells its neighbors about his assignment

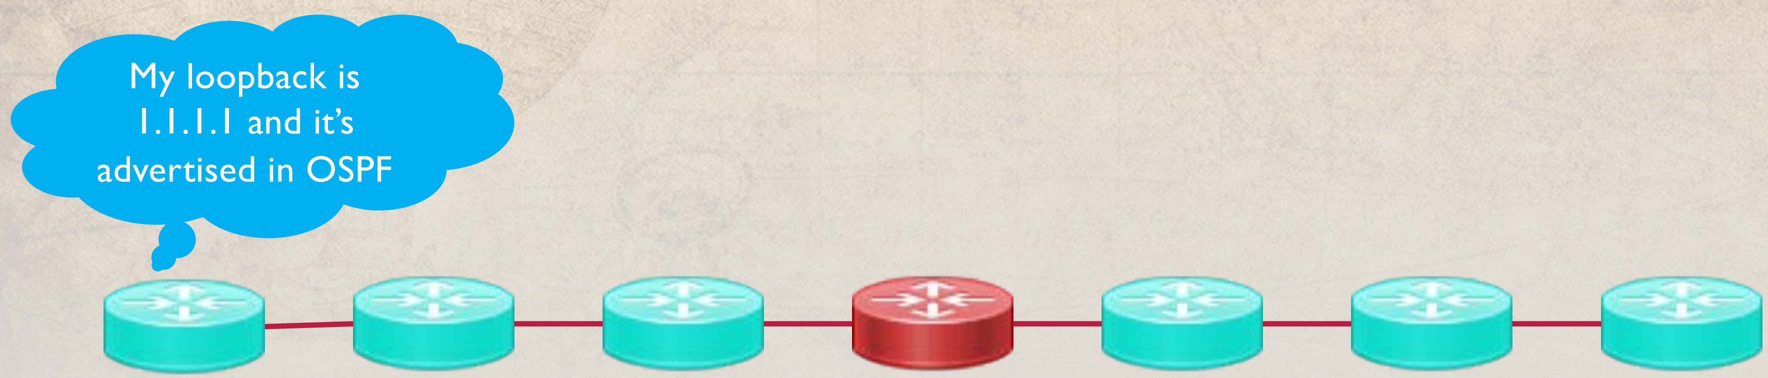

A BIT OF MPLS INTERNALS

- Each router assigns labels to its reachable destinations
- And tells its neighbors about his assignment

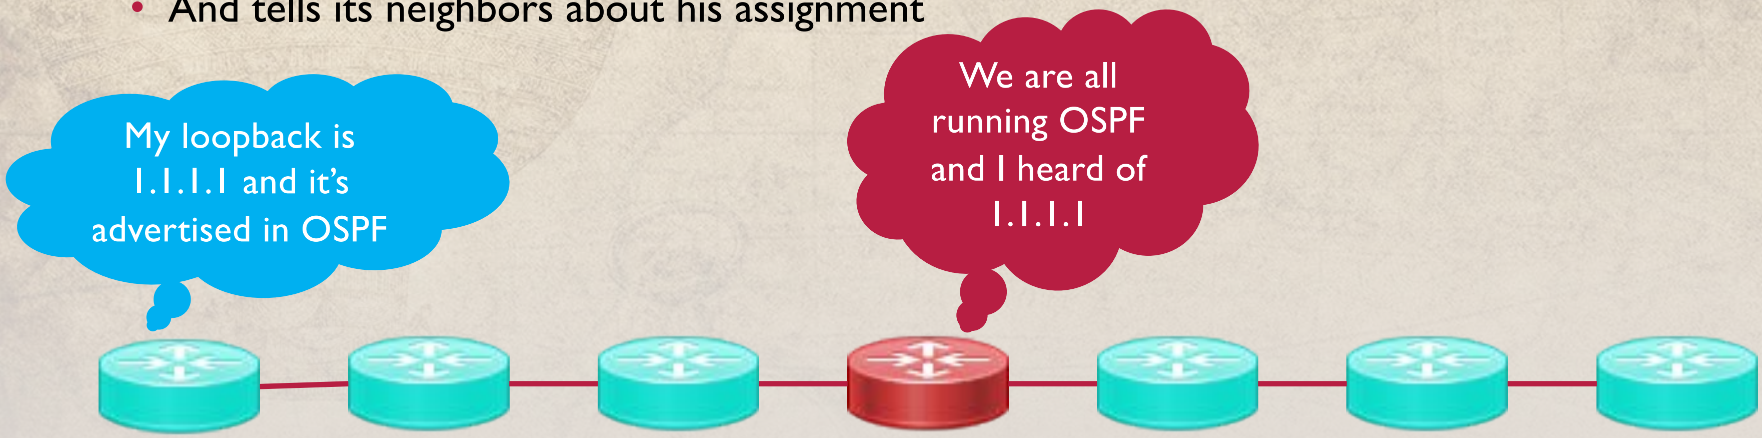

A BIT OF MPLS INTERNALS

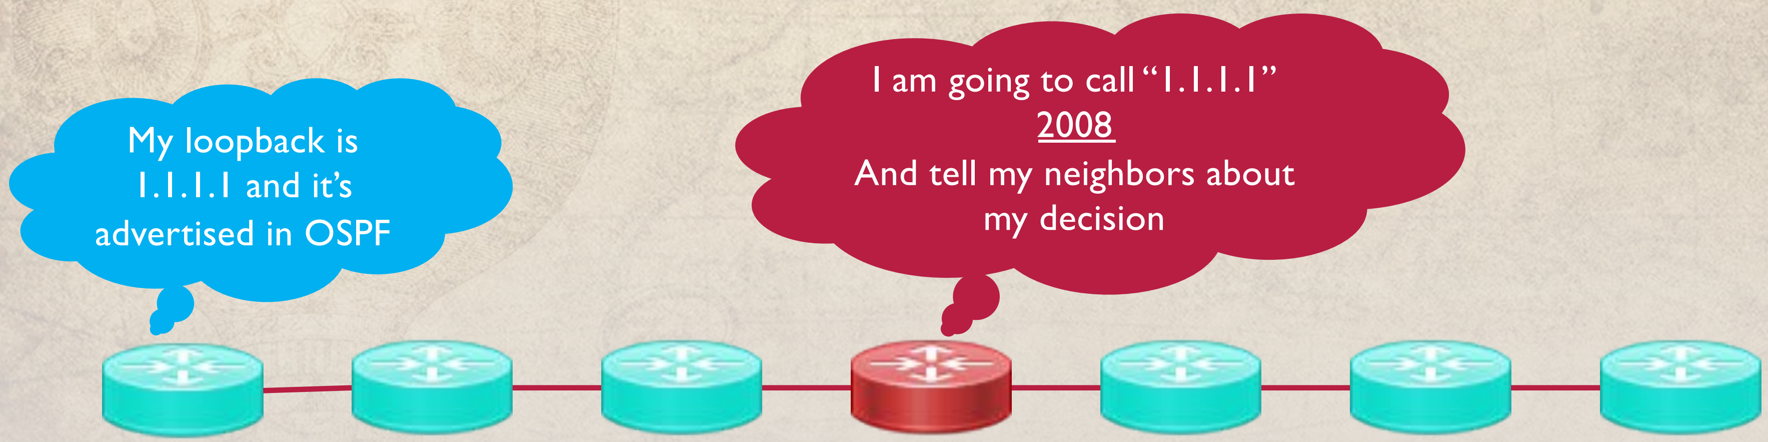

A BIT OF MPLS INTERNALS

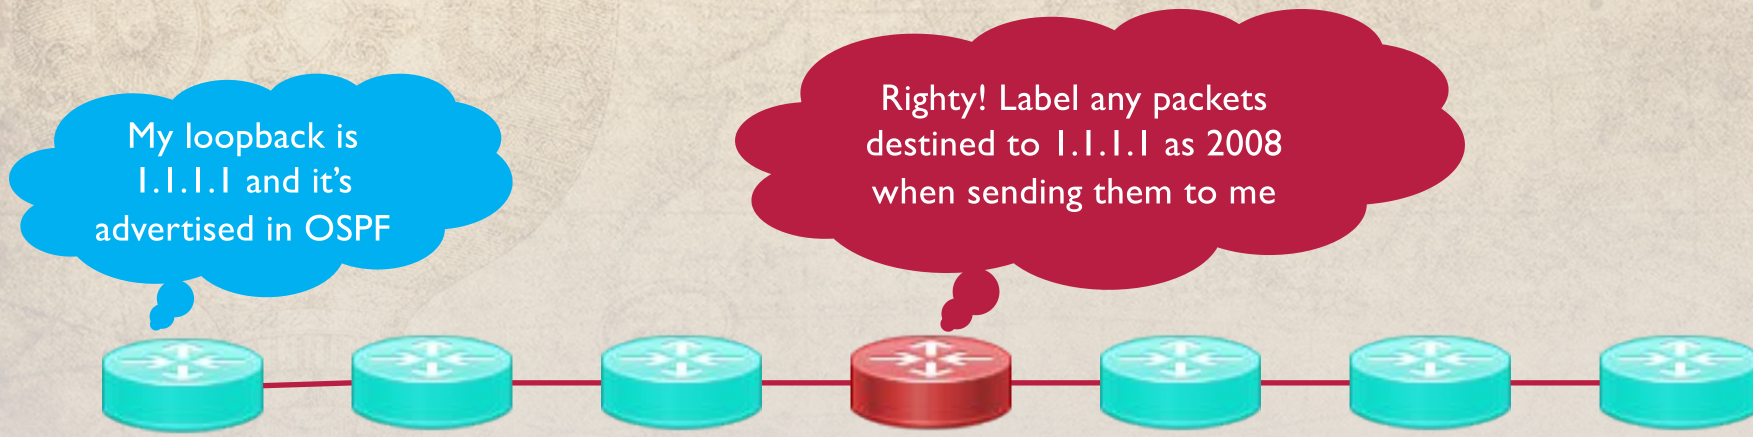

A BIT OF MPLS INTERNALS

A BIT OF MPLS INTERNALS

Everybody else does the same. So every router knows about their label assignments of their neighbors. So they label packets accordingly when sending packets to them.

LDP (Label Distribution Protocol) is one way of doing it

A BIT OF MPLS INTERNALS

IP Packet — IP packet destined for 1.1.1.1

— MPLS label

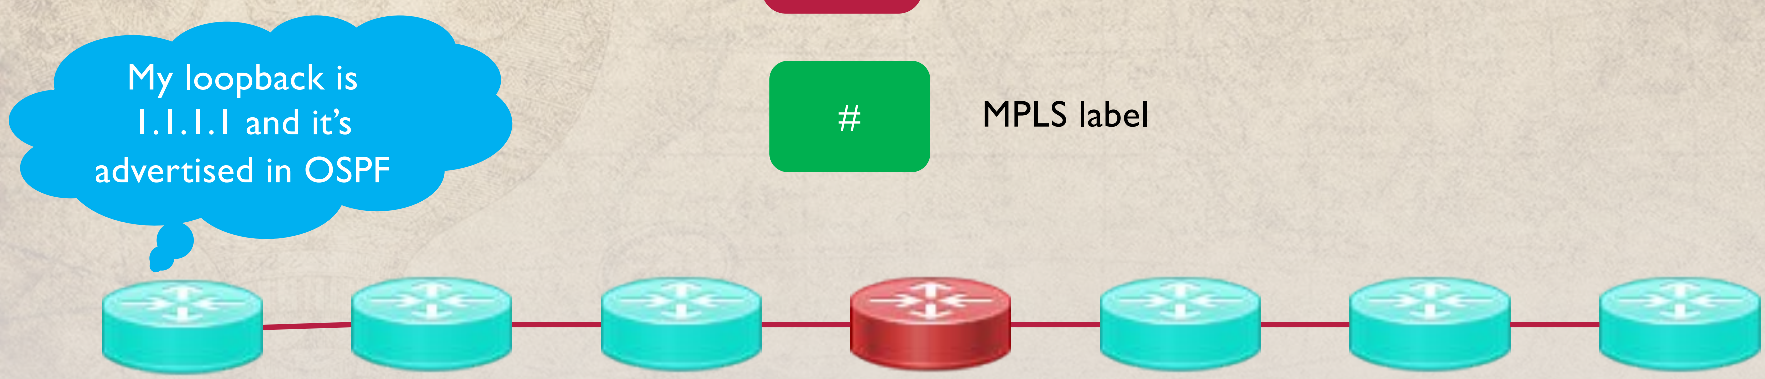

Loop0: 1.1.1.1/32

A BIT OF MPLS INTERNALS

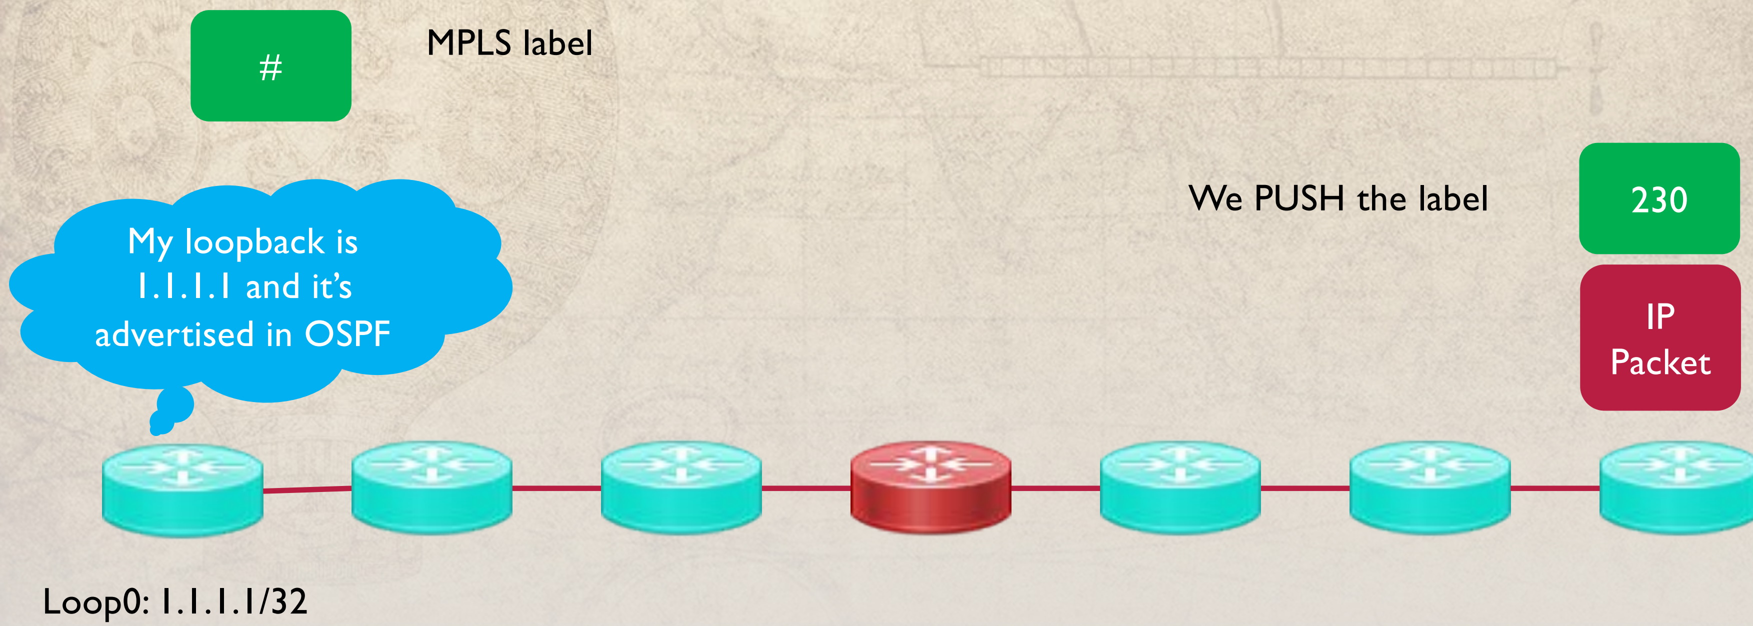

A BIT OF MPLS INTERNALS

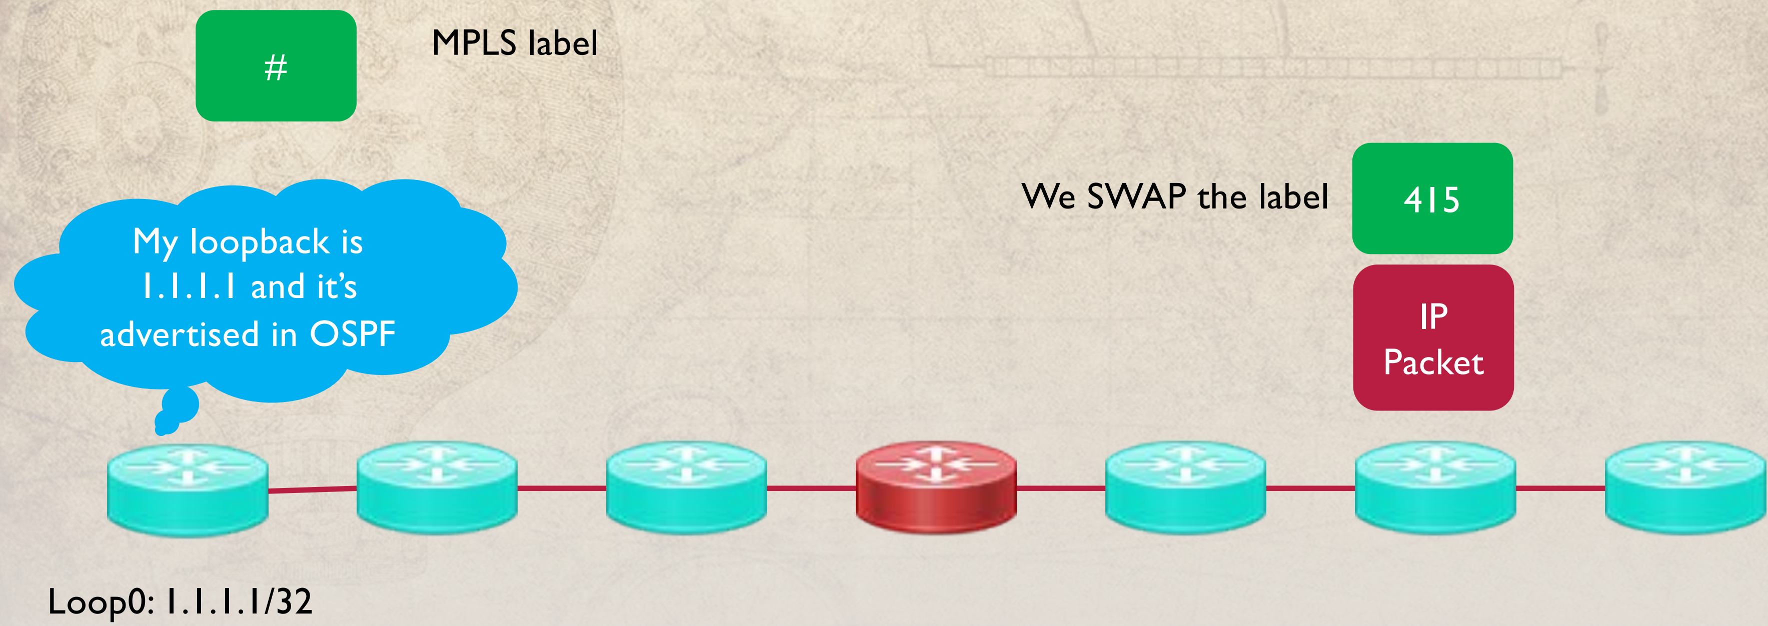

A BIT OF MPLS INTERNALS

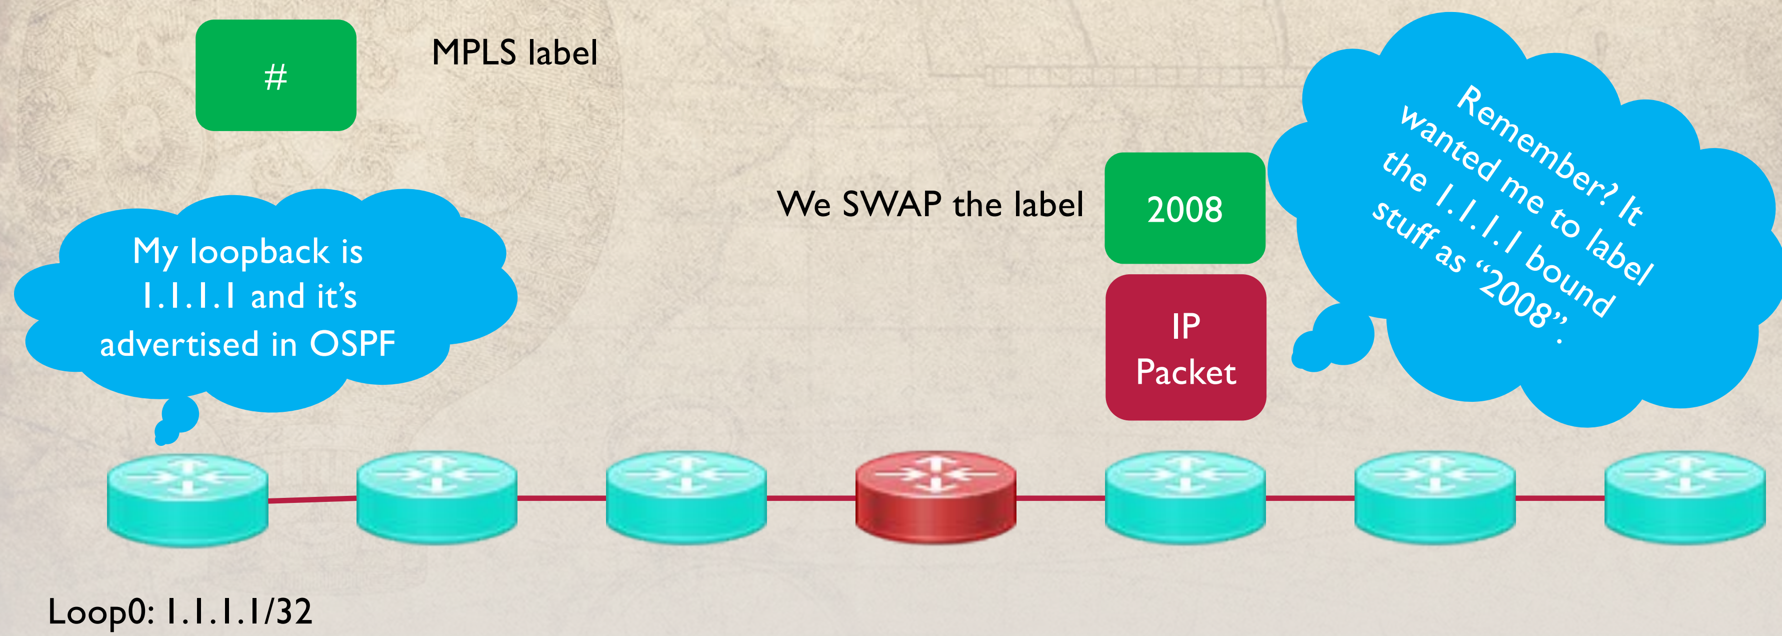

A BIT OF MPLS INTERNALS

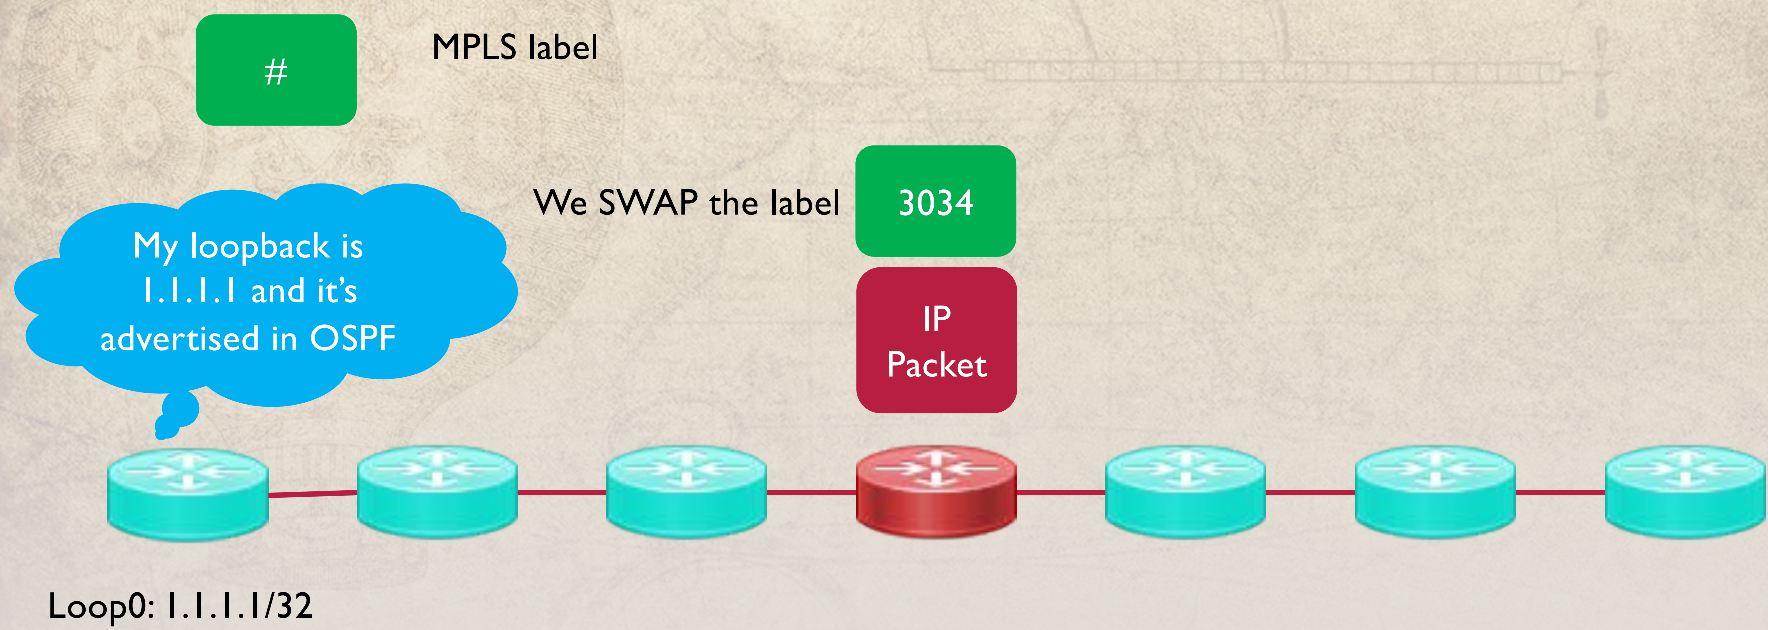

A BIT OF MPLS INTERNALS

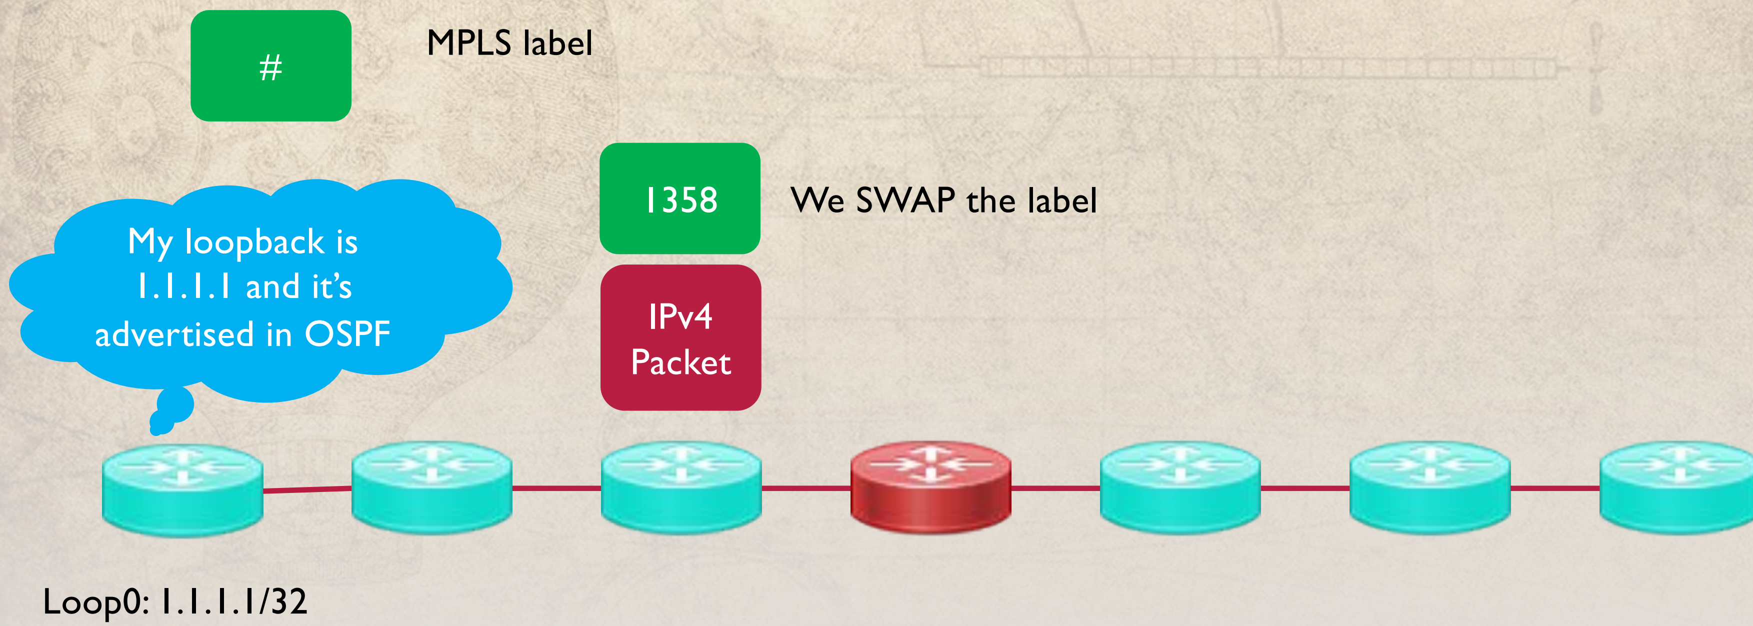

A BIT OF MPLS INTERNALS

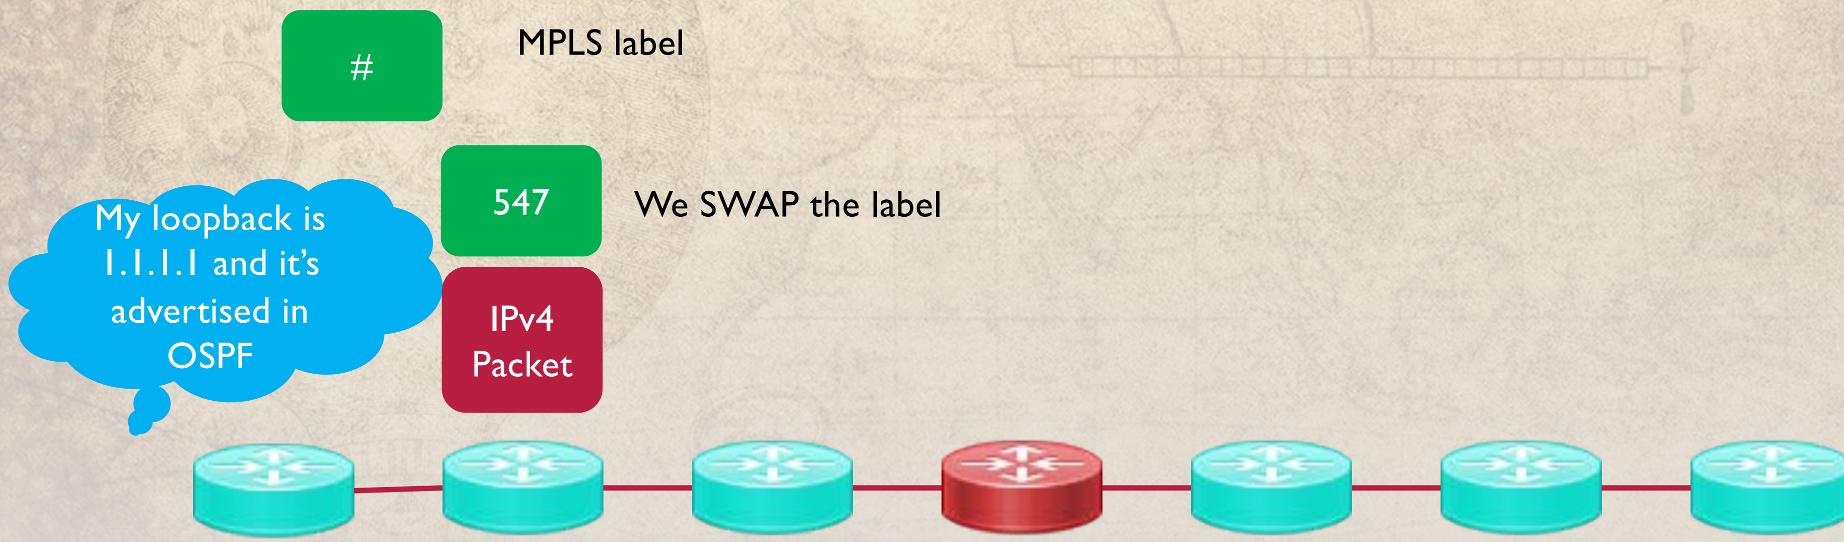

Loop0: 1.1.1.1/32

A BIT OF MPLS INTERNALS

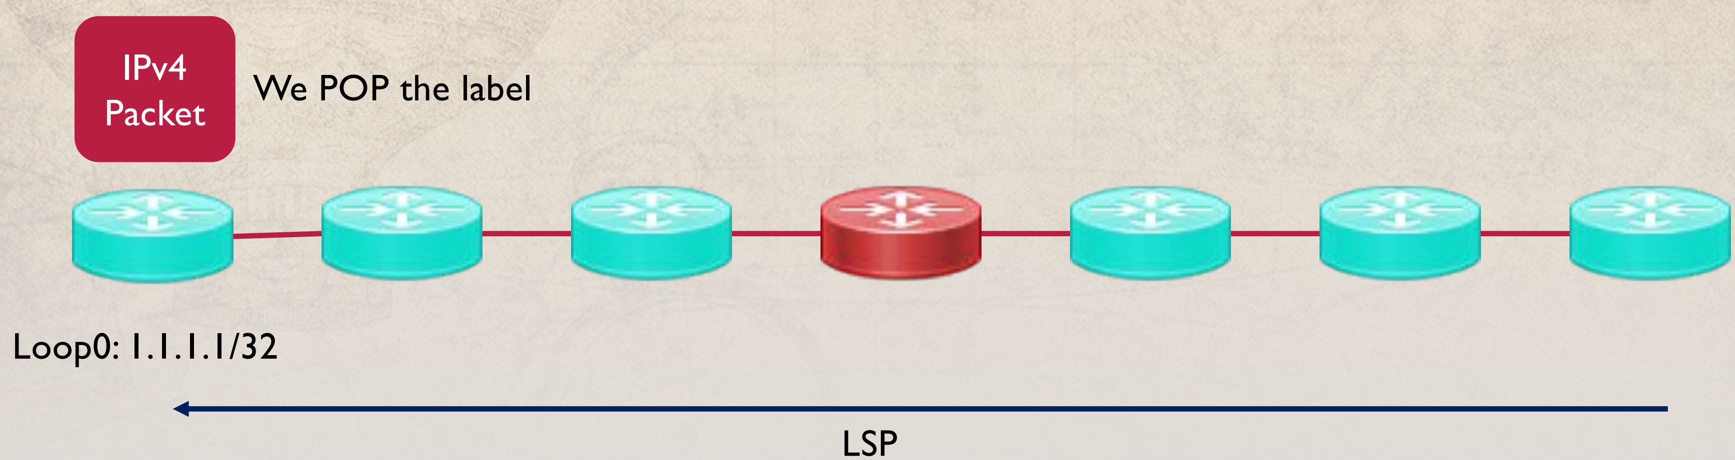

A BIT OF MPLS INTERNALS

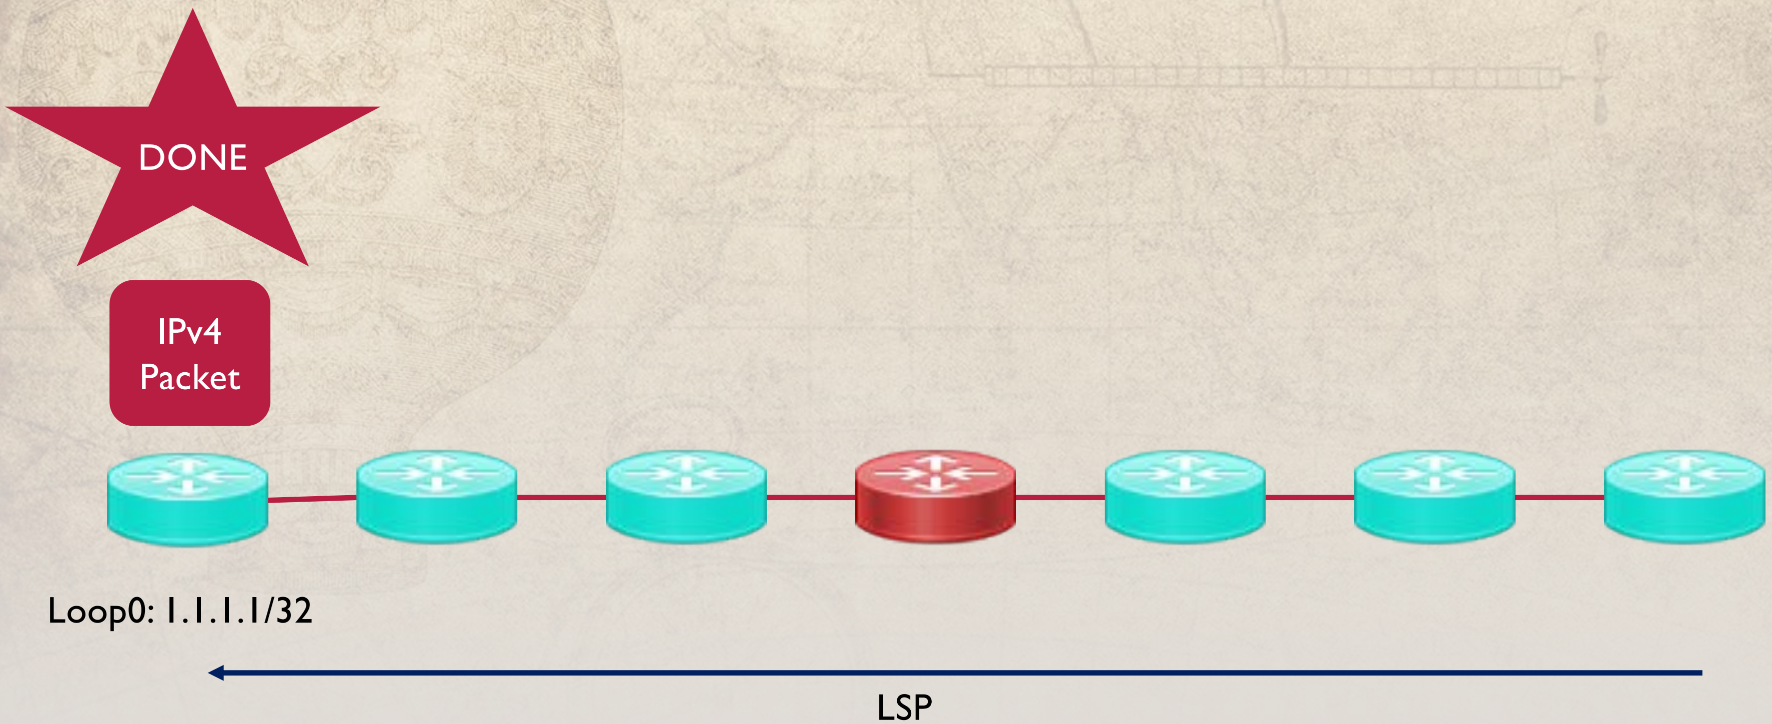

A BIT OF MPLS INTERNALS

- It does sound like a tunnel but its technically an LSP.
  - LSP: label-Switched Path
  - LSP's are one way

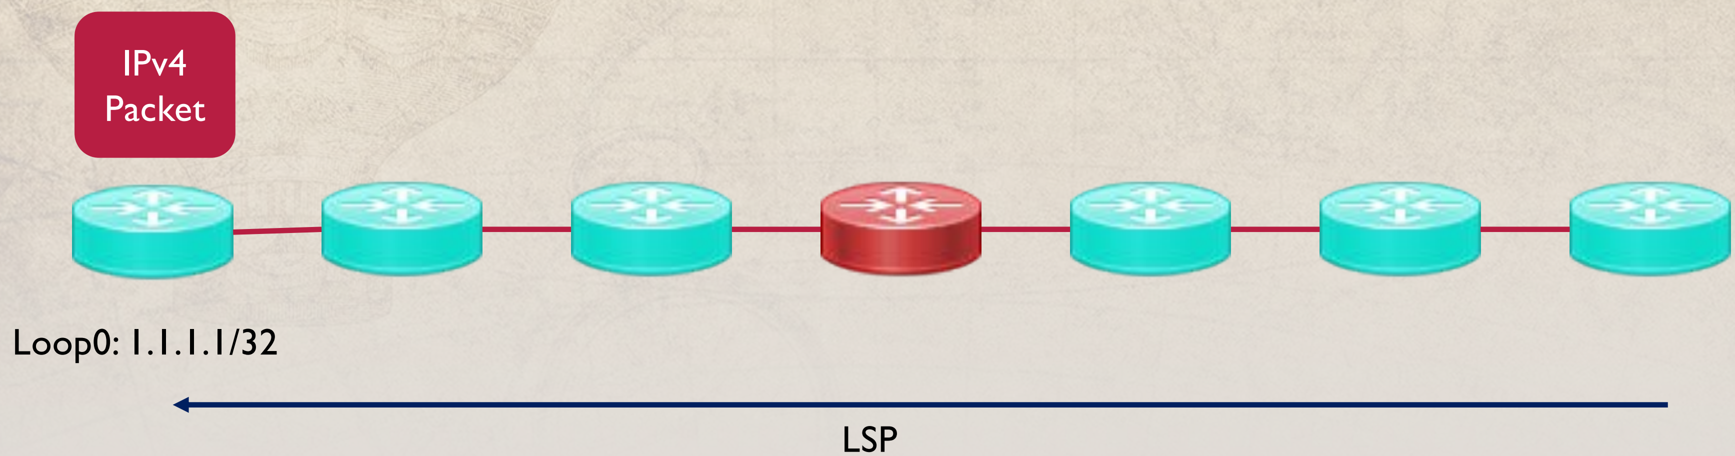

A BIT OF MPLS INTERNALS

Wasn't it cool?

We didn't look at the IP header and payload.
It could have been anything: IPv4 packet, IPv6 packet, even bunch of beans, nails, coins etc.
All we cared about, looked at and processed were the labels

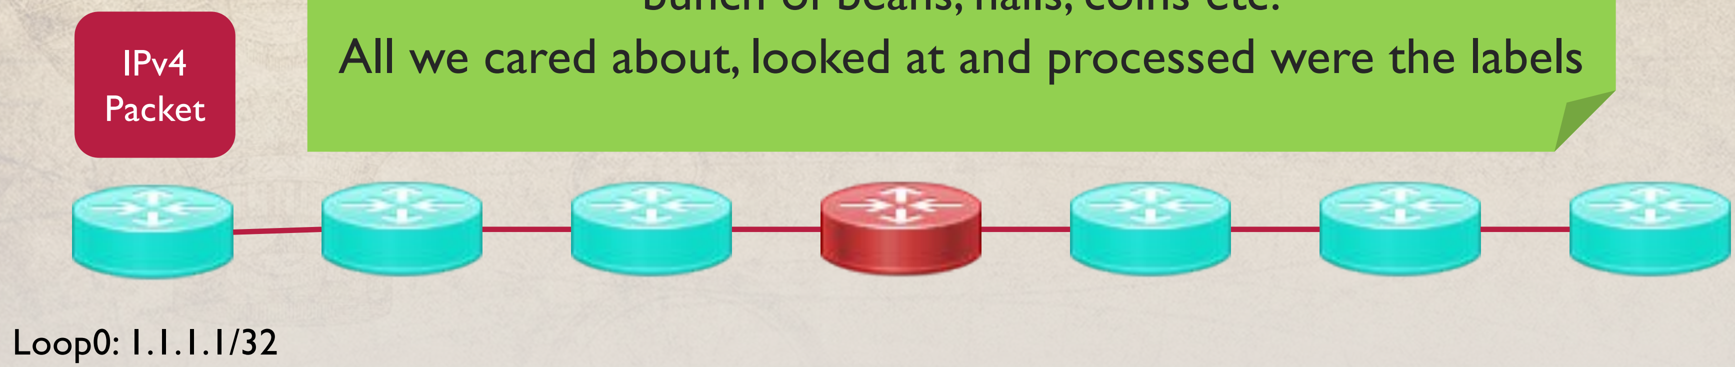

WHAT TO CALL IT?

- The label operations
  - PUSH
  - SWAP
  - POP

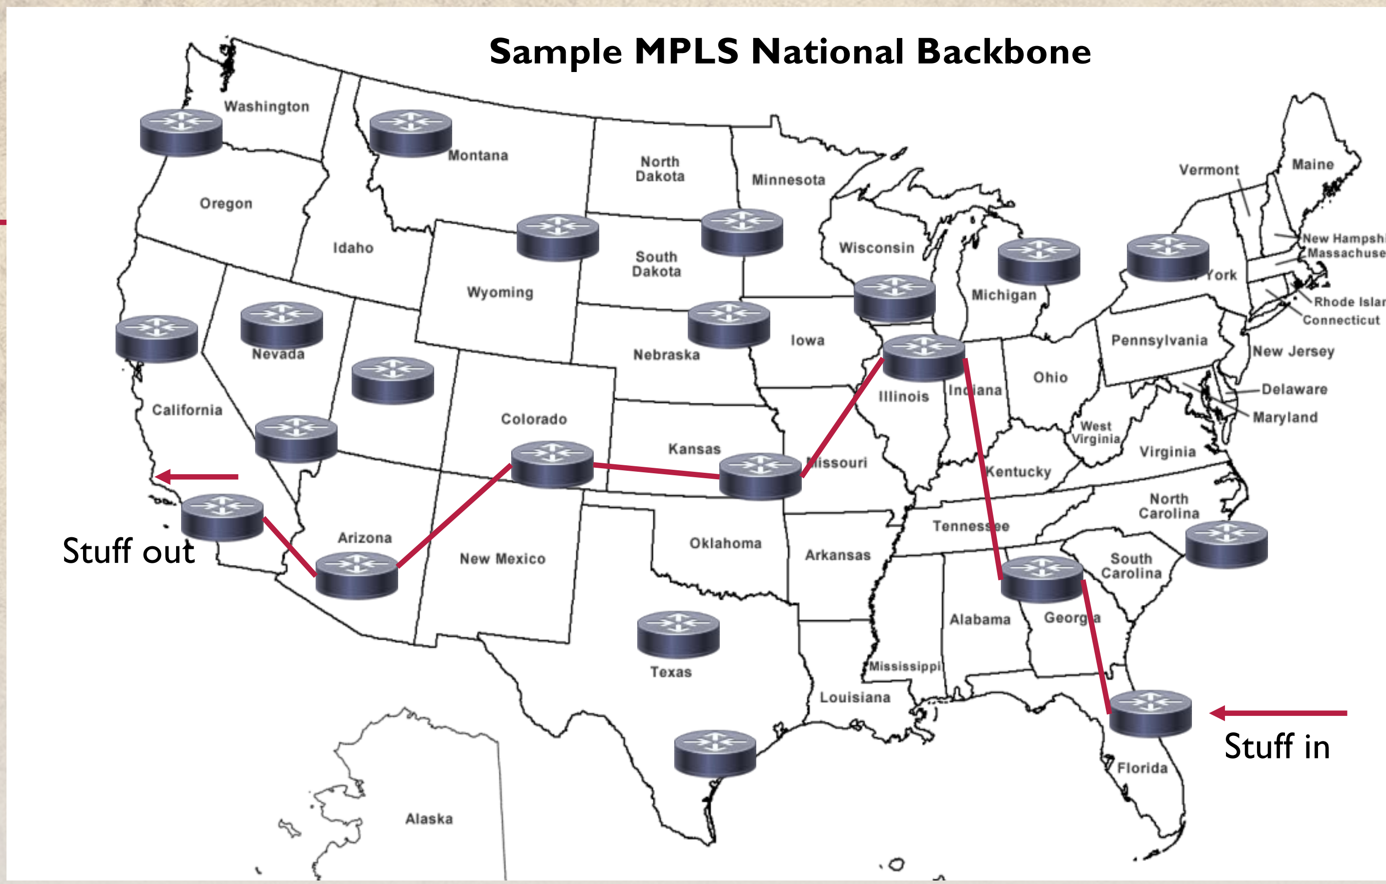

MCO: PUSH
ATL: SWAP
ORD: SWAP
MCI: SWAP
DEN: SWAP
PHX: POP
LAX: Receive

CHECK POINT

- Any questions on why we might need MPLS?
- We know we have a problem and a potential solution

ALL IN ALL

- Each router assigns labels to its reachable destinations
- Going by the passport example we care about sources and destinations
- Sources and destinations are usually edge (aka PE devices)
- Such addresses are usually loopbacks advertised in IGP

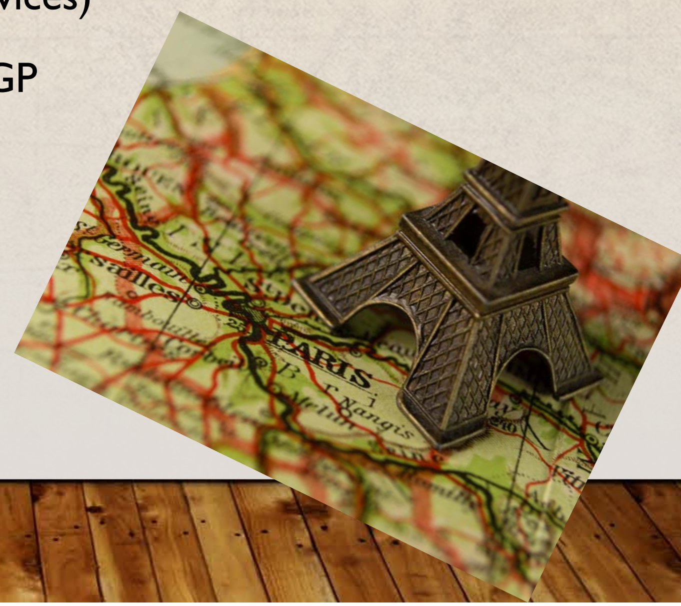

LET'S INTRODUCE MPLS

- Seems like MPLS can solve our problem(s) here

LET'S INTRODUCE MPLS

- Simple configuration to get the basic feature started
- Interface-level config for the labeling to work in regular IOS
- You can always tweak it a lot more including LDP settings

Cisco IOS snippet:

```
Rtr(config-if)#int g0/0
Rtr(config-if)#mpls ip
```

MPLS SET-UP

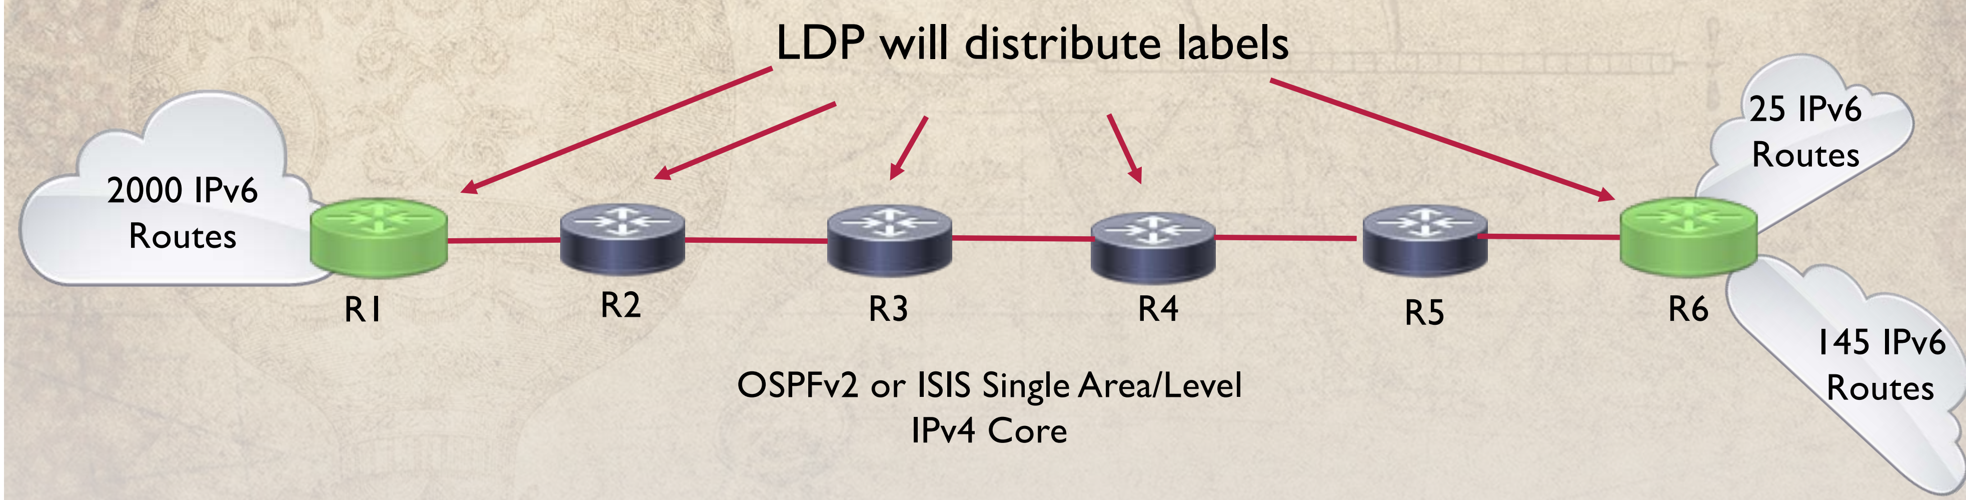

So all the routers along the path will soon tell their neighbors:

What they call the loopback addresses of R1 and R2

aka: The source and destination

THINGS LOOK GOOD…ALMOST

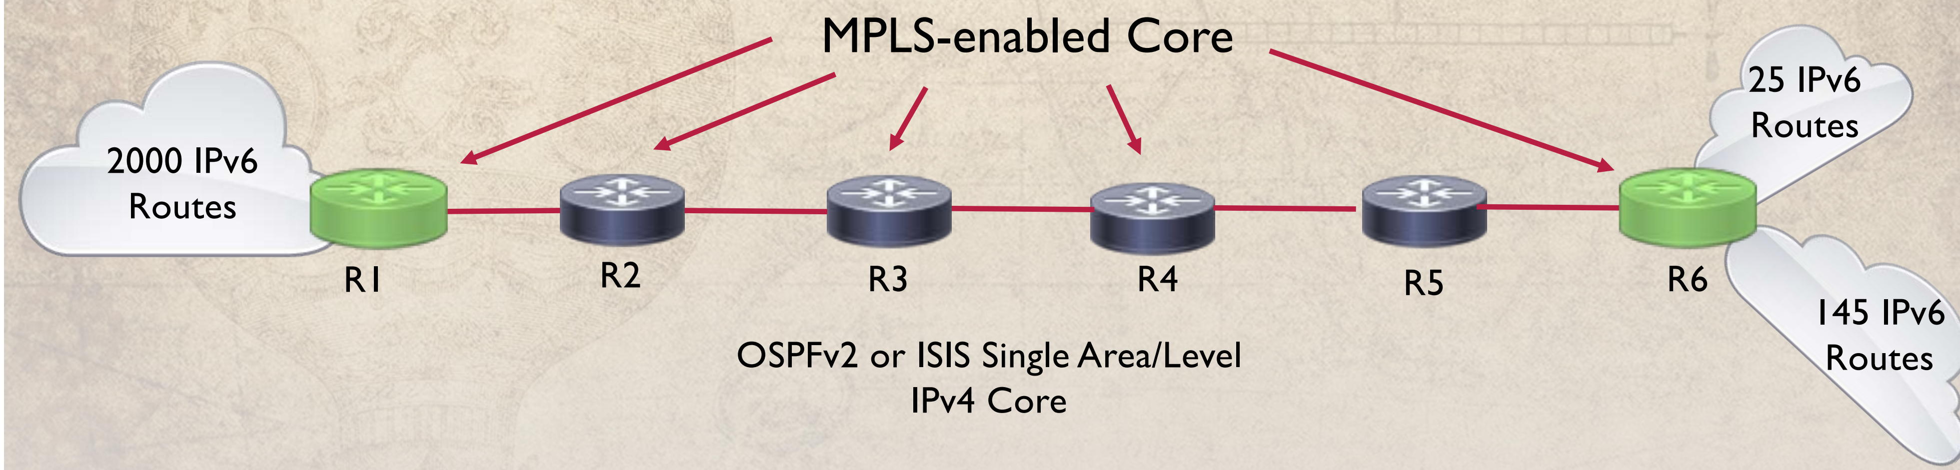

THINGS LOOK GOOD…ALMOST

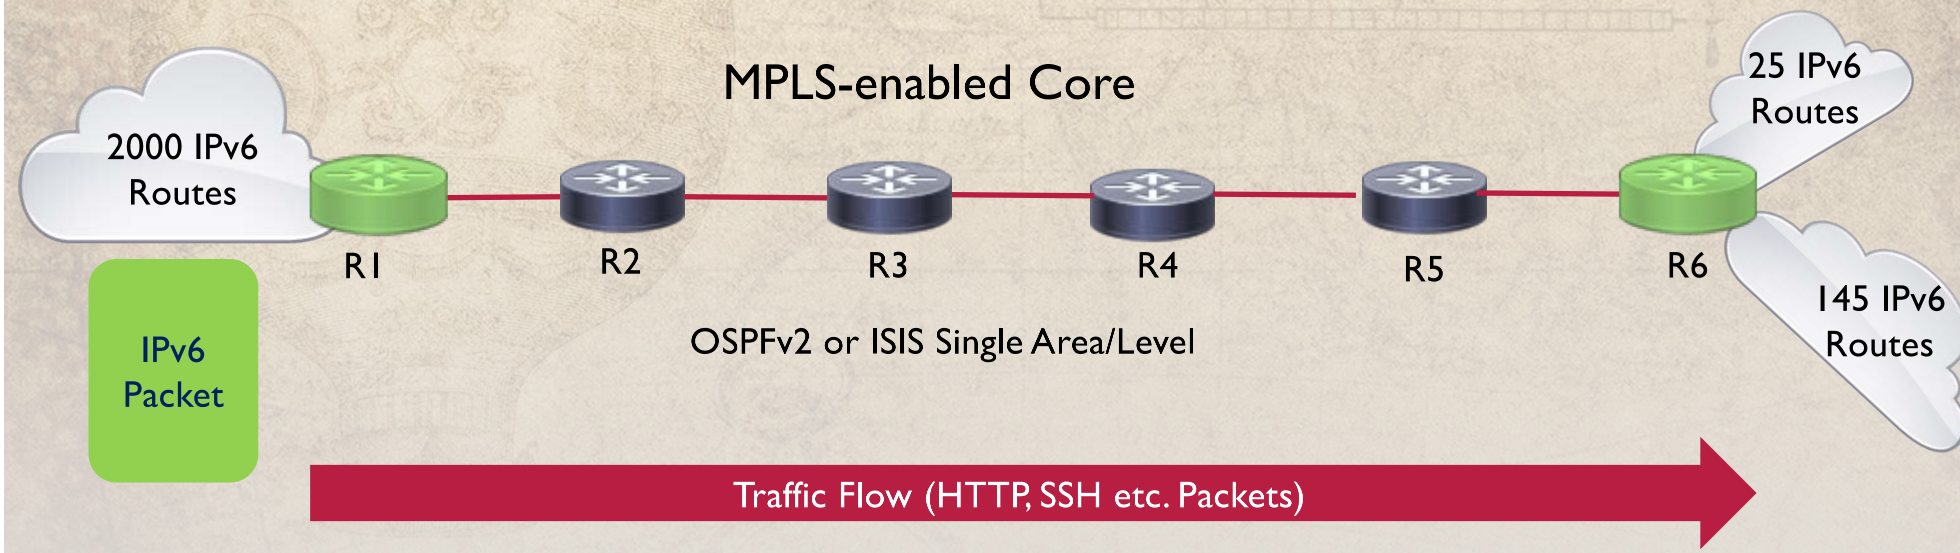

THE IPV4 NETWORK WITH IPV6 ISLANDS

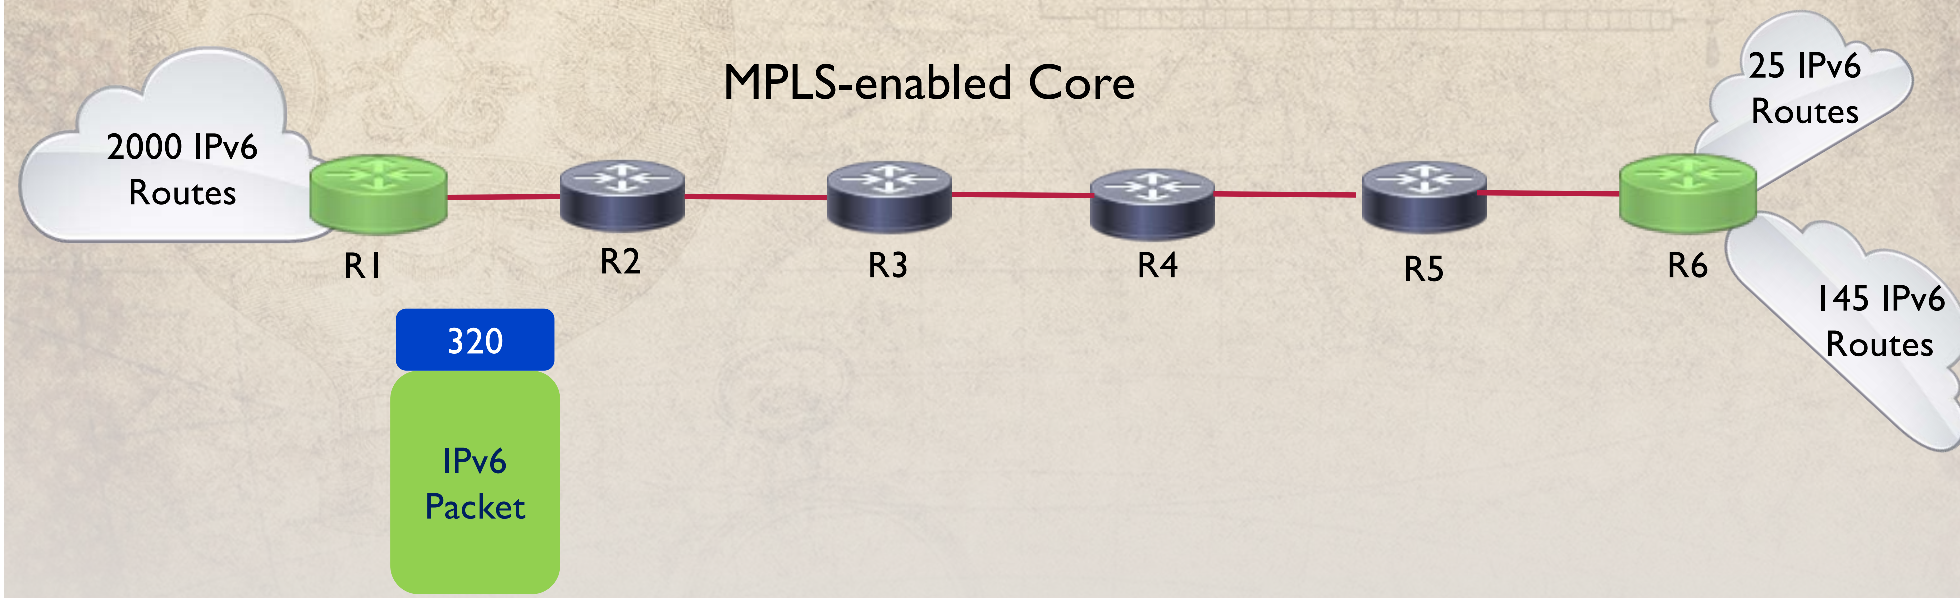

THE IPV4 NETWORK WITH IPV6 ISLANDS

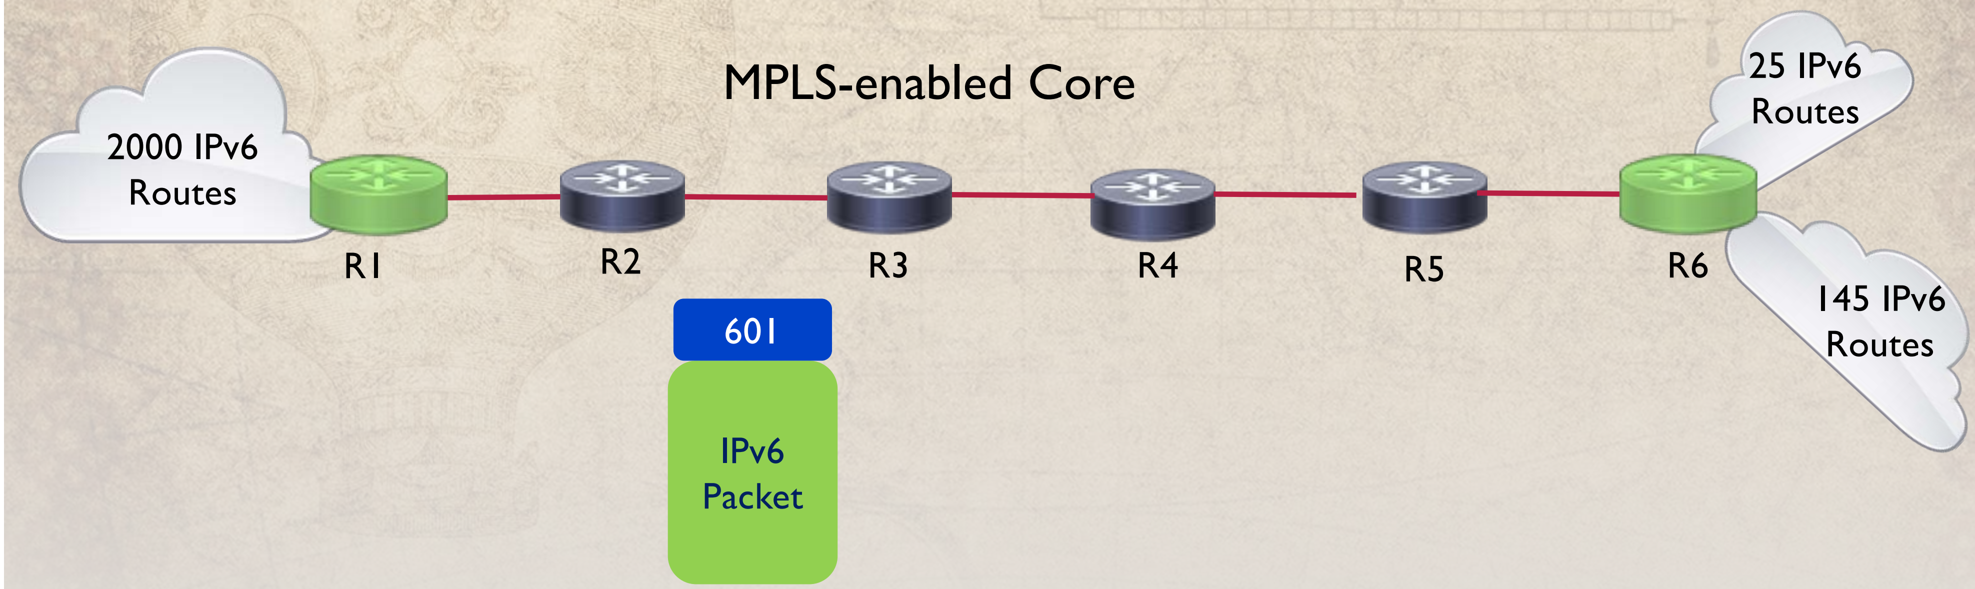

THE IPV4 NETWORK WITH IPV6 ISLANDS

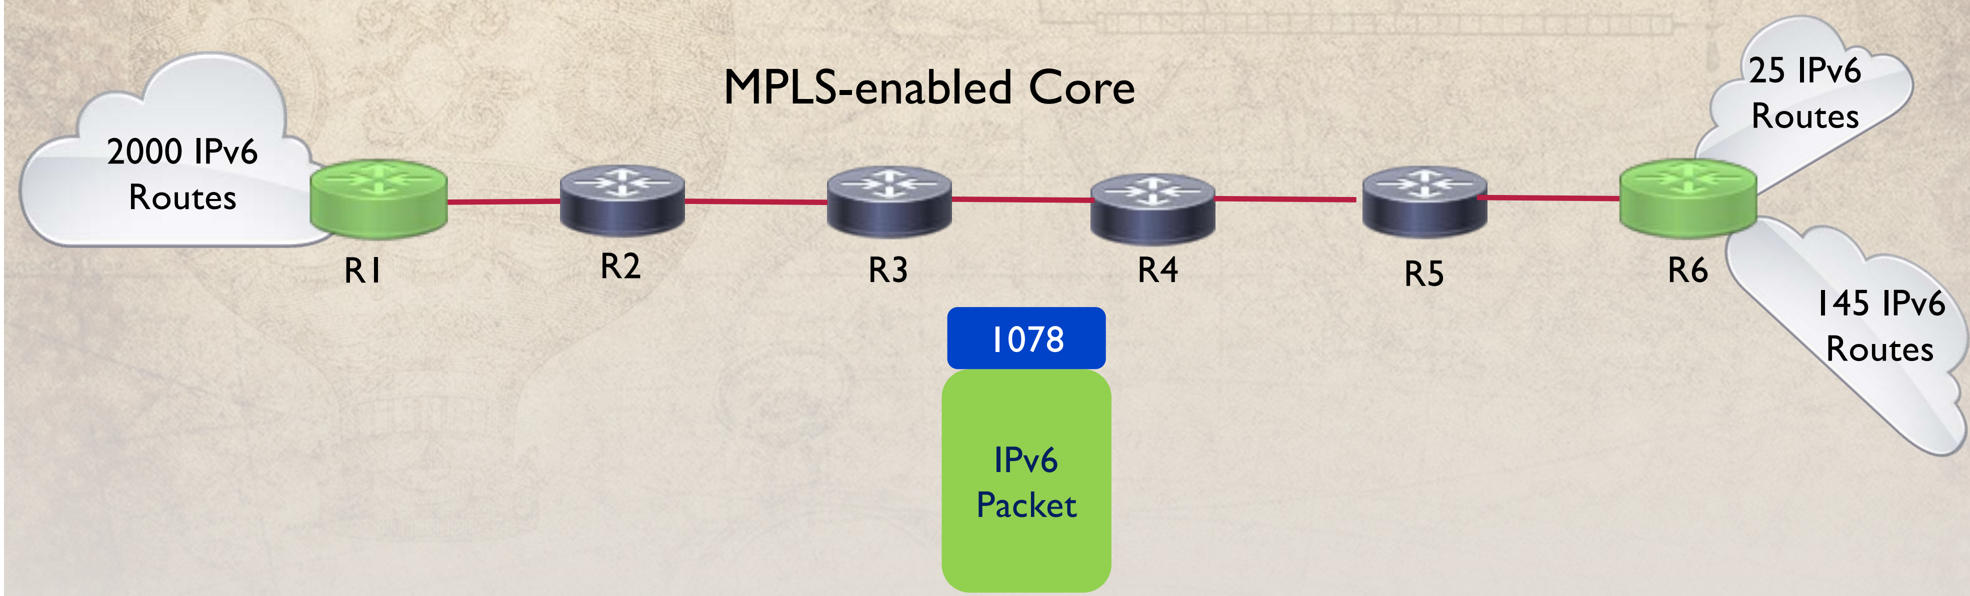

THE IPV4 NETWORK WITH IPV6 ISLANDS

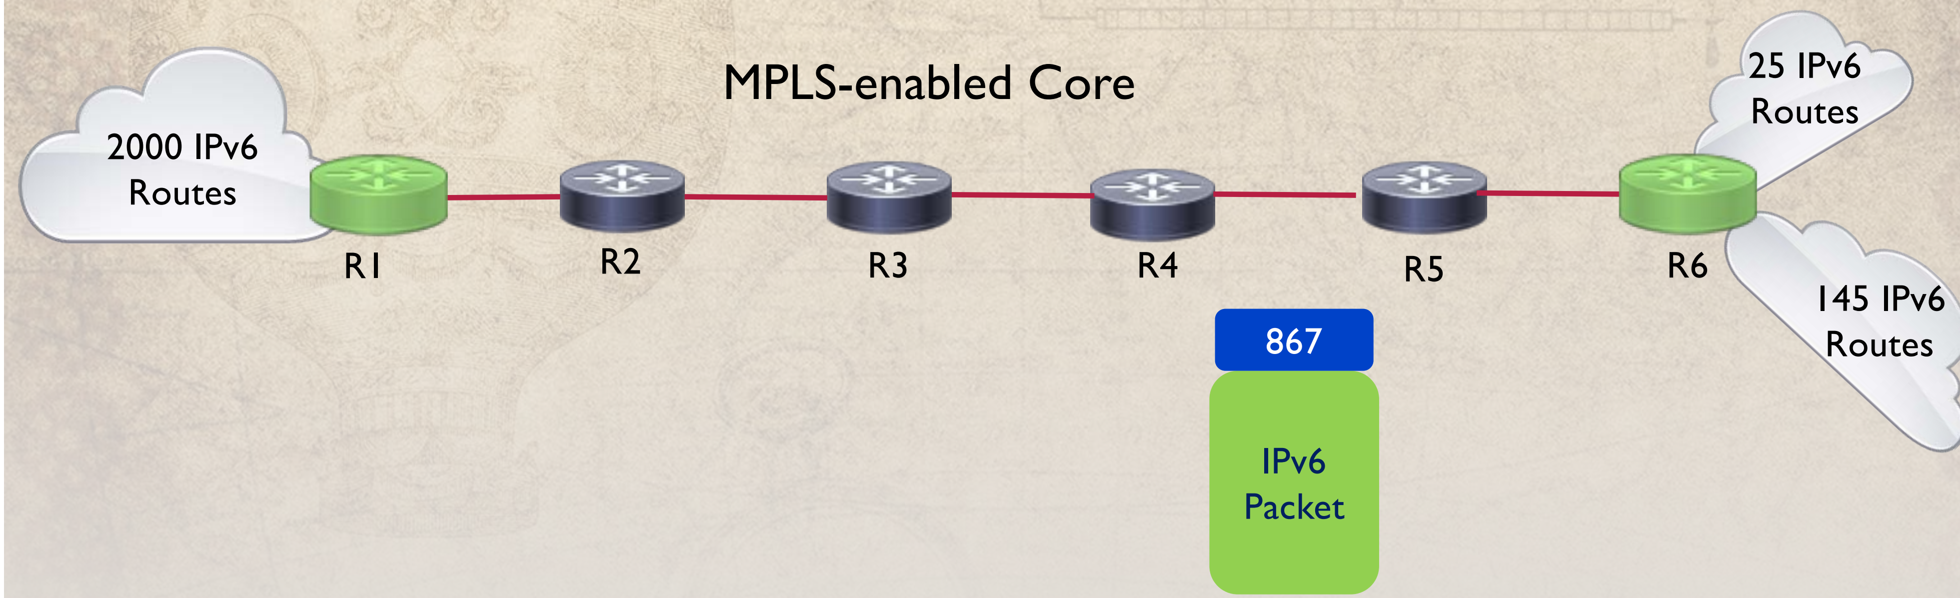

THE IPV4 NETWORK WITH IPV6 ISLANDS

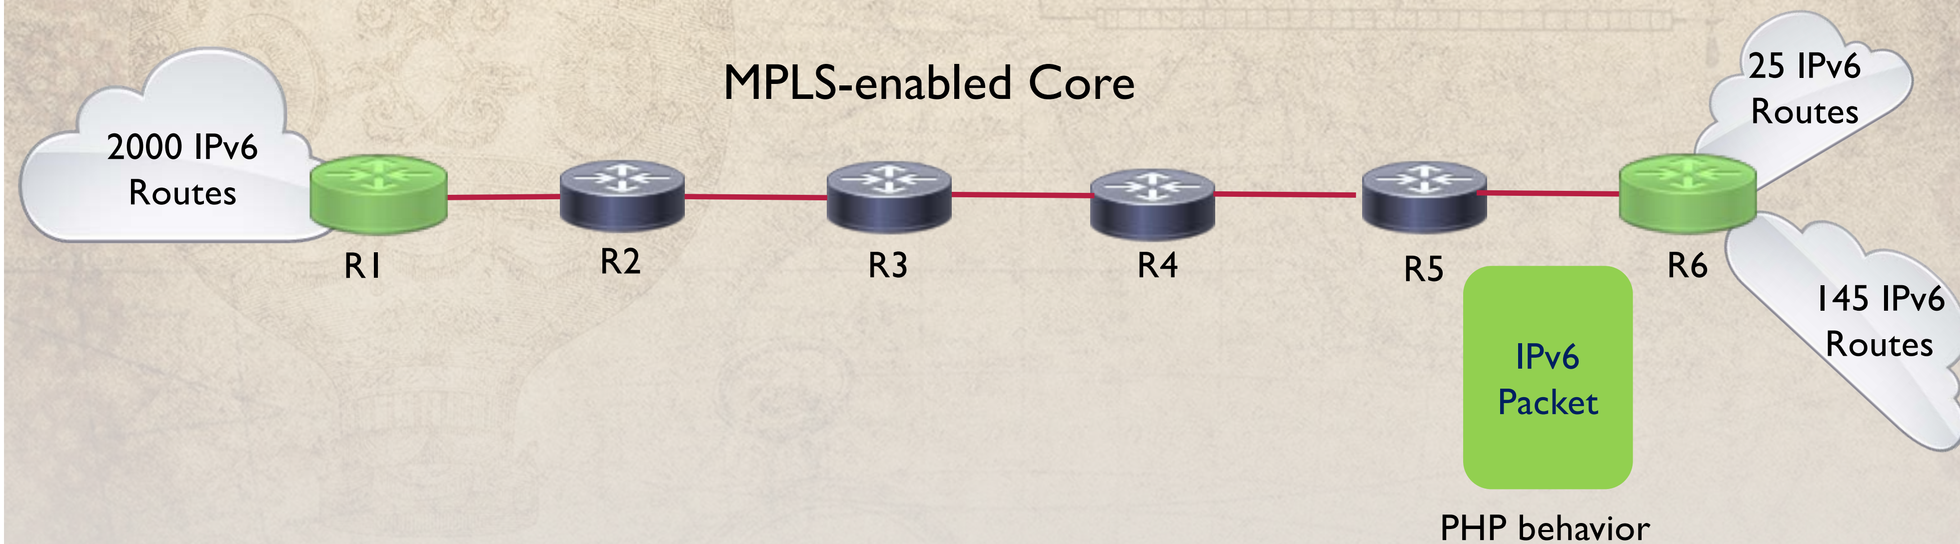

REMEMBER…

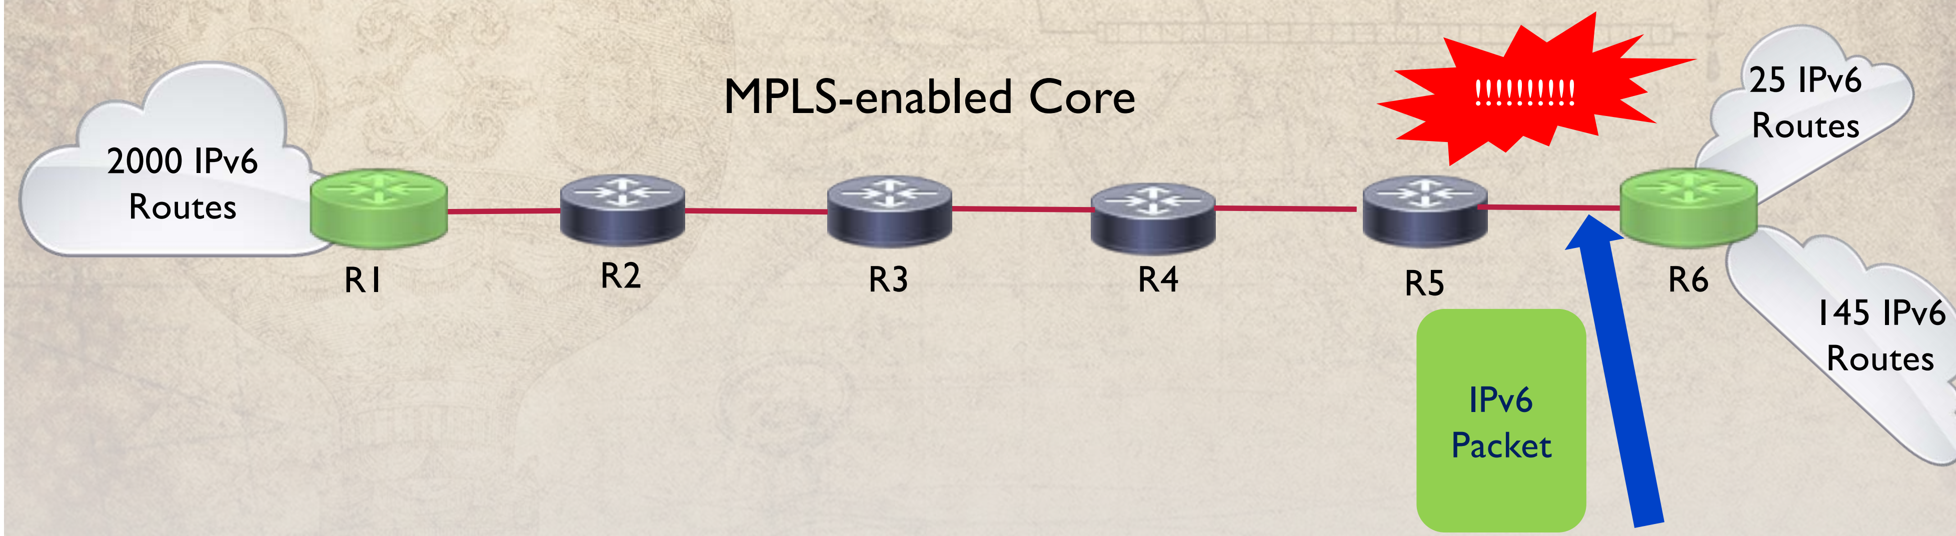

This segment is still part of the IPv4 ONLY core!

REMEMBER…

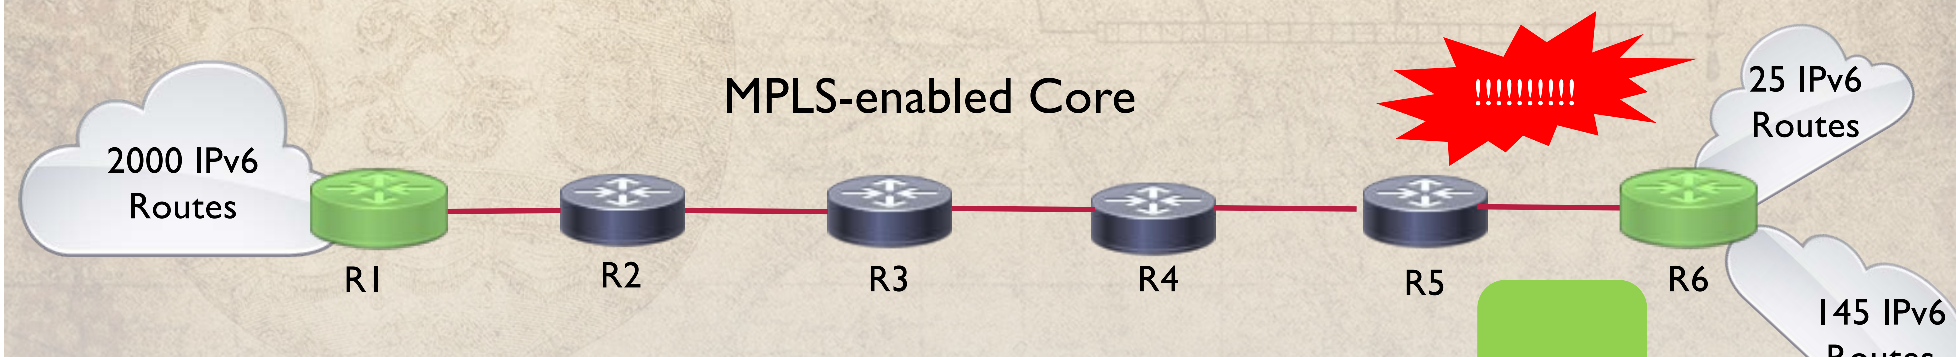

A few issues can happen including:

R5 now has to use the IPv6-specific Ethernet header (L2 header; 0x86DD) while it is not IPv6-aware and potentially does not support IPv6 at all.

LET'S FIX IT

IDEAS?

LABEL STACKING

- In most real world cases we have more than one label
- For instance:
  - **Top label:** How we pass through the transit network (the "staple" solution)
  - **Bottom label:** When packets arrive at the destination where to put it next.
    - E.g. What downstream VRF to go to?
    - E.g. What IPv6 network to pass onto?

REMEMBER?

Two MPLS labels!

```
> Frame 65: 82 bytes on wire (656 bits), 82 bytes captured (656 bits) on interface 0
> Ethernet II, Src: ca:03:0b:f8:00:08 (ca:03:0b:f8:00:08), Dst: ca:02:00:08:00:1c (ca:02:00:08:00:1c)
> MultiProtocol Label Switching Header, Label: 17, Exp: 6, S: 0, TTL: 253
> MultiProtocol Label Switching Header, Label: 25, Exp: 6, S: 1, TTL: 255
> Internet Protocol Version 6, Src: 2006::6, Dst: 2001::200
> Transmission Control Protocol, Src Port: 80, Dst Port: 60121, Seq: 1, Ack: 11, Len: 0
```

Bottom Label
Stays the same all the way

Top Label (Transport label)
Changes hop-by-hop

SMART SOLUTION: 2 LABELS

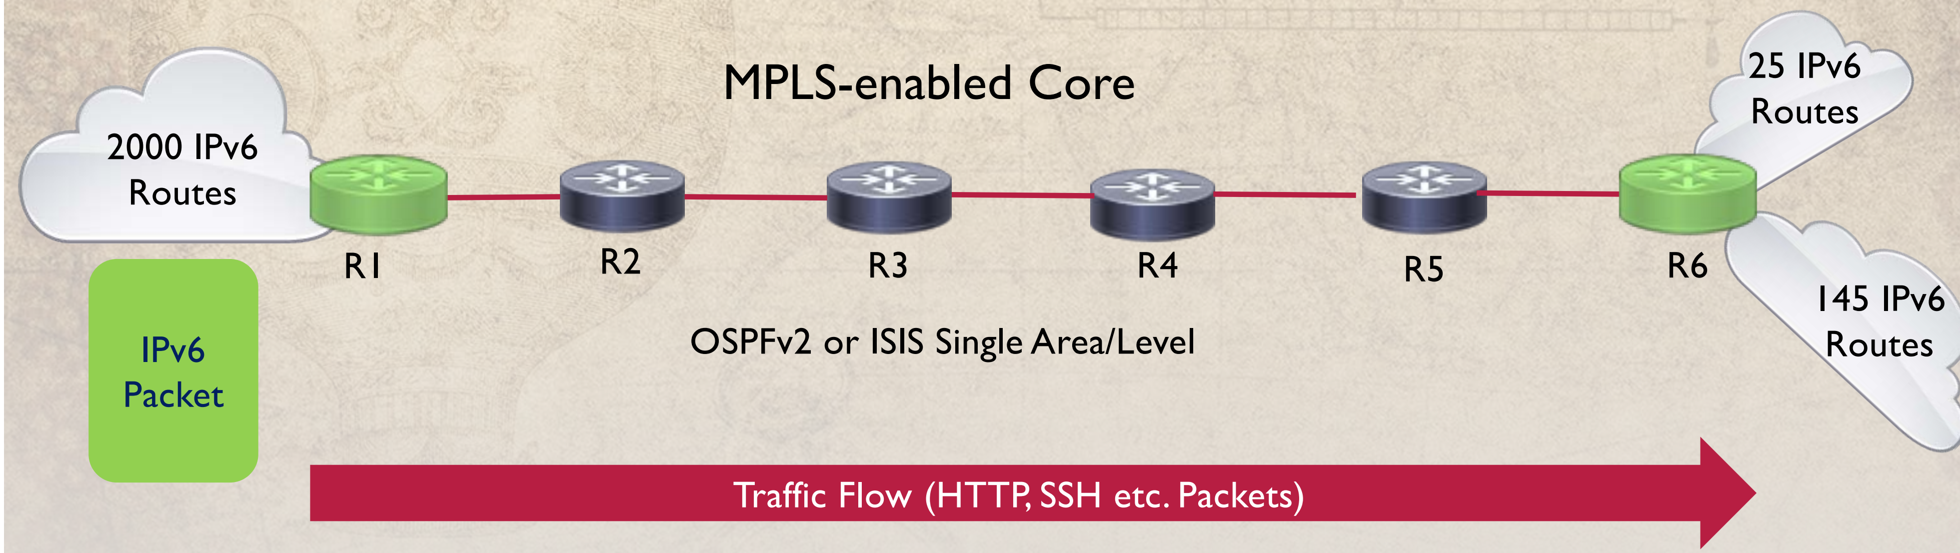

SMART SOLUTION: 2 LABELS

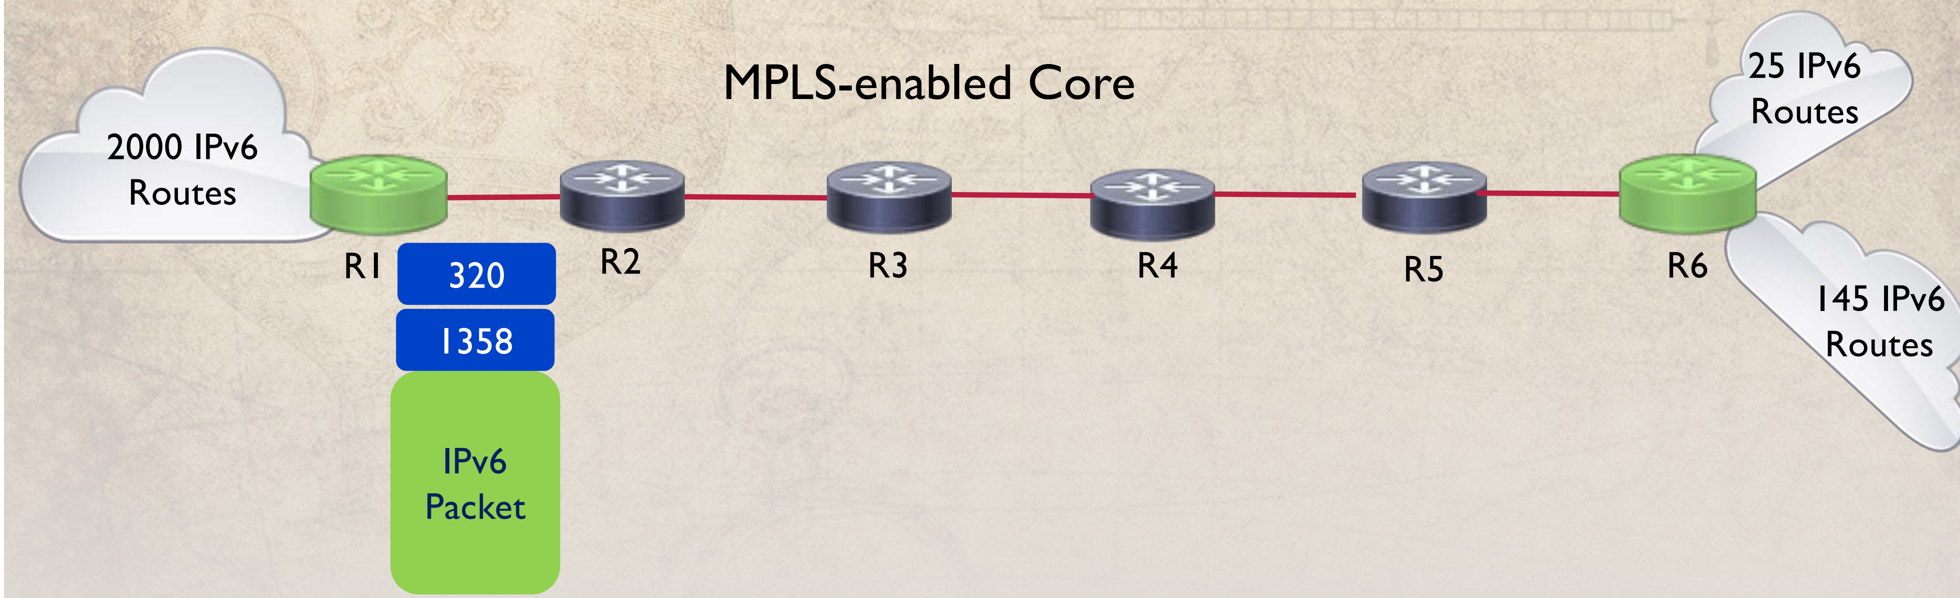

SMART SOLUTION: 2 LABELS

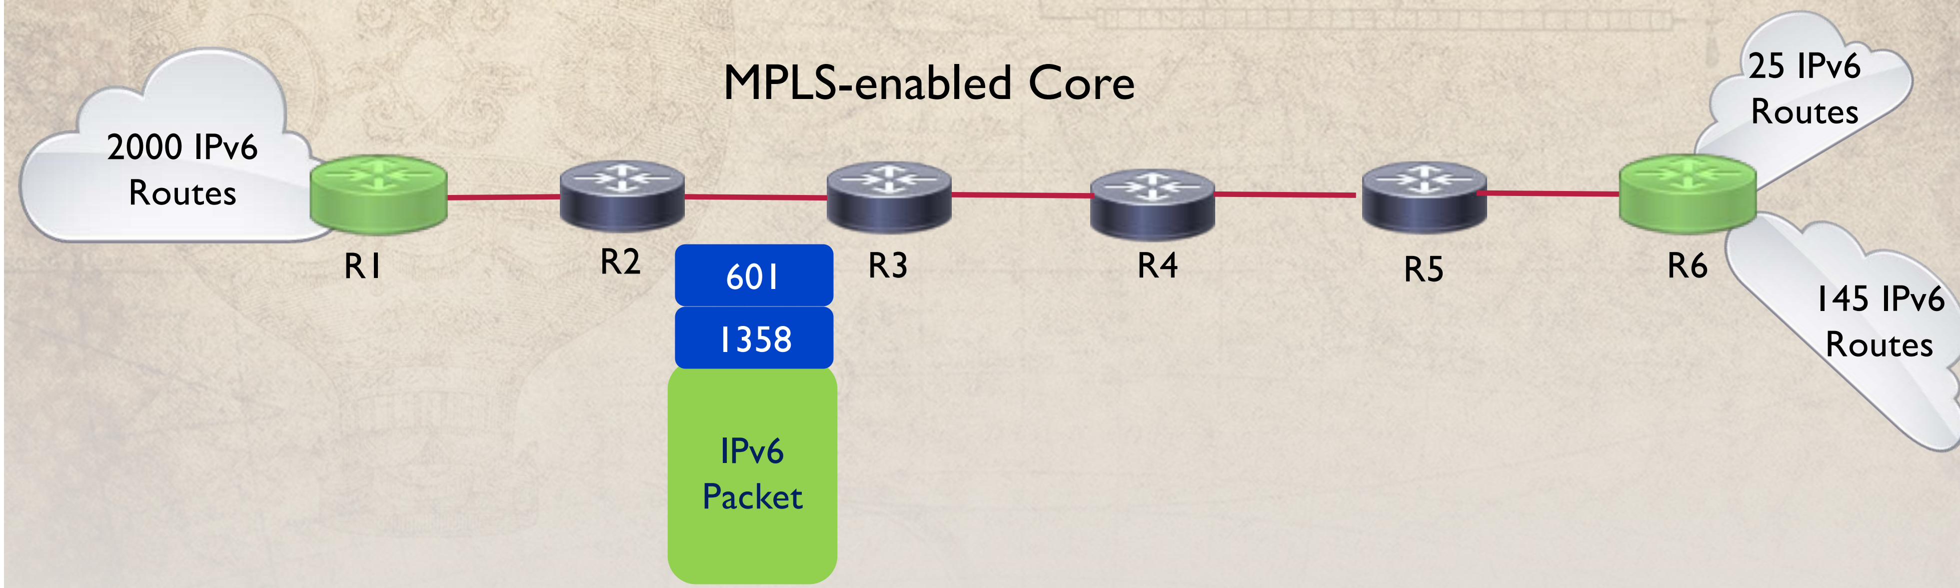

SMART SOLUTION: 2 LABELS

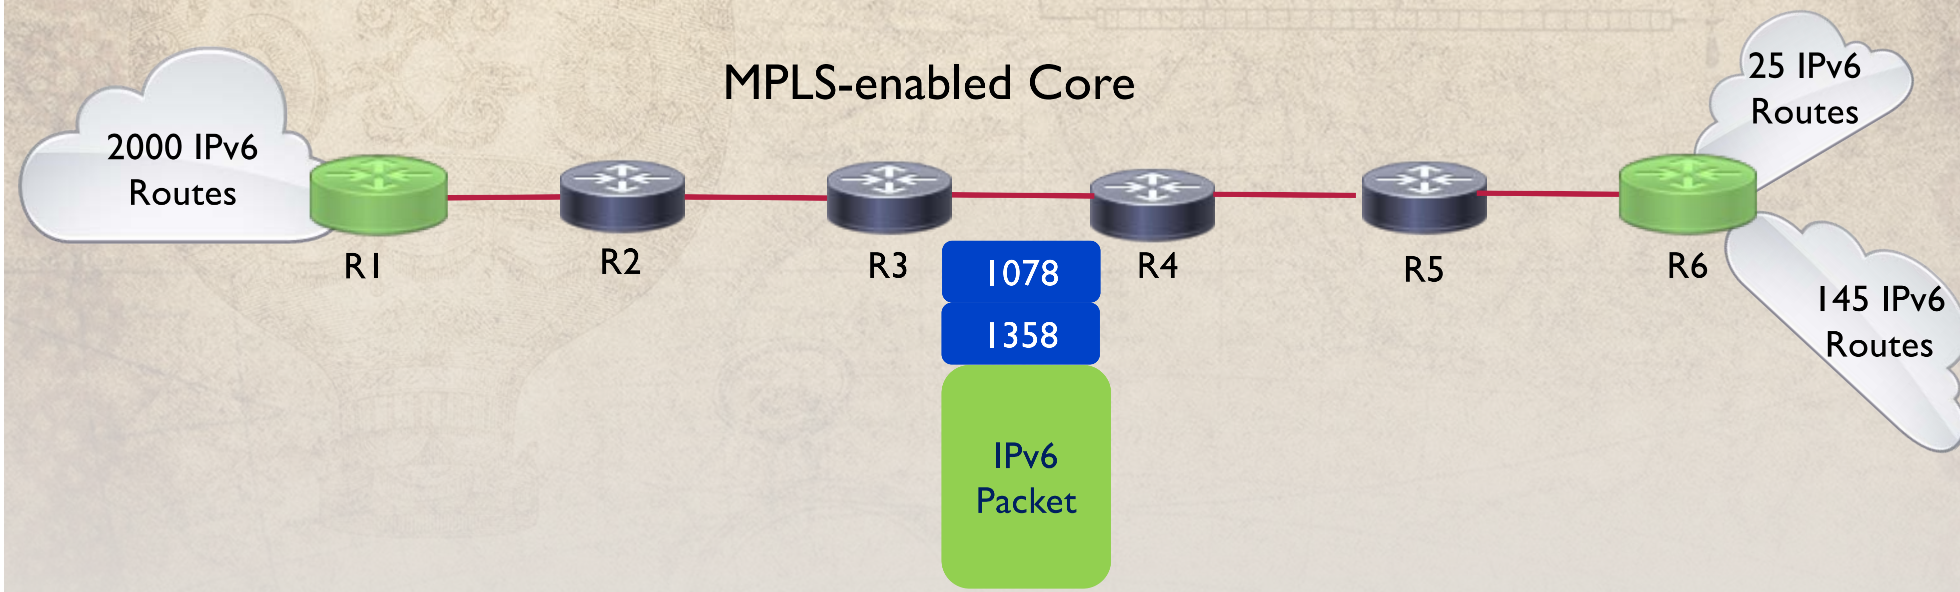

SMART SOLUTION: 2 LABELS

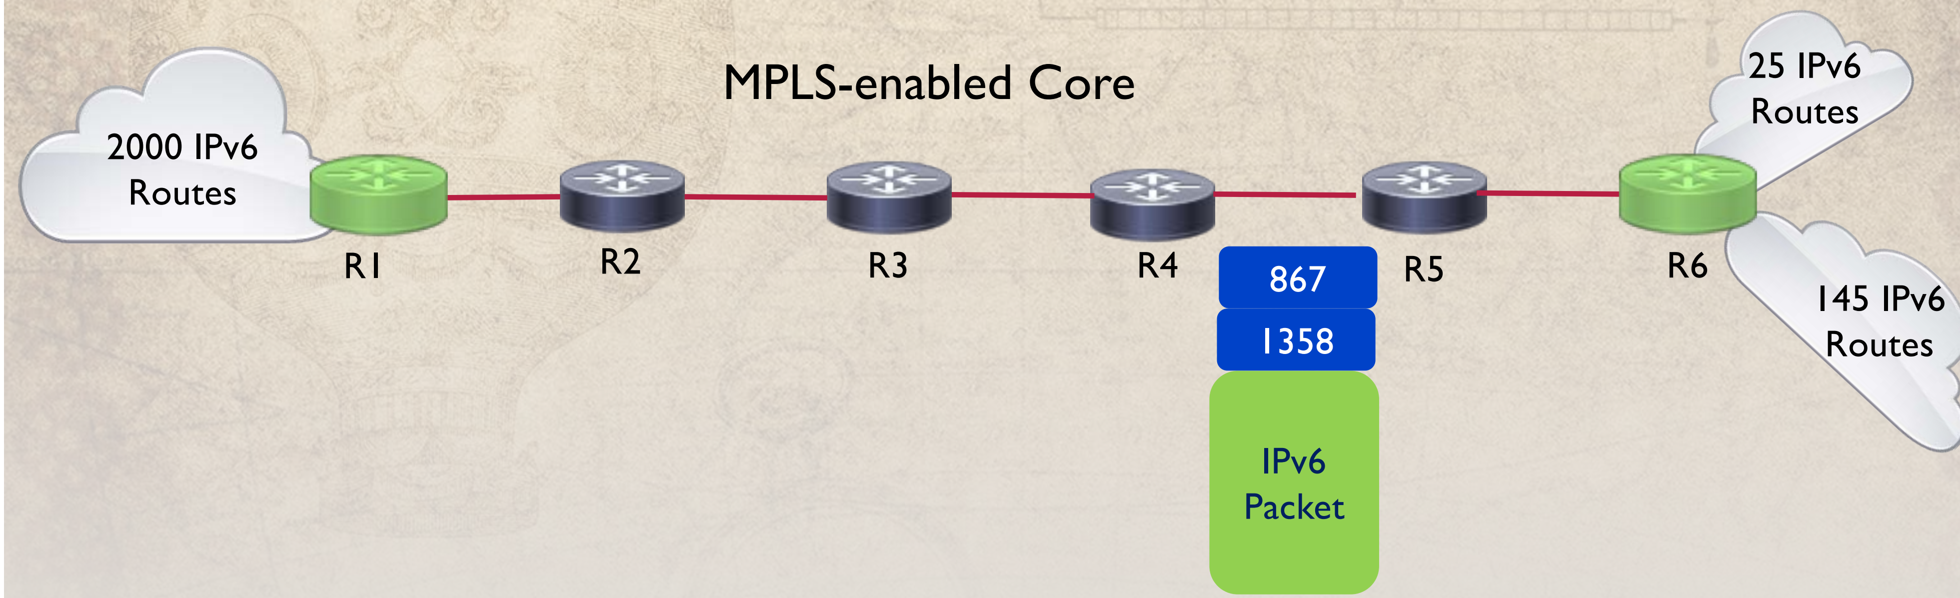

SMART SOLUTION: 2 LABELS

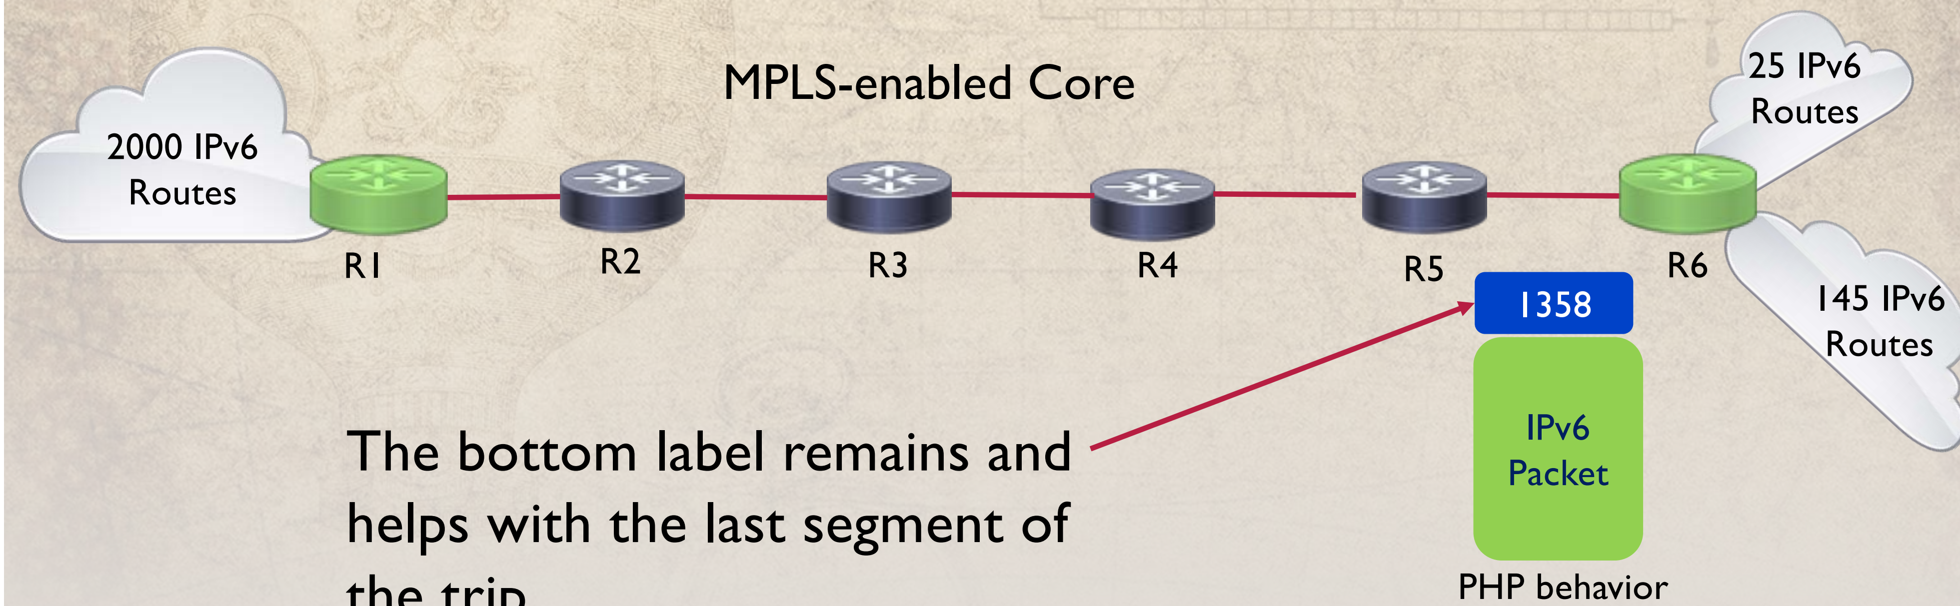

The bottom label remains and helps with the last segment of the trip.

HOW DID THE LABEL DISTRIBUTION WORK?

1- **The TOP label**: PE finder aka next-hop finder

Everybody uses LDP here. So every router knows about the label assignments of their neighbors.
They label packets accordingly when sending packets to them.

HOW DID THE LABEL DISTRIBUTION WORK?

The BOTTOM label:

- Bottom labels are assigned to IPv6 prefixes on the two ends
- They are assigned by each edge device (PE)
- But the labels have to be shared with other edges of the network

HOW DID THE LABEL DISTRIBUTION WORK?

2- **The BOTTOM label**: BGP updates can transfer MPLS labels for each IPv6 prefix from one edge to another

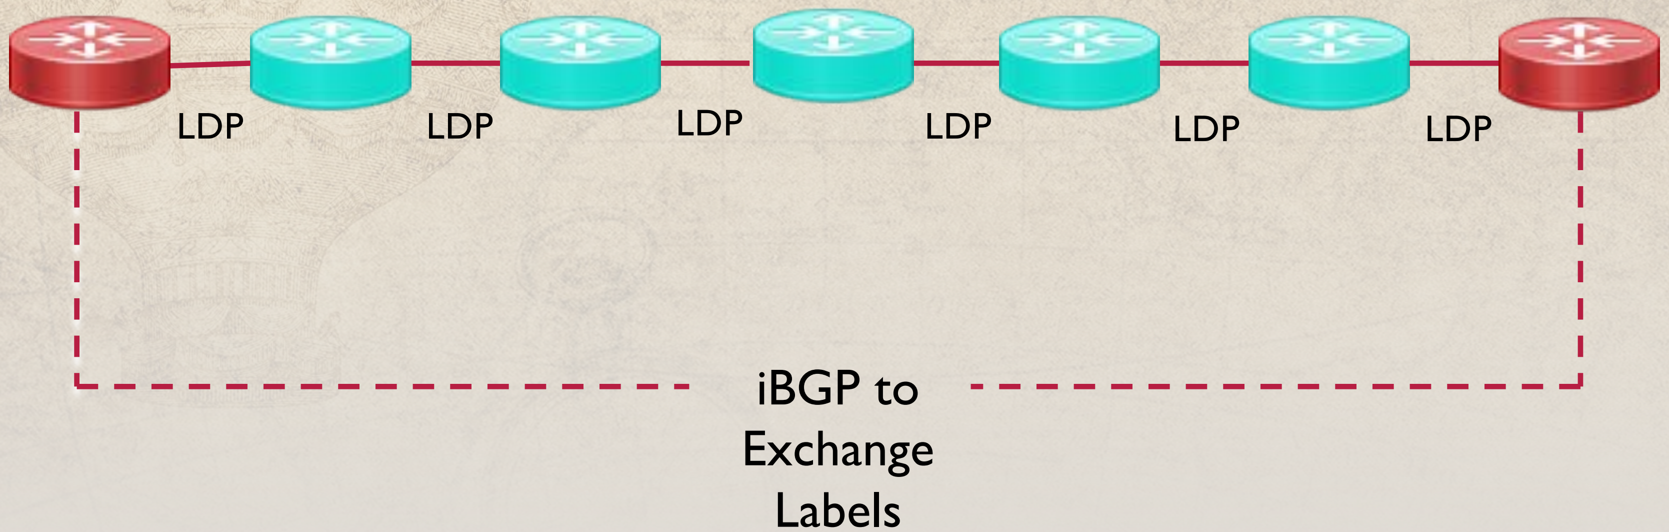

HOW DID THE LABEL DISTRIBUTION WORK?

Sample configuration in regular IOS on a Cisco router

PE 1 – R1

```
router bgp 100
 neighbor 6.6.6.6 remote-as 100
 neighbor 6.6.6.6 update-source Loopback0
  !
 address-family ipv6
  neighbor 6.6.6.6 activate
  neighbor 6.6.6.6 send-label
 exit-address-family
```

PE 2 – R6

```
router bgp 100
 neighbor 1.1.1.1 remote-as 100
 neighbor 1.1.1.1 update-source Loopback0
  !
 address-family ipv6
  neighbor 1.1.1.1 activate
  neighbor 1.1.1.1 send-label
 exit-address-family
```

LET'S ADD AN IPV6 PREFIX TO AN ISLAND

Sample configuration in regular IOS on a Cisco router

PE1 – R1

```
router bgp 100
 neighbor 6.6.6.6 remote-as 100
 neighbor 6.6.6.6 update-source
Loopback0
  !
 address-family ipv6
  neighbor 6.6.6.6 activate
  neighbor 6.6.6.6 send-label
  network 2031::31/128
exit-address-family
```

LET'S ADD AN IPV6 PREFIX TO AN ISLAND

Sample configuration in regular IOS on a Cisco router

PE 1 – R1

```
router bgp 100
 neighbor 6.6.6.6 remote-as 100
 neighbor 6.6.6.6 update-source
Loopback0
  !
 address-family ipv6
  neighbor 6.6.6.6 activate
  neighbor 6.6.6.6 send-label
  network 2031::31/128
exit-address-family
```

- Now PE1 R1 tells the other PE what label it's decided to assign to 2031::31/128.
- **Let's say 18**

BGP UPDATE – BOTTOM LABEL TRANSFER

```
> Frame 1223: 146 bytes on wire (1168 bits), 146 bytes captured (1168 bits) on interface 0
> Ethernet II, Src: ca:00:10:94:00:08 (ca:00:10:94:00:08), Dst: ca:01:2f:14:00:08 (ca:01:2f:14:00:08)
> MultiProtocol Label Switching Header, Label: 404, Exp: 6, S: 1, TTL: 255
> Internet Protocol Version 4, Src: 1.1.1.1, Dst: 6.6.6.6
> Transmission Control Protocol, Src Port: 179, Dst Port: 20219, Seq: 970, Ack: 503, Len: 88
v Border Gateway Protocol - UPDATE Message
    Marker: ffffffffffffffffffffffffffffffff
    Length: 88
    Type: UPDATE Message (2)
    Withdrawn Routes Length: 0
    Total Path Attribute Length: 65
  v Path attributes
    > Path Attribute - ORIGIN: IGP
    > Path Attribute - AS_PATH: empty
    > Path Attribute - MULTI_EXIT_DISC: 0
    > Path Attribute - LOCAL_PREF: 100
    v Path Attribute - MP_REACH_NLRI
      > Flags: 0x80, Optional, Non-transitive, Complete
        Type Code: MP_REACH_NLRI (14)
        Length: 41
        Address family identifier (AFI): IPv6 (2)
        Subsequent address family identifier (SAFI): Labeled Unicast (4)
      > Next hop network address (16 bytes)
        Number of Subnetwork points of attachment (SNPA): 0
      v Network layer reachability information (20 bytes)
          Label Stack=18 (bottom), IPv6=2031::31/128
```

LET'S REVIEW THE FLOW ONE LAST TIME

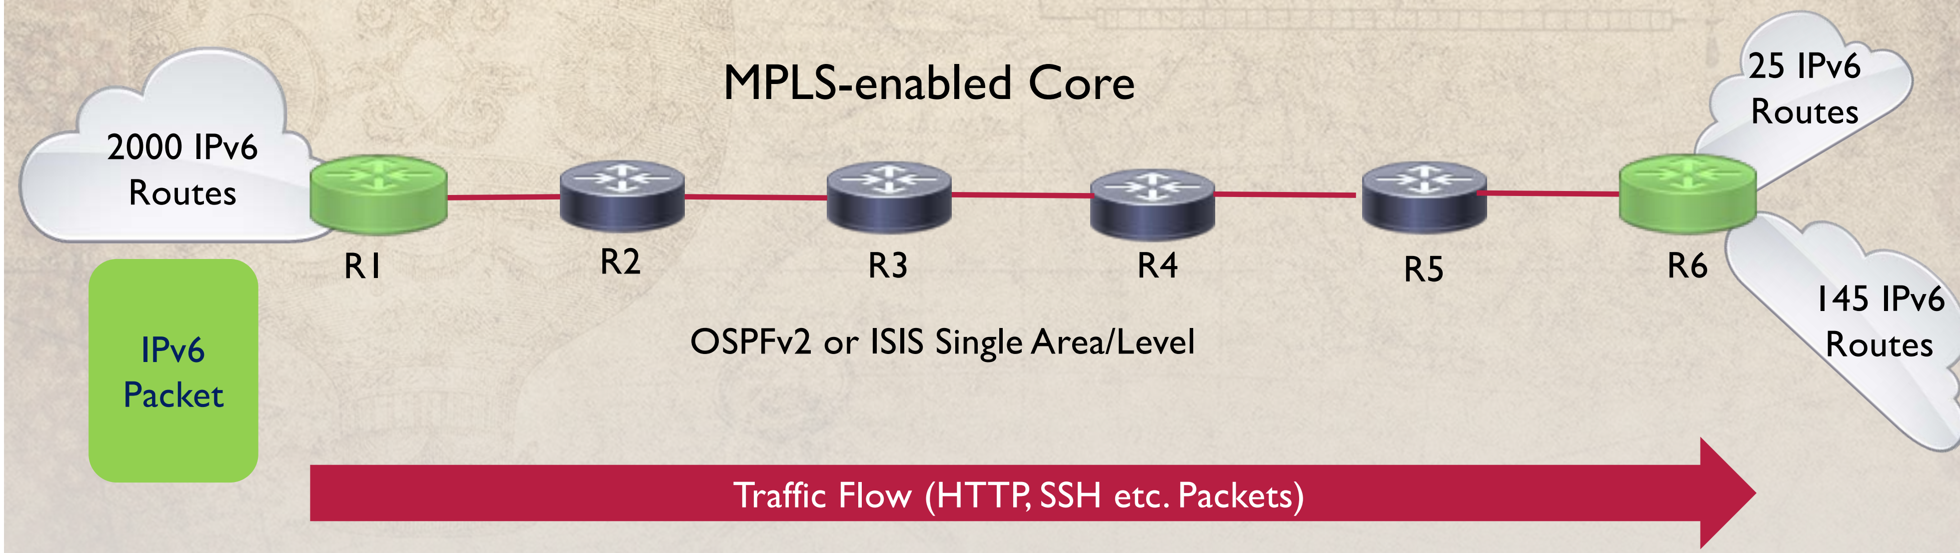

LET'S REVIEW THE FLOW ONE LAST TIME

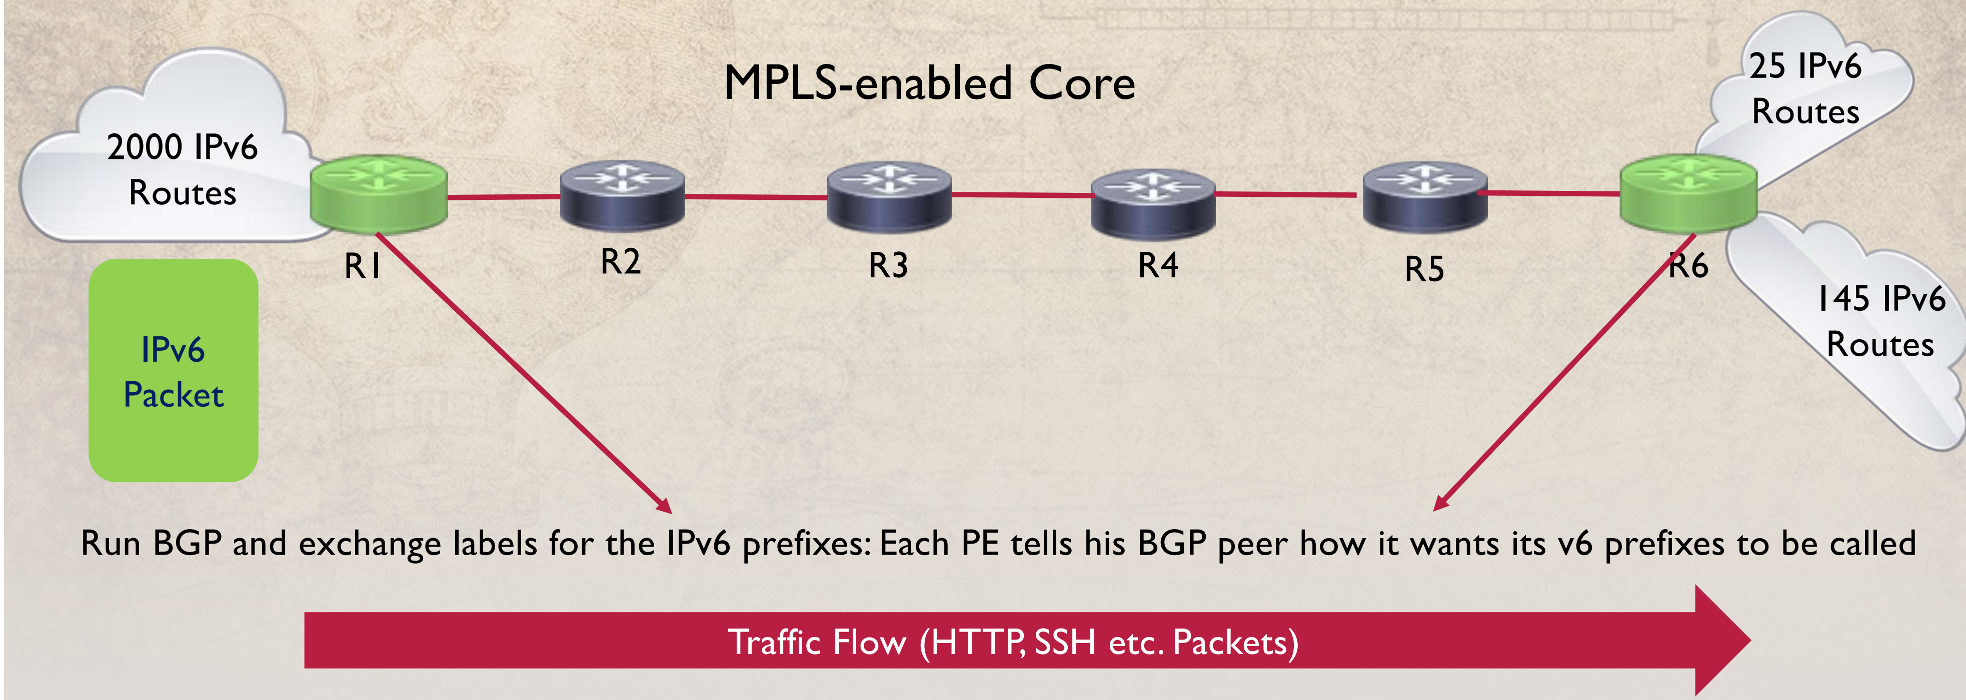

END TO END FLOW

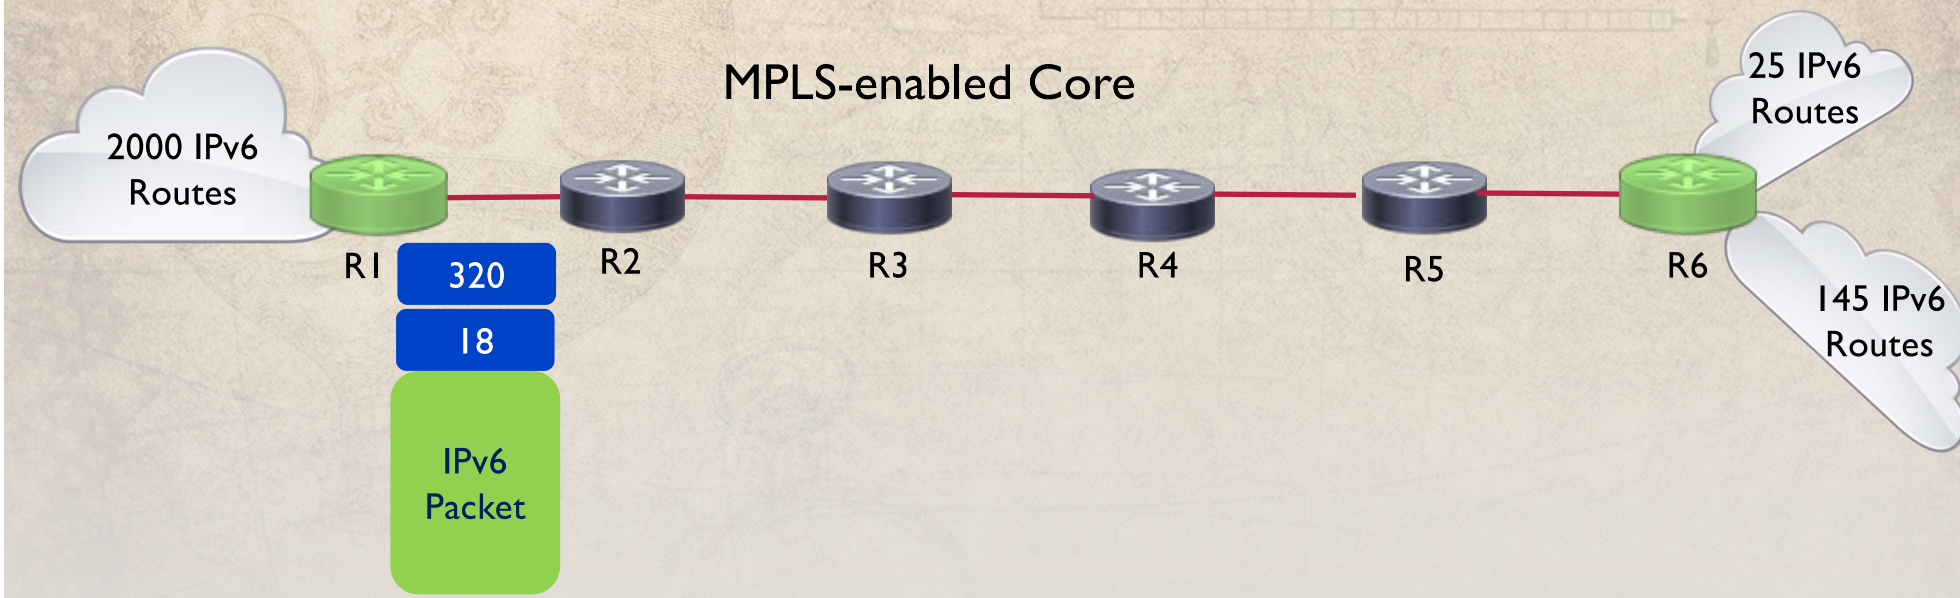

END TO END FLOW

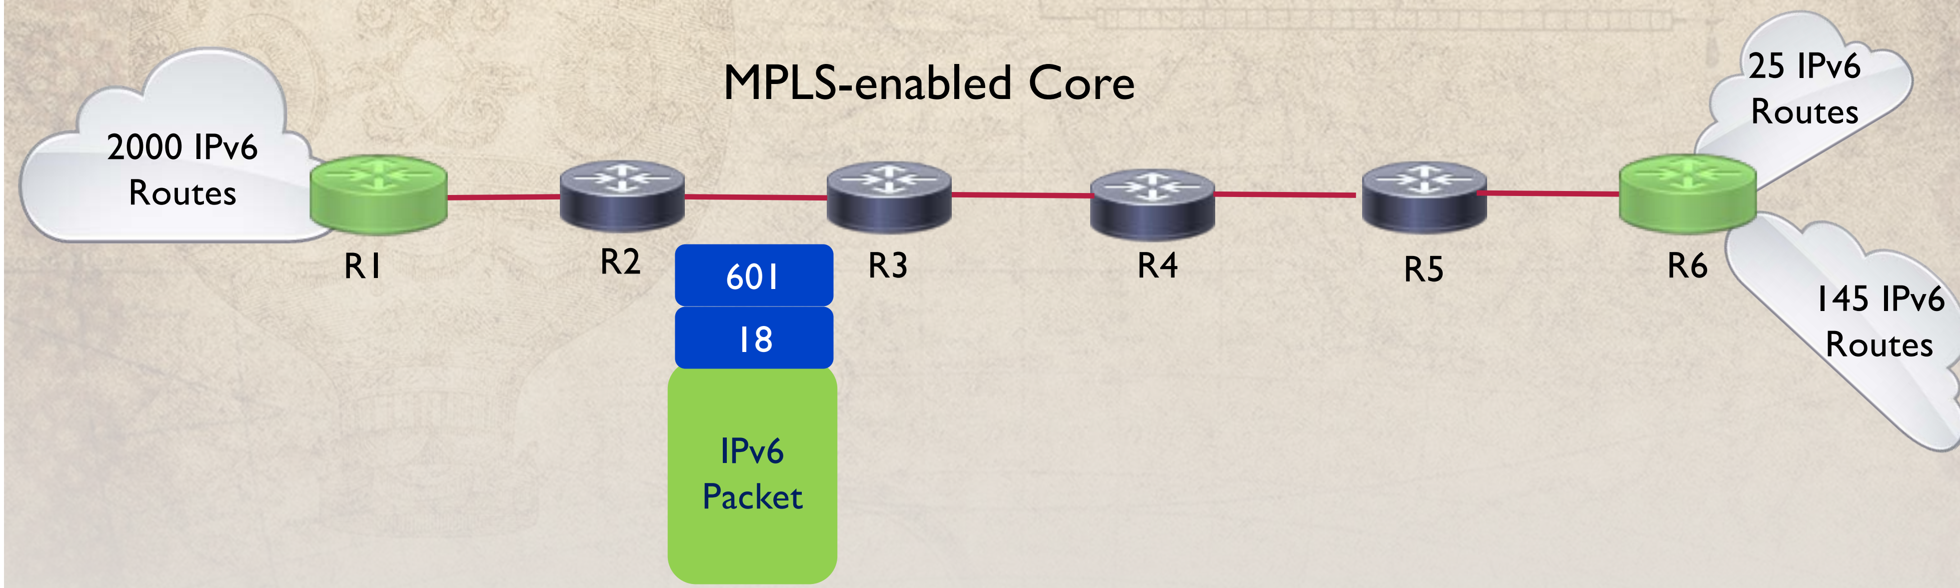

END TO END FLOW

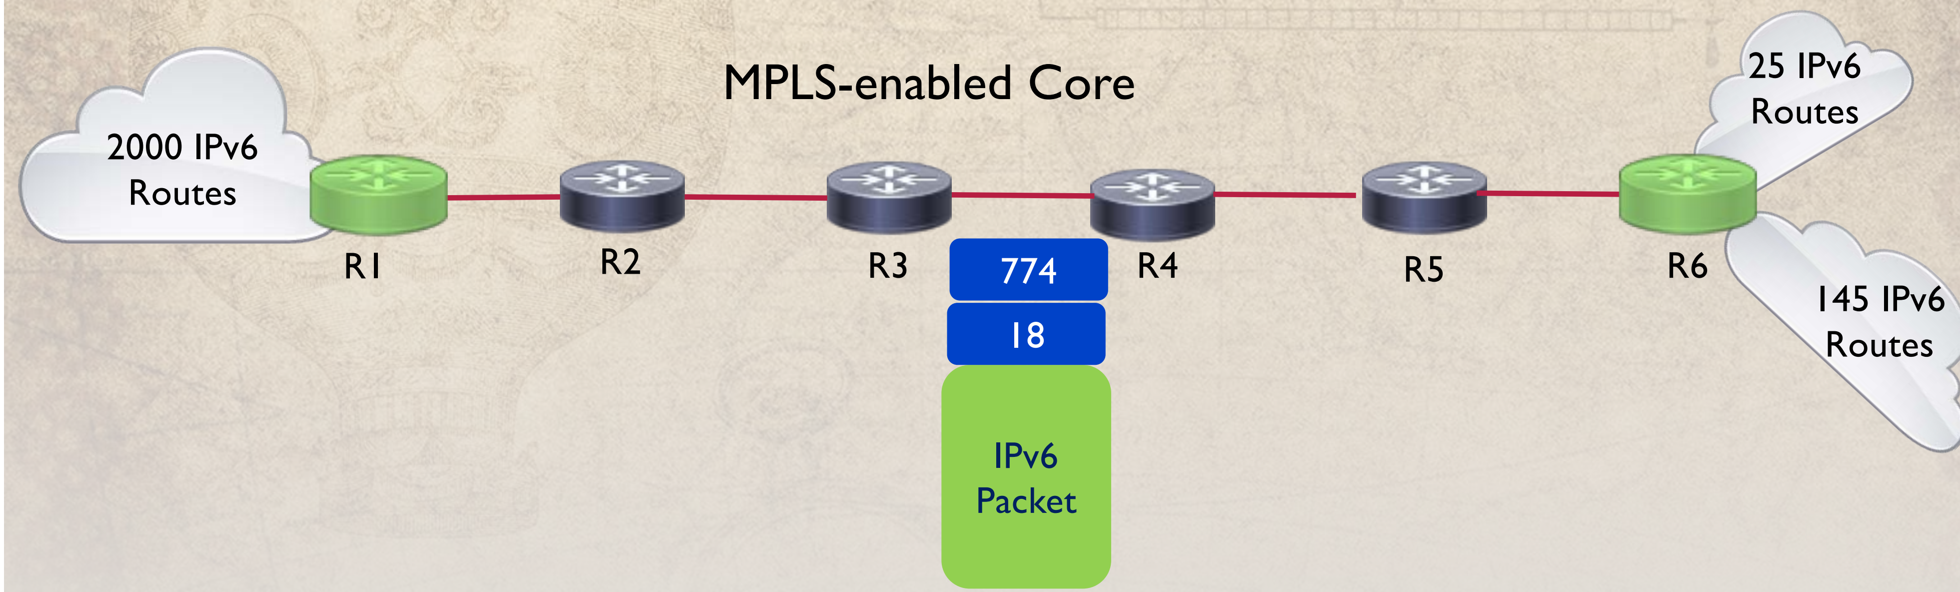

END TO END FLOW

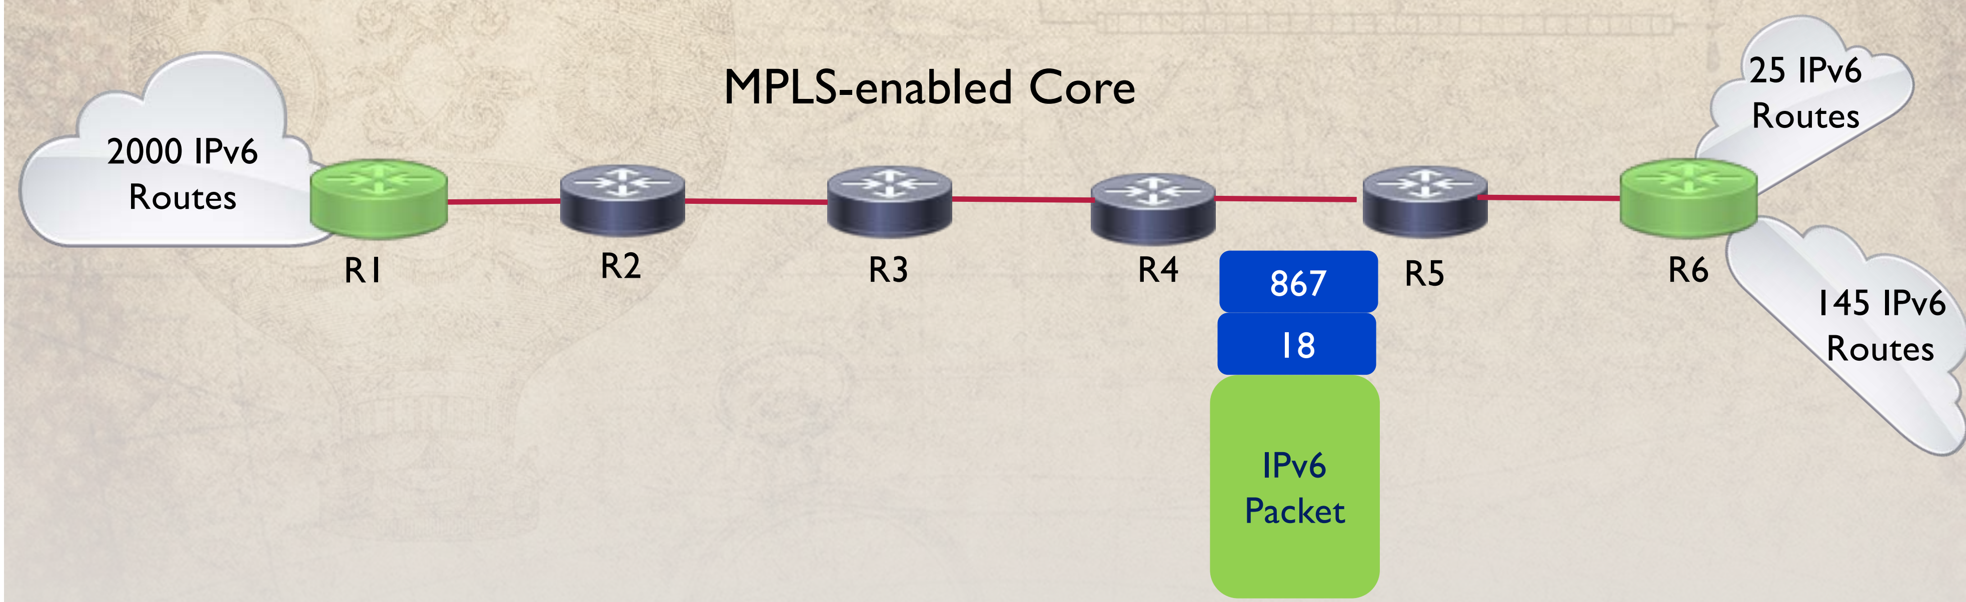

END TO END FLOW

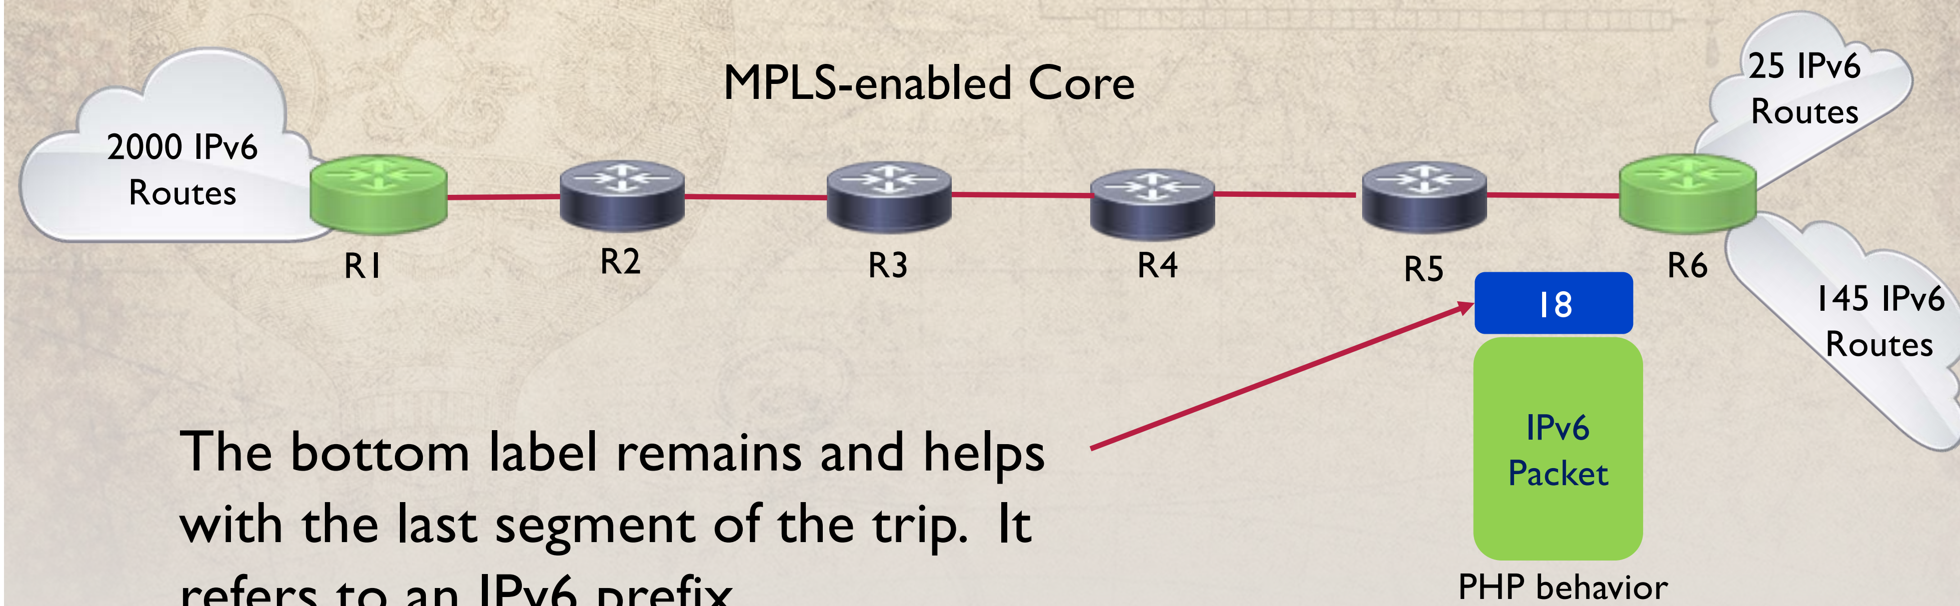

The bottom label remains and helps with the last segment of the trip. It refers to an IPv6 prefix

USE CASE

- Ready to discuss a real-world scenario?
  - Any questions before we move on?

USE CASE – COMPANY A

- Company A is a laptop manufacturer with around 15000 employees. They receive parts from their partners and assemble in their facilities in Taiwan.
- They have 150 offices including 100 small sales offices, 20 engineering or Q/A sites and 30 corporate only and customer service locations.
- The engineering teams in Europe and NA, choose hardware, design and Q/A.
- Users have all sort of PC's and laptops plus a lot of different machines including SCADA, 6-axis robots, PLC's, scanners, engineering printers and thin-clients.

THE CURRENT NETWORK

- Their existing network includes a sprawling global all-IPv4 WAN with some single points of failure. Although most of these points had redundant devices in the original designs yet due to the organic growth both are actively in use.
- BGP is used for WAN connectivity with mighty route reflectors
- IT has concerns about quite a few tier-2 WAN routers and their remaining memory and CPU horsepower.
- The WAN core also includes several POP's and was mostly built between 2012 and 2014
- While the leadership understands the future trends of IPv6 yet they prefer to keep it out of the core until it's either proven or they've upgraded the WAN core to more up-to-date hardware.

SAMPLE PROJECT REQUEST

Dear IT,

Although we understand no corporate application currently needs IPv6, yet; 3 of our engineering labs are in urgent need to IPv6 connectivity. The number and scope might increase with little notice. Those labs are in Santa Clara, Frankfurt and Vancouver.

Our new hardware (mix of everything; IoT and stations) will arrive in 30 business days and we should be prepared to start testing using IPv6 almost right after.

Any delay on your side can delay a major product shipping.

John G.

Head of Engineering Connectivity

KEY POINTS ABOUT THE DESIGN

COMPANY A

- Unknown machine types and IoT
  - Rules out machine to machine tunneling solutions
- Expediated deadline
  - Rules out invasive and high risk changes such as dual stack everywhere
- Vendor support
  - They might have different vendors and IPv6 tunneling support is not consistent
- Stability is key
  - Keep in mind you don’t want to risk/destabilize the entire corporation to support 3 labs
- Scope can grow
  - Solution must scale

LET'S RULE OUT SOME DESIGN OPTIONS

- IPv6 only sites
- Client to client tunnels
- Site to site tunnels
- Dual stack everywhere

A POSSIBILITY…

- 6PE and MPLS in the core:
  - The 3 sites can run both IPv4 and IPv6 as needed
  - Stable and pretty "beaten up" by the largest operators in the world
  - Gradual learning steps aka learn as-needed (v6-free core)
  - Reduced number of new "features" and protocols in the core
    - Load and risks of bugs
  - Quick new expansion
  - Can easily be automated and monitored compared to a dual stack IPv4 IPv6 global core
  - MPLS is not as "new" as v6 in the core and chances are low that you have to upgrade codes
  - Benefits of larger users community: MPLS in the core vs IPv6 in the core
  - And it scales well

USE CASE – LOCAL BGP ROUTE REFLECTORS

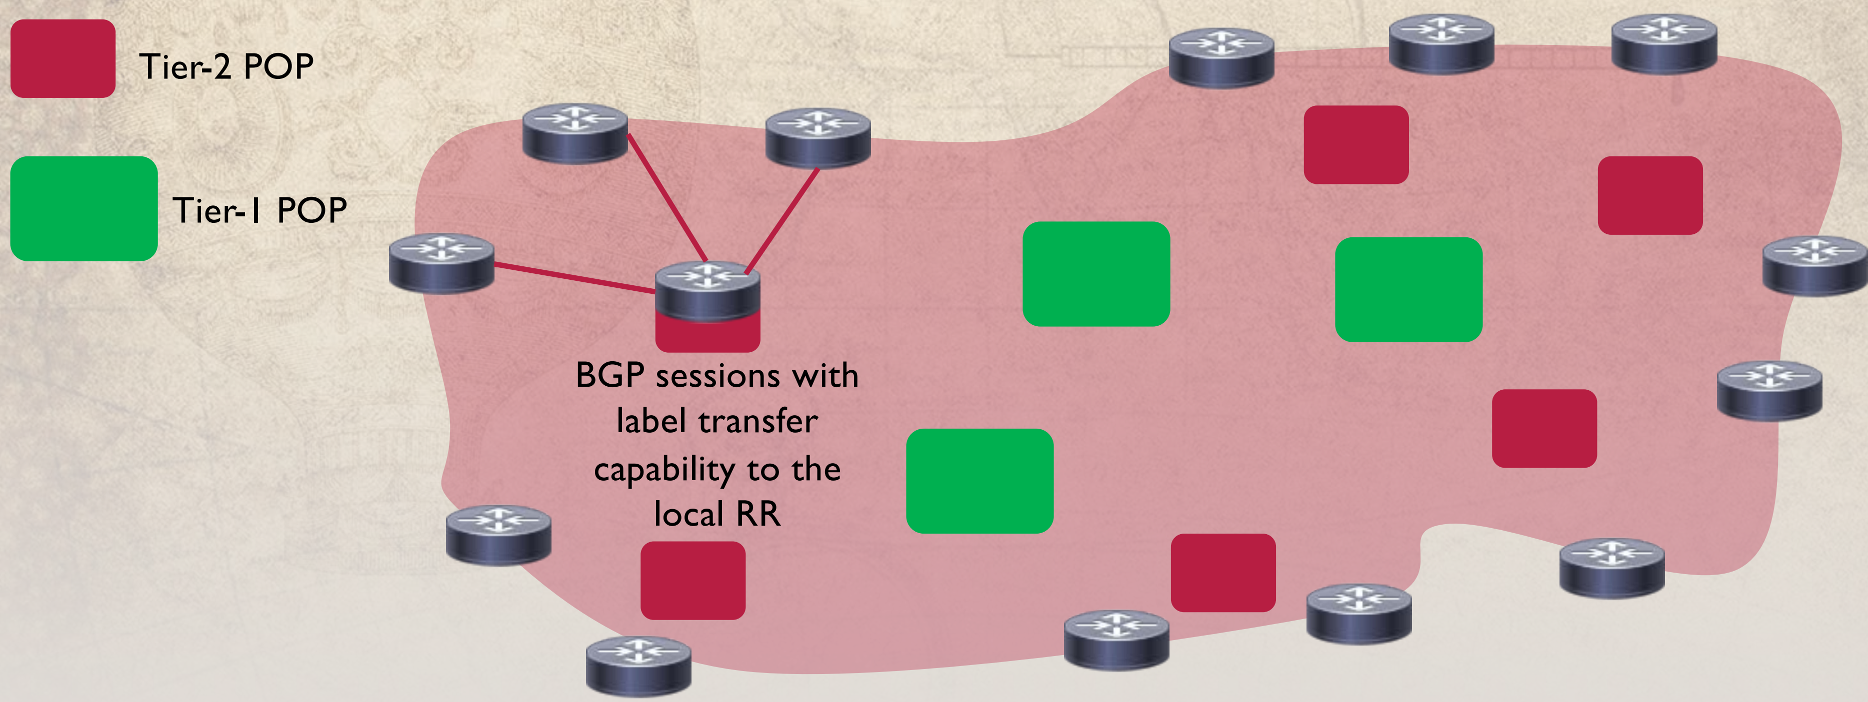

KEY TAKEAWAYS

- The core does not "understand" IPv6 unless it directly participates in the control panel
  - Such as the route reflectors and edge devices
- This design opens doors to other enhancements and bonus points
  - Traffic Engineering
  - Fast Reroute (~50ms convergence)

KEY TAKEAWAYS

- What if Company A decides to roll out virtual routing instances or VRF's and treats each site/lab as a segmented customer?
  - The same concept can be used and its called 6VPE.
    - Can be a future presentation

KEY TAKEAWAYS

- 6PE is a yet another tool in your toolbox

Q/A