(respectful) Change is hard

Shaqayeq Saleh,
George Michaelson,
Tashi Phuntsho

APNIC

Bicsi

About APNIC

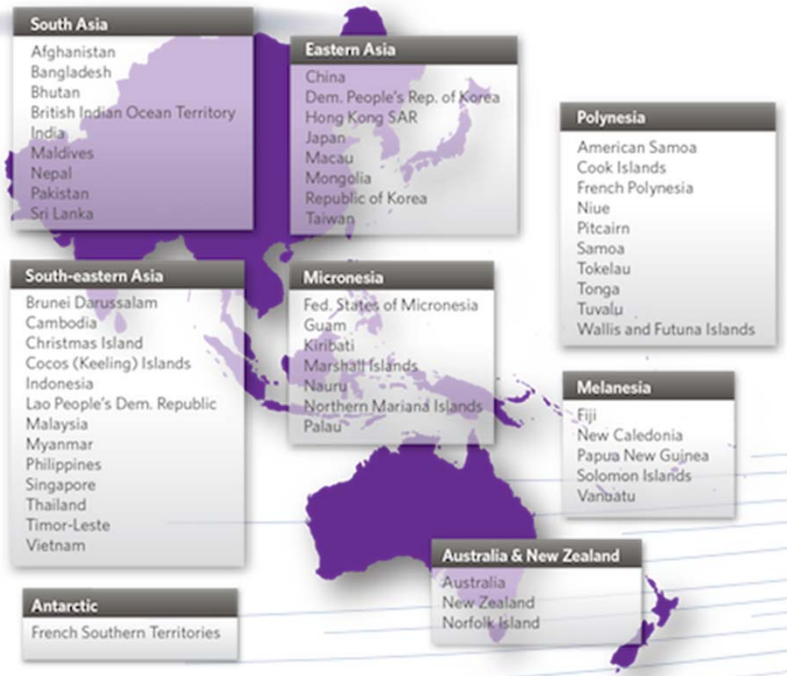

- One of five Regional Internet Registries (RIR) worldwide
  - Vision: **"A global, open, stable, and secure Internet that serves the entire Asia Pacific community"**
  - **Regional structure: 56 economies, with some National Internet Registries (NIR)**
- **Responsible management of the public interest in Internet Numbers**
  - **IPv4 addresses**
  - **IPv6 addresses**
  - **Autonomous System (AS) numbers**
- **Member body, consensus based processes**

BICSI

About the industry

- Hosting, DC is our second-largest membership

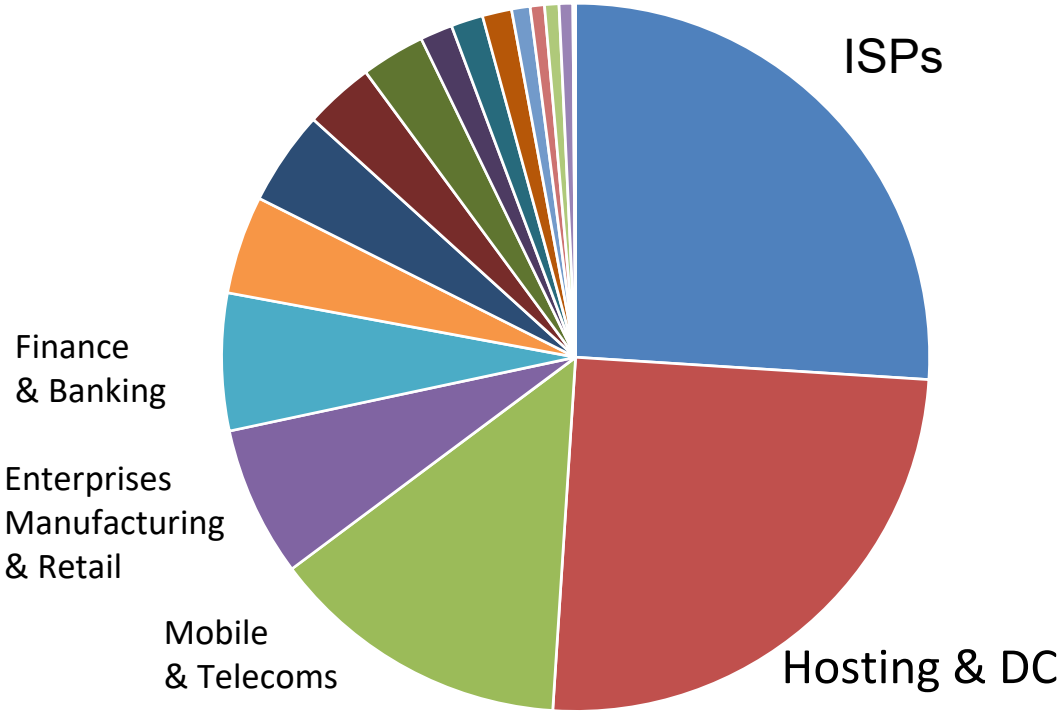

- Internet service provider (ISP)
- Hosting/Data centre
- Telecommunications/Mobile operator
- Enterprise/Manufacturing/Retail
- Banking/Financial
- Software vendor
- Academic/Educational/Research
- Government/Regulator/Municipality
- Media/Entertainment
- Industrial (construction, mining, oil)
- Infrastructure (transport/hospital)
- Other
- Internet exchange point (IXP)
- Domain name registry/Registrar
- Hardware vendor
- Non-profit/NGO/Internet community
- NREN/Research network

Bicsi

Hosting/DC and addresses

- Most of the other large membership categories are customers of these services
- …So Trends and dynamics in the area relating to Internet Number Resources (INR) in the DC, or for long-haul communications are very important to us
- How is this industry sector tracking?

Hosting/DC by INR type

ASN only	IPv4 only	ASN and IPv4	IPv4 and IPv6	IPv6 only	ASN and IPv6	ASN, IPv4 and IPv6
10	435	361	211	12	2	563

796 DC members with address holdings (49%) have no IPv6

2 members have the bulk of IPv6 holdings. Most members have A small footprint and may need to discuss future IPv6 needs

Bicsi

Change is coming

- World IPv6 capability is tracking at 16%
  - Asia-Pacific at 18%
  - Japan at 25%
  - USA at 40%
  - G20 at 21%
  - G8 at 28%

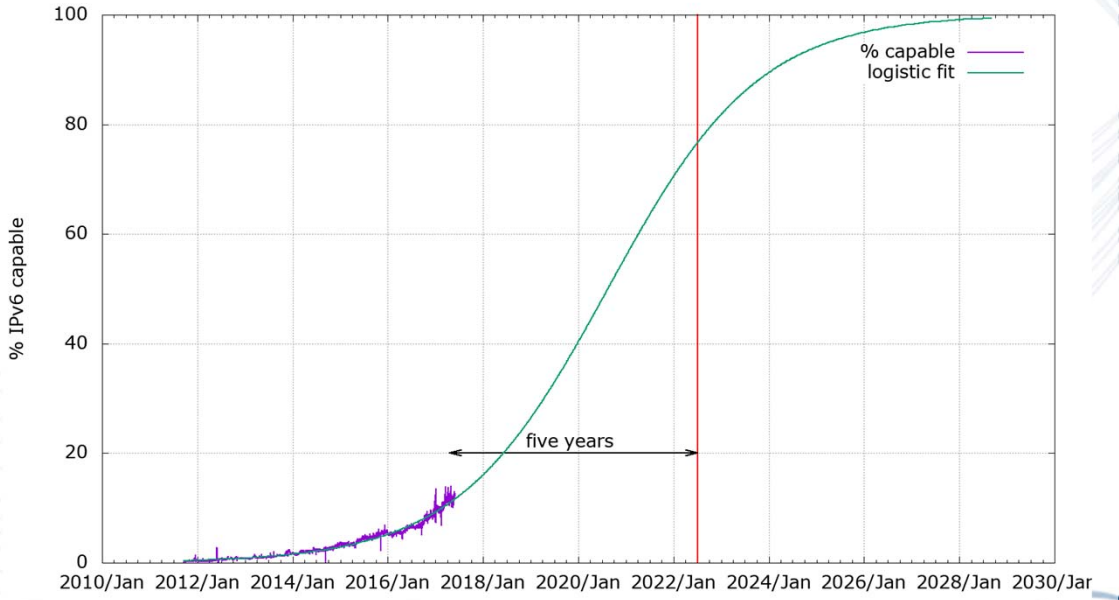

IPv6 in the DC

- Emerging RFCs of interest:
  - RFC7755 Stateless IP/ICMP Translation for IPv6 Data Center Environments (Feb 2016)
  - End Work on IPv4 draft-ietf-sunset4-ipv6-ietf-01 (work in progress)
  - Gap Analysis for IPv4 Sunset draft-ietf-sunset4-gapanalysis-09 (work in progress)

- DC deployments have many assumptions about IPv4
  - The switching fabric, network management, associated plant.. Is it IPv6 enabled?
  - Are you ready for a future where you run your DC over IPv6 only networks?

BICSI

Dual Stack is harder

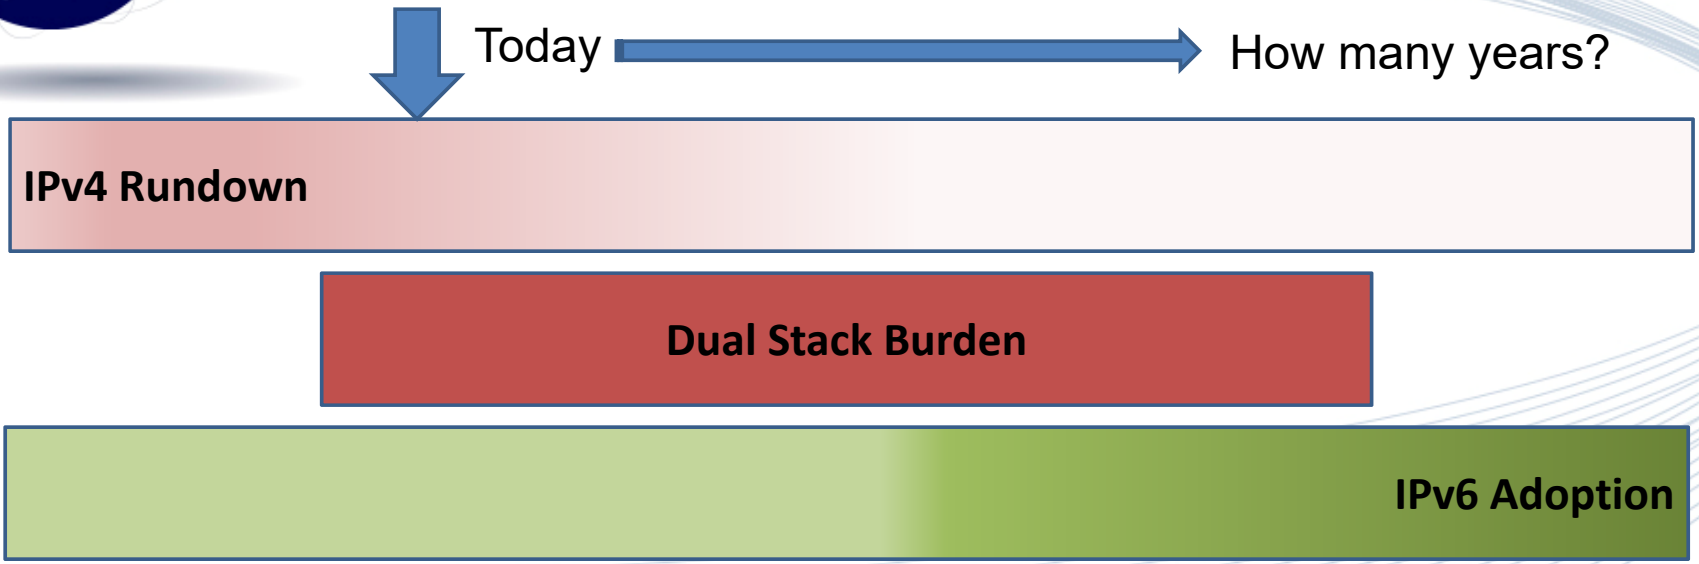

- We did not architect the IPv6 transition for a long duration, and the cost of dual-stack is a persisting burden. How long? We don't know.

Bicsi

We need the industry to change

- We think its going to be necessary for the DC industry sector to engage actively in IPv6
  - Check supply chains for IPv6 only operations
  - Source product with IPv6, dualstack enabled
  - Migrate operations to IPv6 in the NMS and provisioning
- How do you do this kind of change?

Bicsi

Why does this matter?

- Data Centers are one of the fastest growing membership categories in the RIR space
- Data center praxis is founded on IPv4 deployment models
  - Rack management
  - Server farm management
- Cloud service models (kubernetes) tend to reflect IPv4 deployment
  - Use of NATs, RFC1918 addressing inside the cloud
  - ACL practices from IPv4
- Everything needs to be reconsidered in an IPv6 world
  - How do we do this, in a respectful, safe manner?

Bicsi

Our fears

- DC has large investment in IPv4
  - DC cannot obtain publicly routable IPv4 at scale
  - DC builds out to RFC1918, deploys CGNAT at edge
  - DC deploys IPv4/6 proxy at edge
- Outcome: DC remains an IPv4 ”island” with all IPv6 mediated through proxies increasing burden

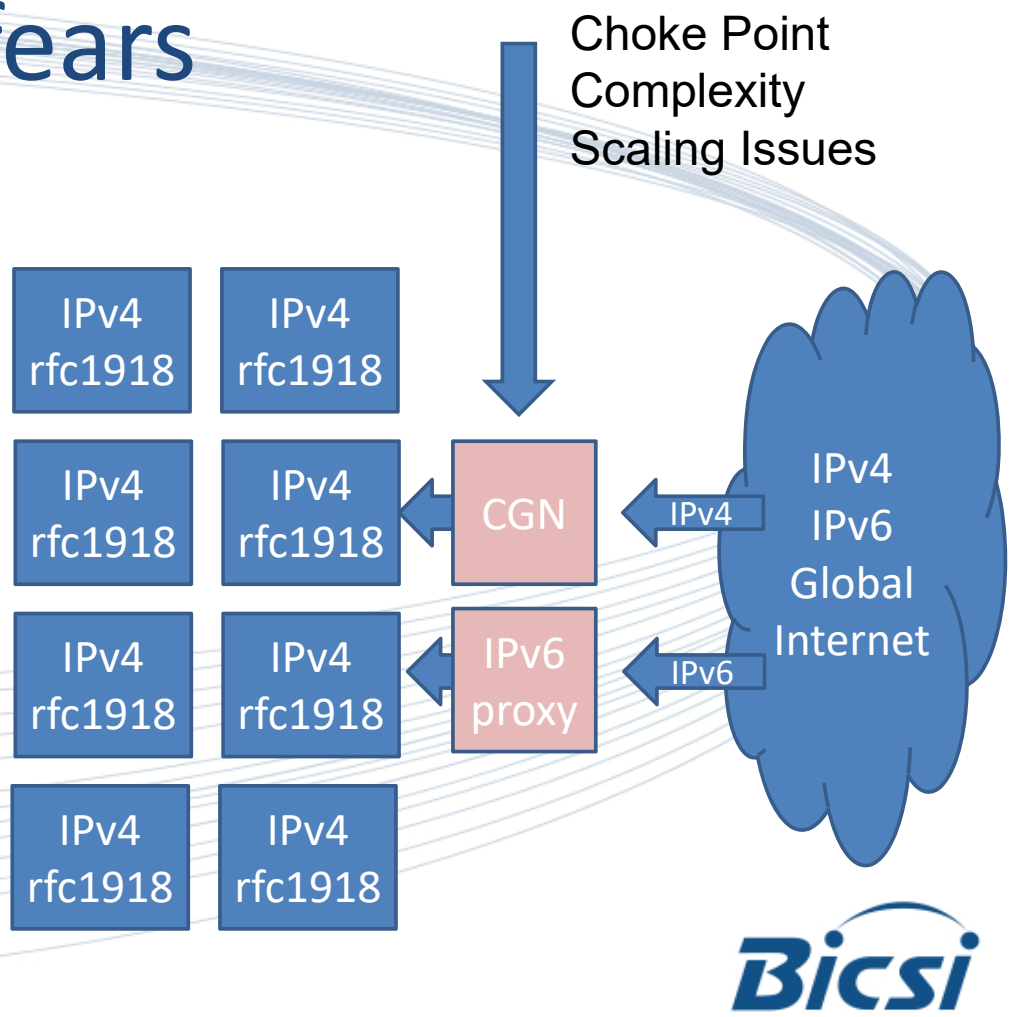

Bicsi

Our Goal

- DC deploys public IPv4 at edge
- DC deploys IPv6 ubiquitously, and builds to IPv6
  - DC deploys V4/V6 proxy to redirect to IPv6 inside

- Outcome: IPv6 services are end-to-end IPv4 proxy mediated but declining burden

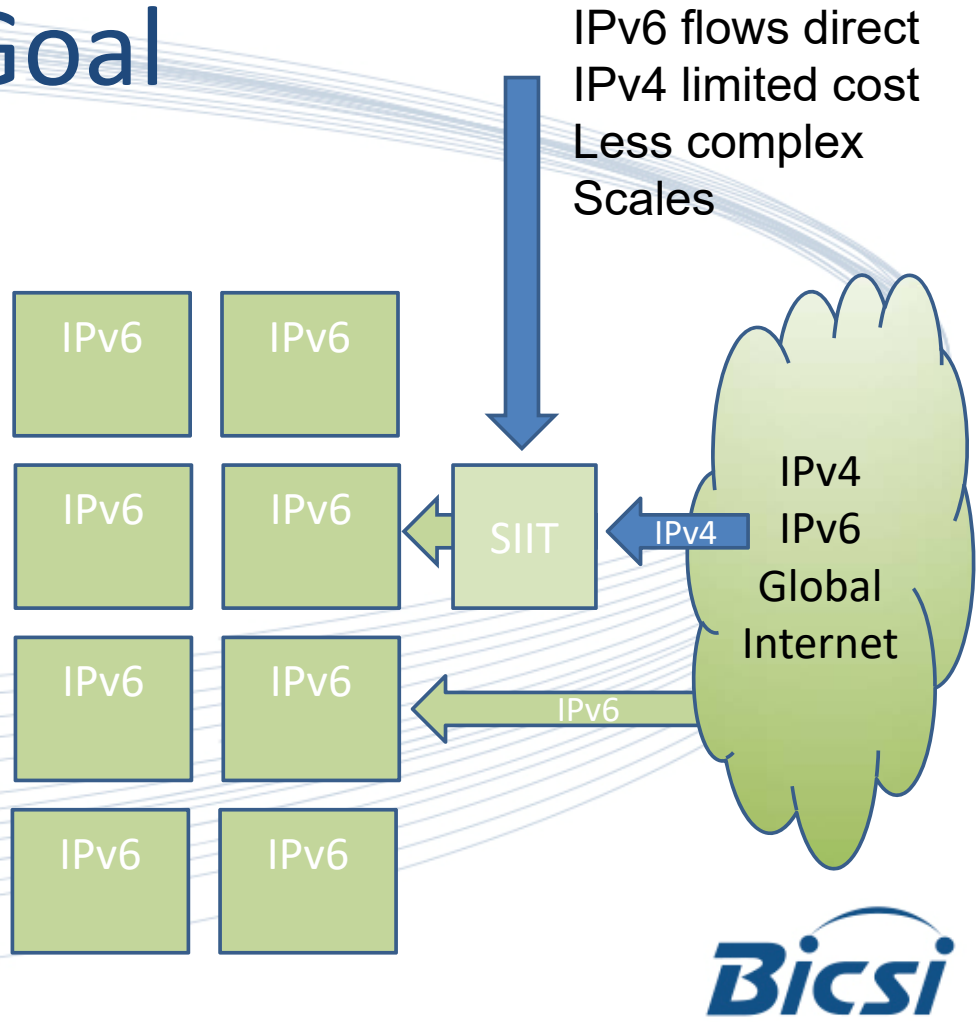

Bicsi

We respect our Elders

Bicsi

The problem: We respect our elders

- Many cultures of the Asia Pacific place high importance on respect for elders
  - Bhutan has a very strong culture of this kind with codified expectations.
  - Malaysia, Singapore, Vietnam, Korea, China, Japan share aspects of this quality in their heritage, e.g. widespread historical role of Confucianism.
- Networking is subject to rapid change, “pivot” moments
  - How do you signal radical technology shifts, in a context of respect for your prior leadership?
  - What's the right way to effect change in a respectful way?
- Here’s a story from Bhutan, about an approach to technology shift.
  - Which worked
  - Which respected and engaged with the existing responsible leadership

Bicsi

Bhutan Telecom experience

Meet Tashi: APNIC Training

10 Years experience in IP and Transmission network design, operation, and maintenance experience, having worked for the incumbent telco in Bhutan.

He has been involved in capacity development in the APNIC community by providing training in number of technical areas, such as, Routing, Switching, Network Design, Network Security, IPv6.

Tashi has a degree in electrical engineering from India. He also has research experience on next generation networks from Japan, complemented by a Master's Degree in network systems from Australia.

Bicsi

Pre-History

- Before 2004 ISP built from
  - 8mbit aggs to SAT link
  - PDH Microwave feeds to copper last mile, WLL
- Voice, Dial-up, ADSL, 2G, IRC Chat
- This technology is a logical step-up from prior Voice service builds with limited data. Its pre-internet.

Bicsi

2005/6

- 12mbit aggs to SAT link
  - SDH STM-1 155mit Microwave feeds to copper last mile.
- Voice, Dial-up, ADSL, 2G, IRC Chat

- NO significant technology change: just more of it.
- The technology basis is getting old.

Bicsi

2007

- DS3 (45Mbps) to HKIX
  - SDH STM-1 155mit Microwave feeds to copper last mile.
- Voice, Dial-up, ADSL, 2G, IRC Chat

- Even though shifted to a DS3 fiber, still using E1 cards to handover to PoPs. The technology basis is now severely behind current best practices

2008

- DS3 (45Mbps) to HKIX, Additional to LINX
  - SDH STM-1 155mit Microwave feeds to copper last mile.
- Voice, Dial-up in all regions, ADSL, 2G, 3G in towns

- Bundling E1s to feed the DS3. This is not scaling. Congestion is building.

Bicsi

2009-2011

- HKIX/LINX moving to STM-4 by 2011
  - SDH/SONET OC192 national fiber reticulation
- Voice, Dial-up in all regions, ADSL, 2G, 3G in towns
  - Fiber to nodes in some Towns and Cities emerging
- Bundling VC fast ethernet to feed the DS3.
- This is Time to try something else.

Bicsi

Lets try something else

- Tashi secures permission to ”play” in un-used pairs of the national Fiber plan.
  - No disruption to deployed services.
  - No visible service bindings until tested to work
- Builds out a cleaner model, tests and deploys.
  - POC of 4-channel (10GbE per channel) DWDM link to the international gateways
  - Gateways in Indian border towns
  - 10GbE between core routers
- Result: simpler network architecture, ready to deploy.

Bicsi

2012-2014

- HKIX/LINX moving to STM-16
  - SDH/SONET OC192 national fiber reticulation
- Voice, Dial-up in all regions, ADSL, 3G, 4G in cities
  - Fiber to building emerging
- Bundling GigE feeding STM-16

2015

- HKIX/LINX STM-16
  - Added STM-16 to Equinix Singapore
  - 4-channel 10GbE DWDM links to all major core POPs in the country
  - Design/tender of nationwide 48-channel (10GbE scalable to 40GbE per channel) DWDM backbone
  - FTTH emerging in the cities, FTTB in major cities
- Simple, scale-able architecture being carried forward
- Legacy transmission architecture now in retirement

BICSI

Ok: lets apply this to DC models

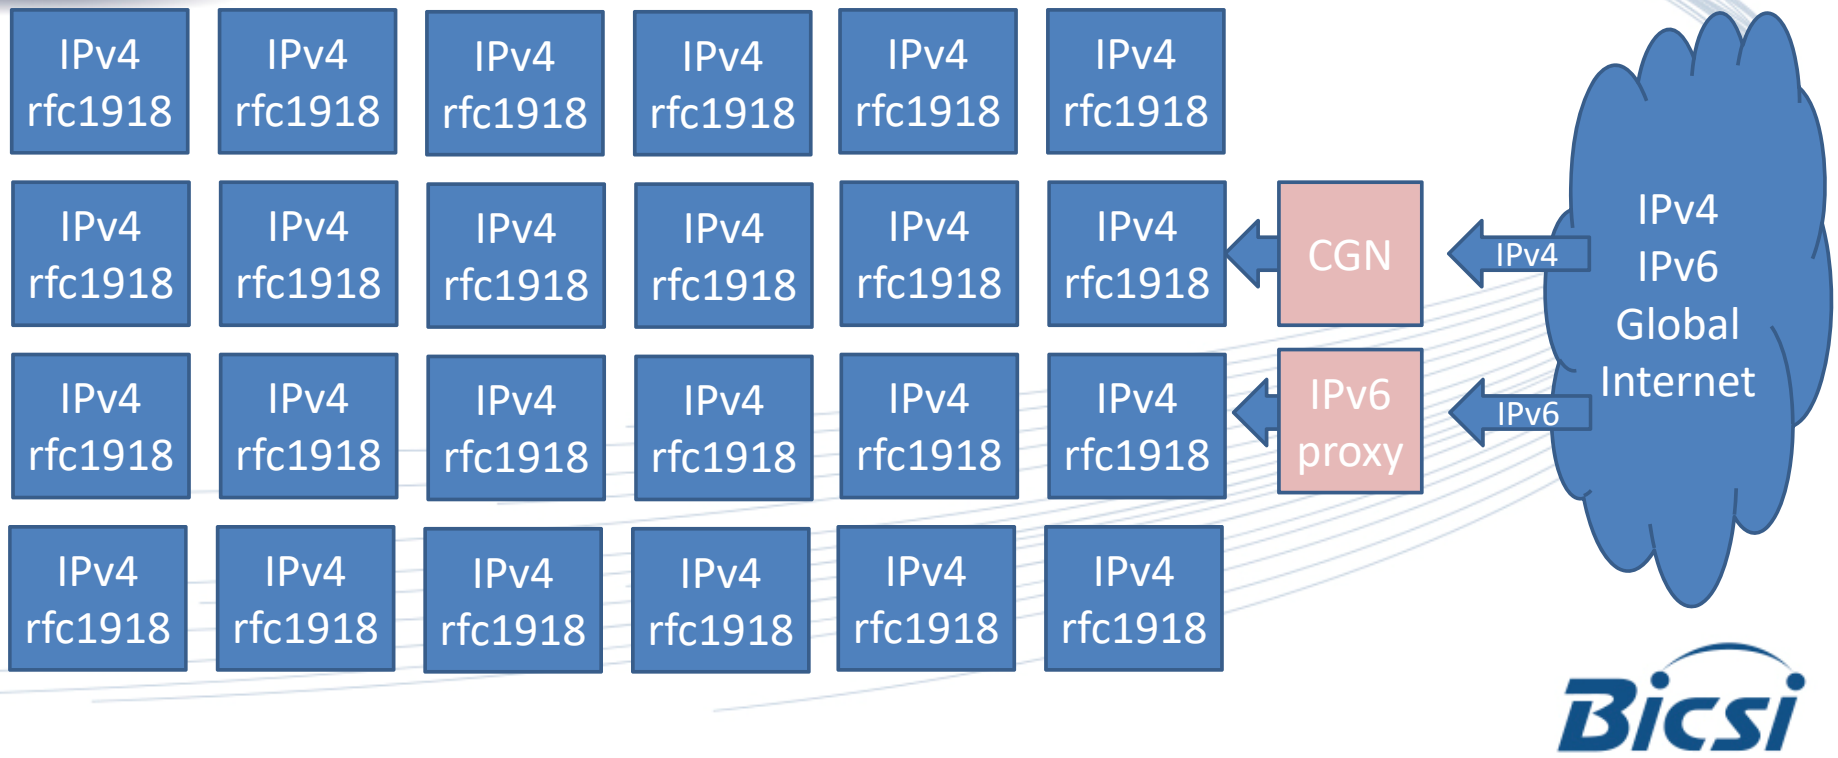

Bicsi

1: Get permission to play

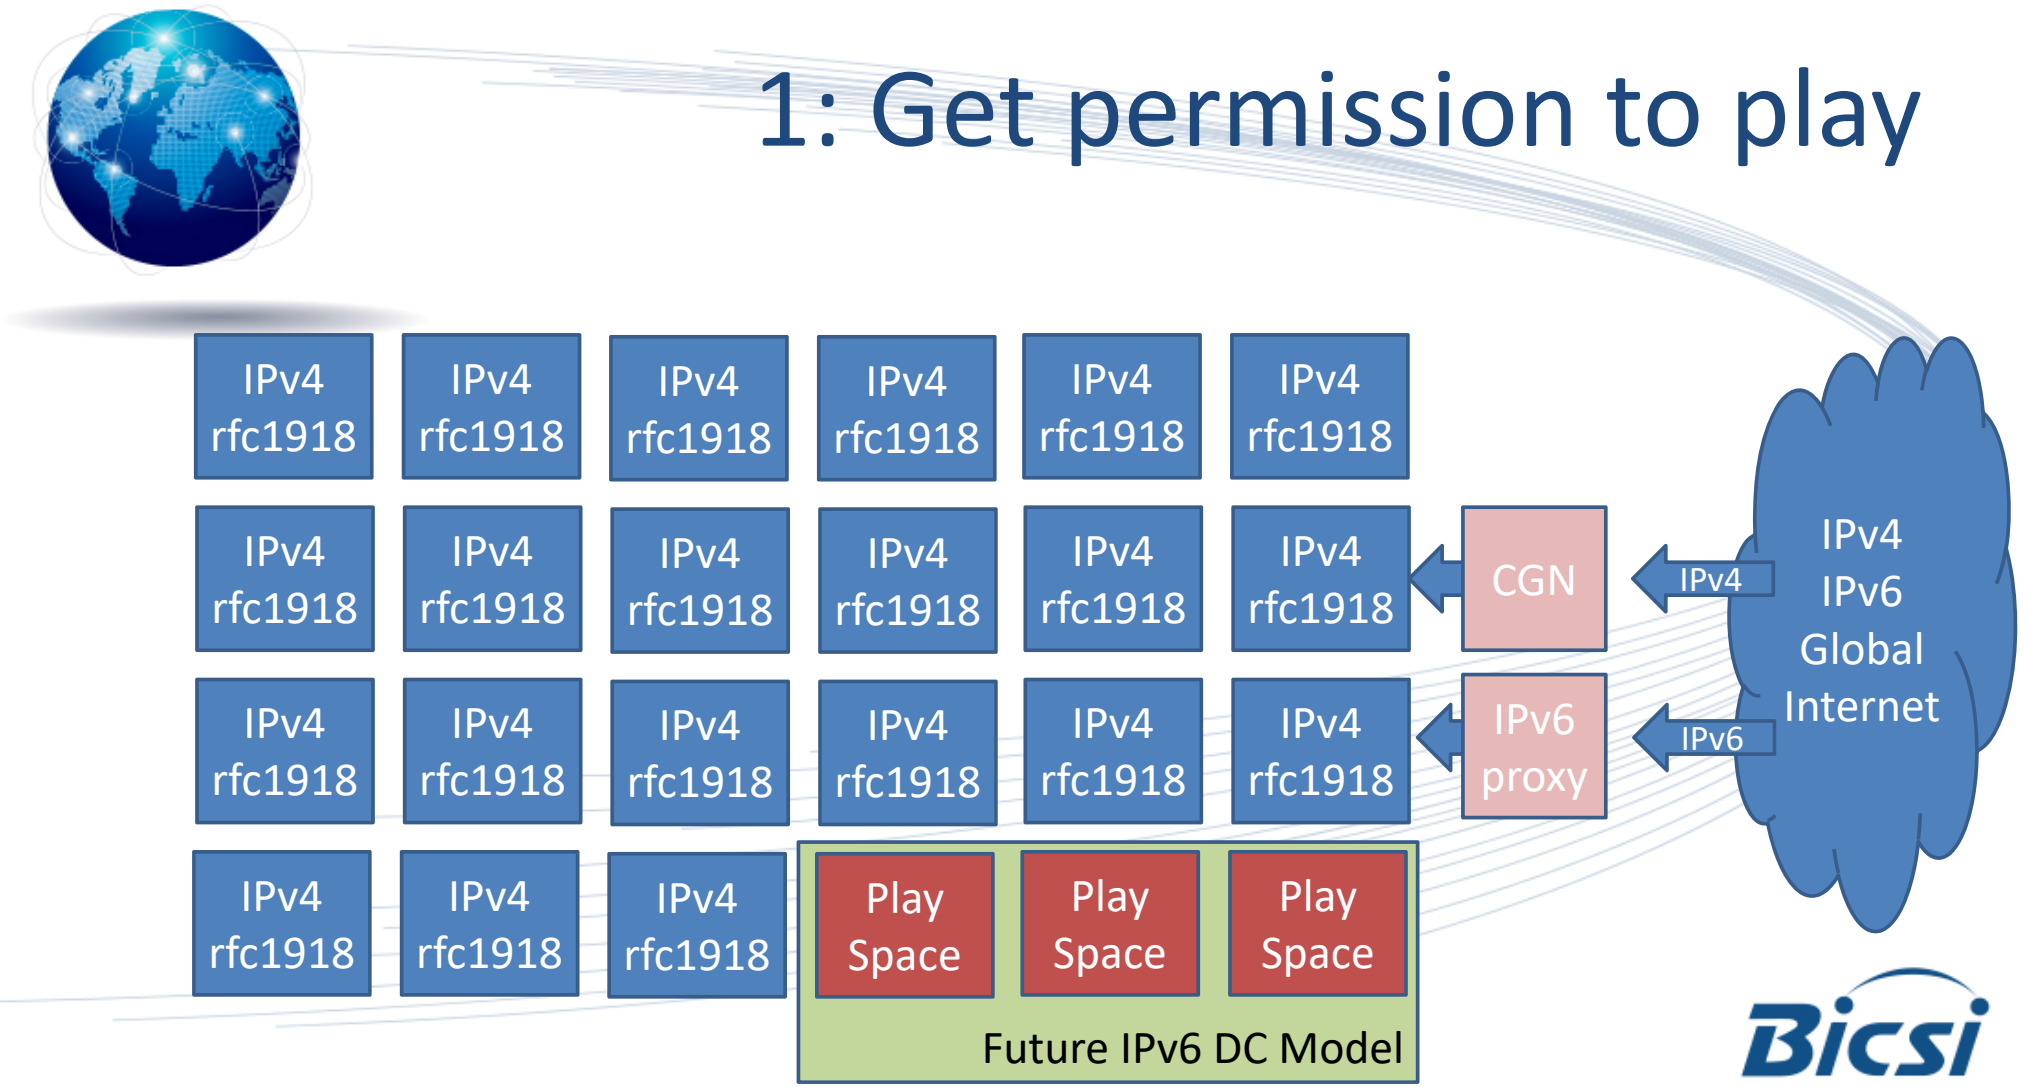

Bicsi

2: Build out Native IPv6 structure

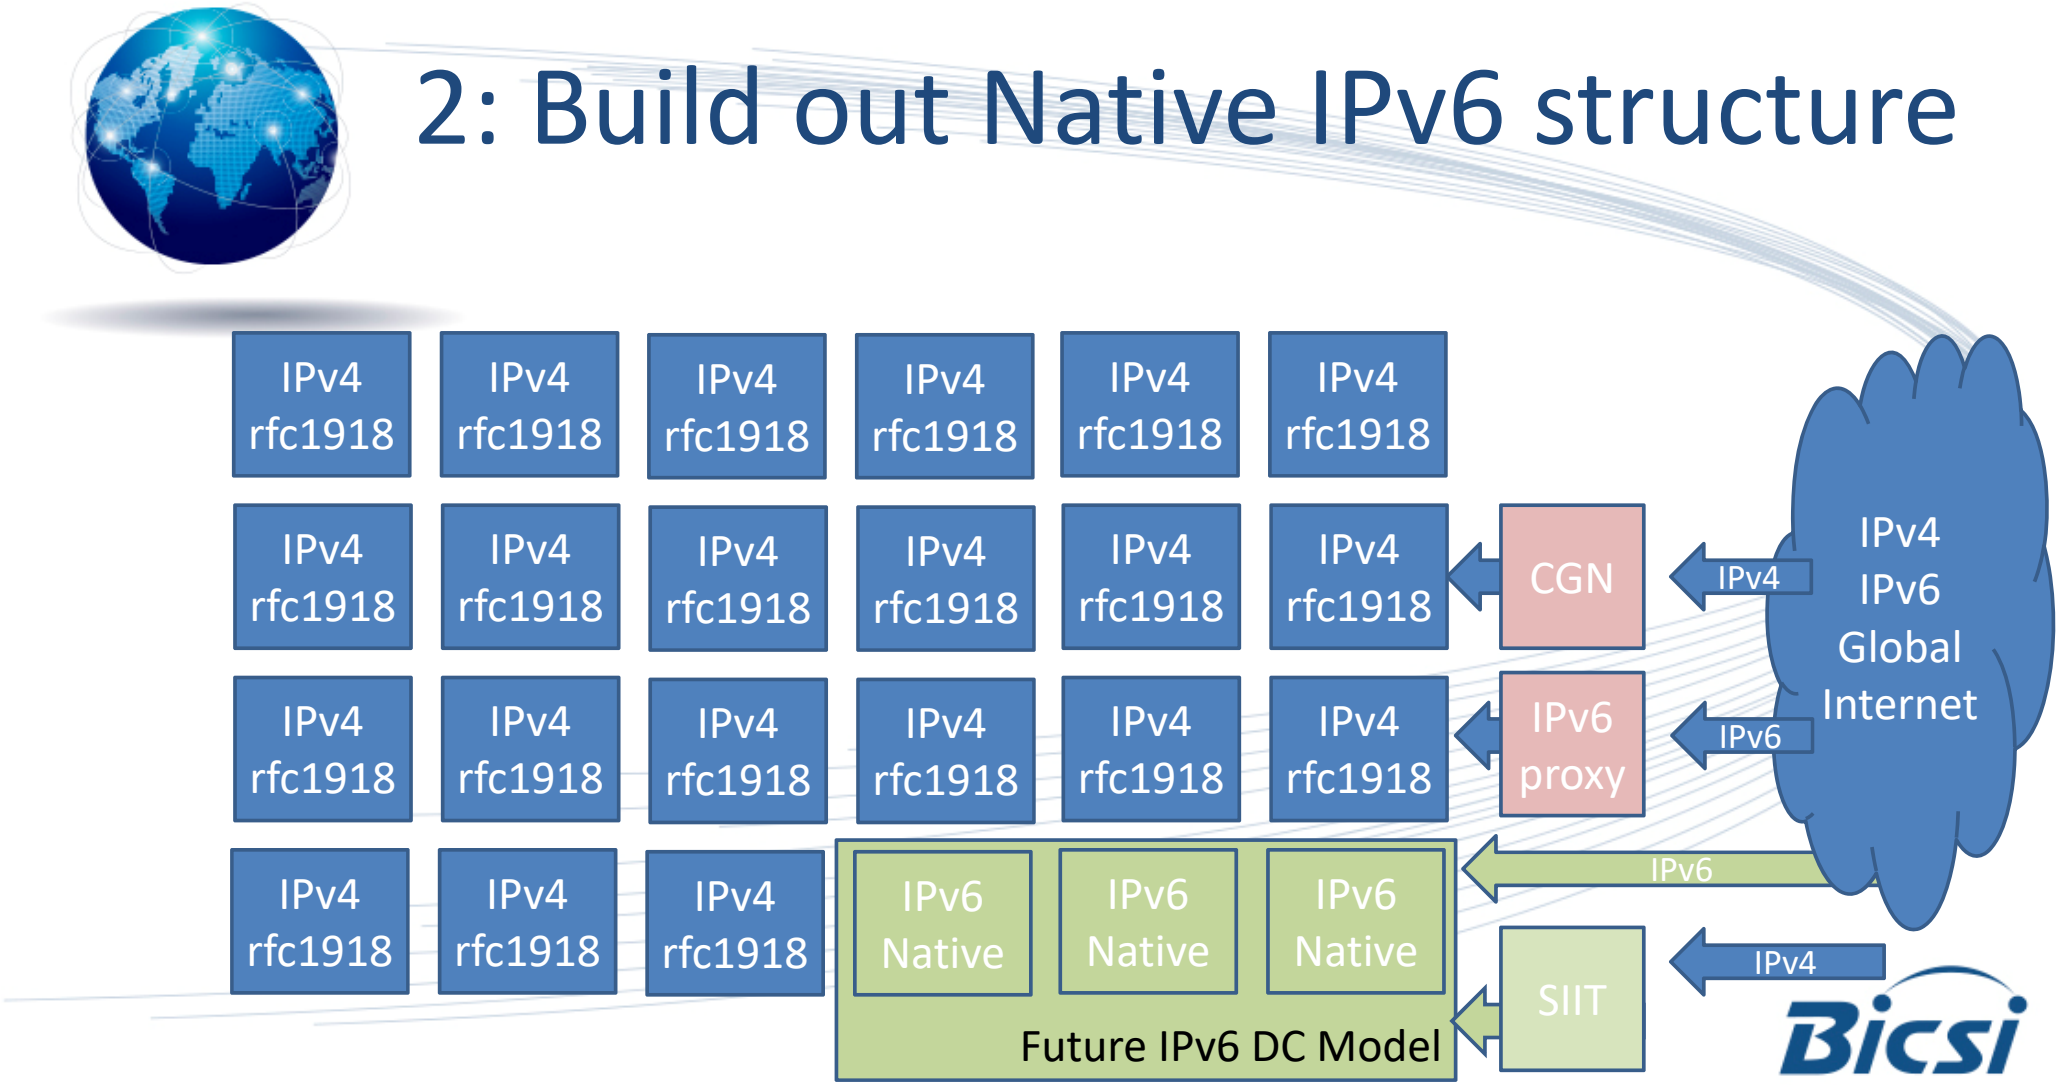

3: Test, Test, Test

- Does your supply chain have native IPv6 out-of-the-box?
- Are your NMS, provisioning s/w systems IPv6 enabled?
- Can you bootstrap full service delivery IPv6 only?
- How much residual traffic flows through SIIT to IPv4 internet?

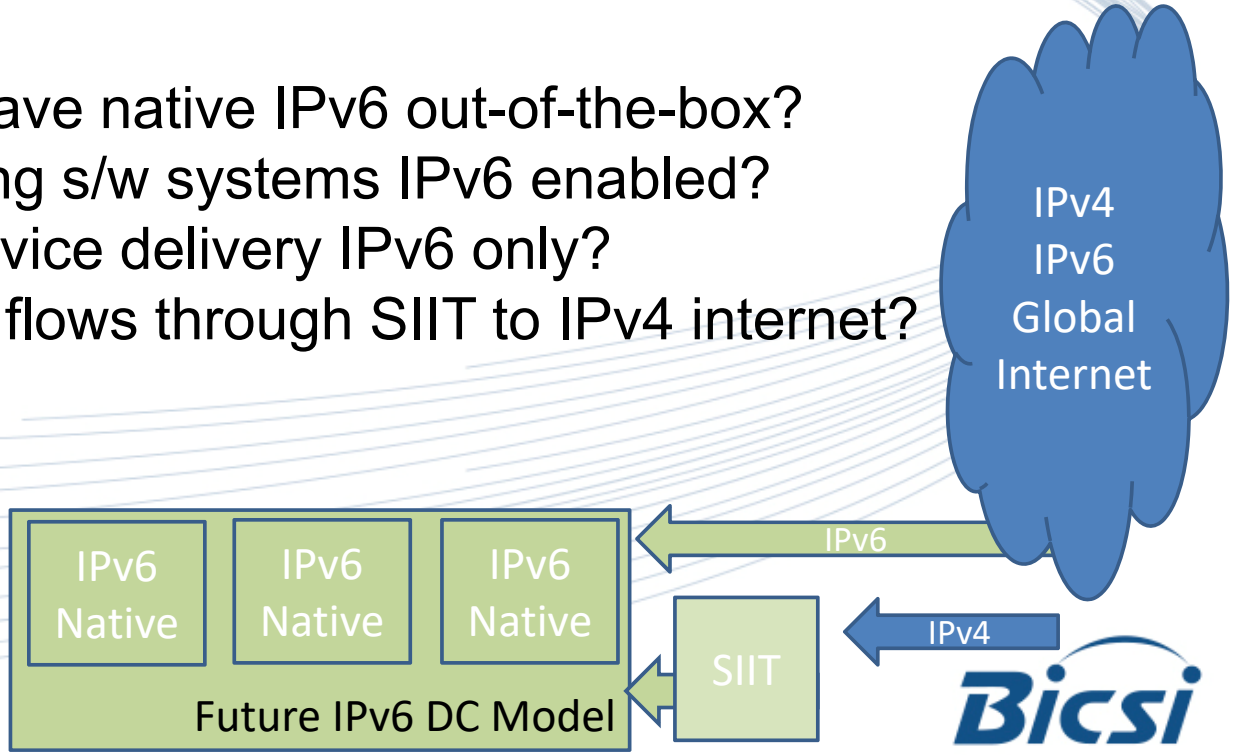

4: Embrace, Extend, Control

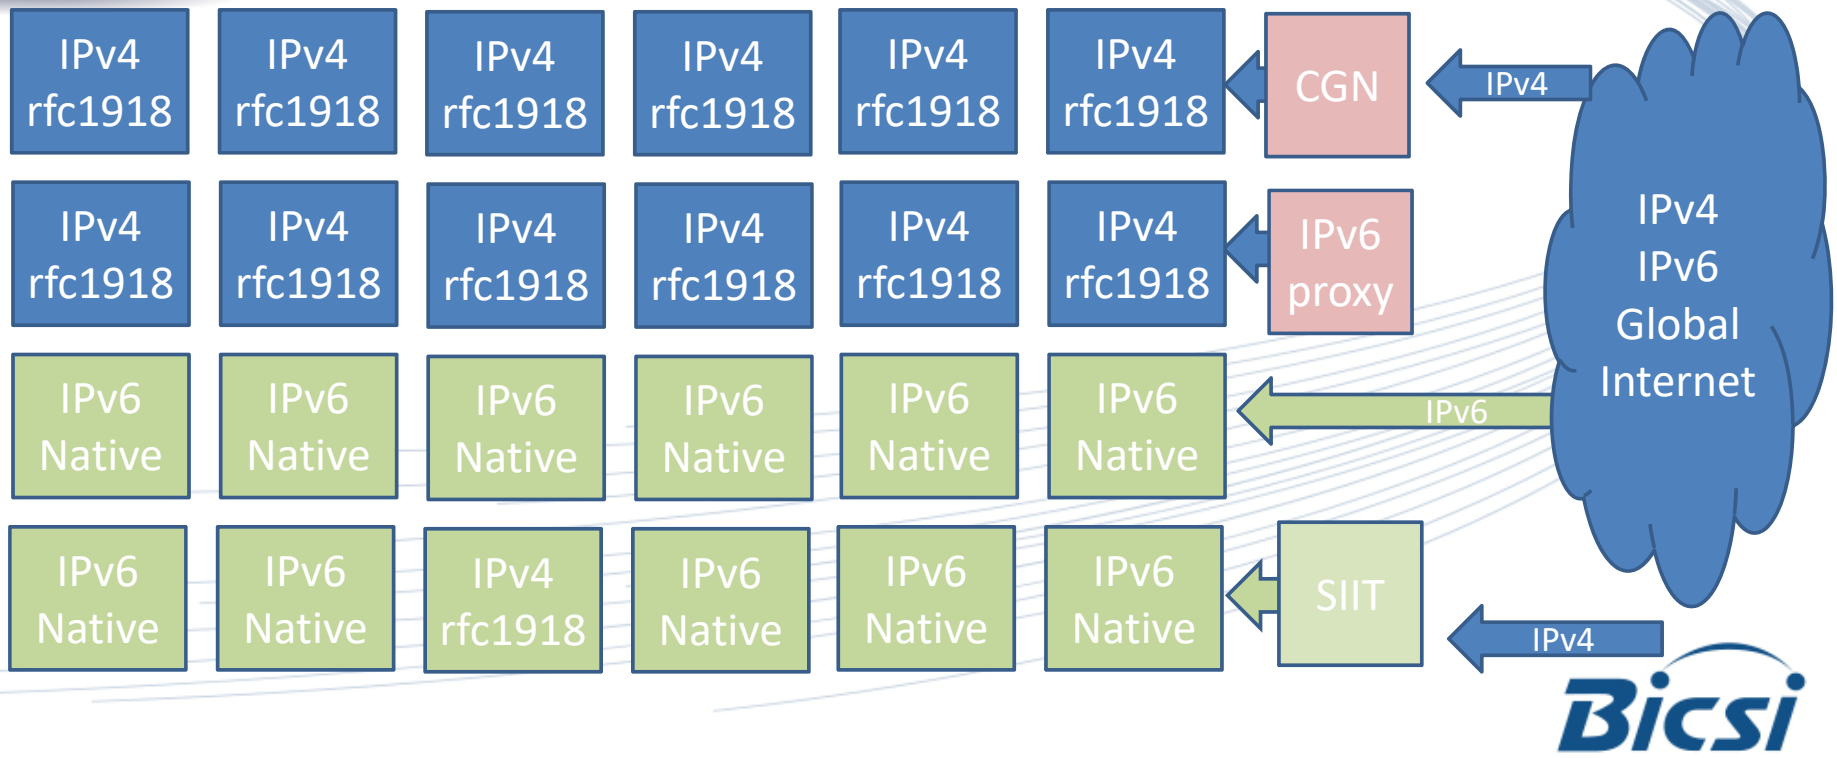

5: IPv4 sunset…

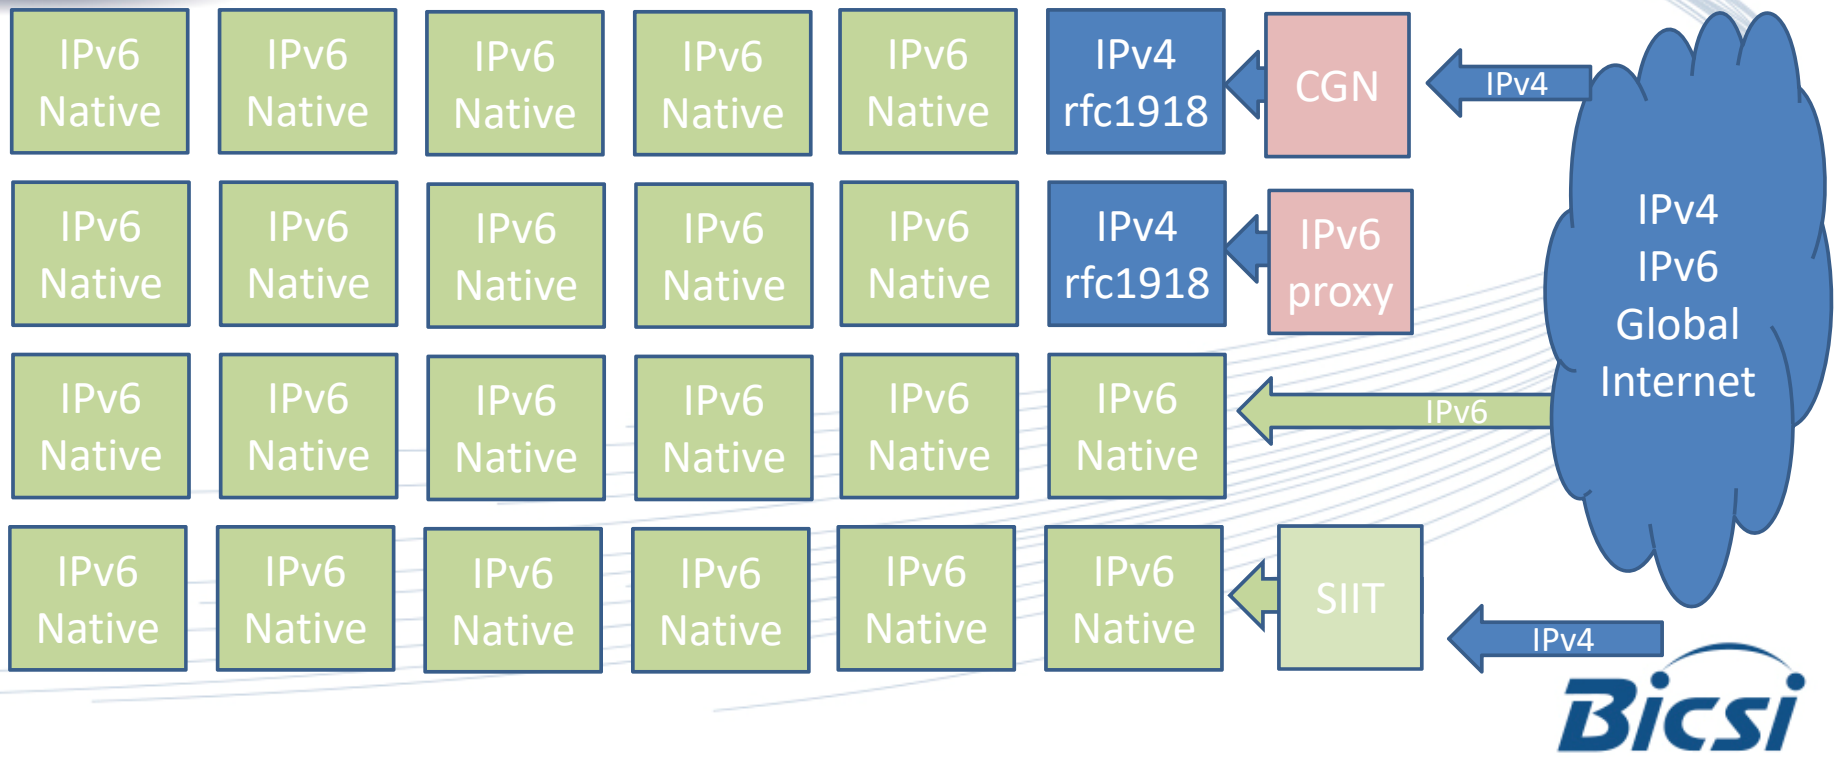

Respectful Change is possible

- Minimise disruption by using ‘permission to play’
- Build real-world deployments, capable of scaling
- Deploy into new services, underpin existing services
- Offer the old model a sensible graceful retirement

- TCO, future scaling, simpler models…

Bicsi

Thank You!

BICSI