Autonomous System (AS)

Collection of networks with same routing policy

Single routing protocol

Usually under single ownership, trust and administrative control

Identified by a unique 32-bit integer (ASN)

External BGP Peering (eBGP)

Between BGP speakers in different AS

Should be directly connected

Never run an IGP between eBGP peers

Internal BGP (iBGP)

BGP peer within the same AS

Internal BGP Peering (iBGP)

Topology independent

Each iBGP speaker must peer with every other iBGP speaker in the AS

Peering to Loopback Interfaces

Peer with loop-back interface

- Loop-back interface does not go down – ever!

AS-Path

Sequence of ASes a route has traversed

Used for:

- Loop detection

AS-Path (with 16 and 32-bit ASNs)

Internet with 16-bit and 32-bit ASNs

- 32-bit ASNs are

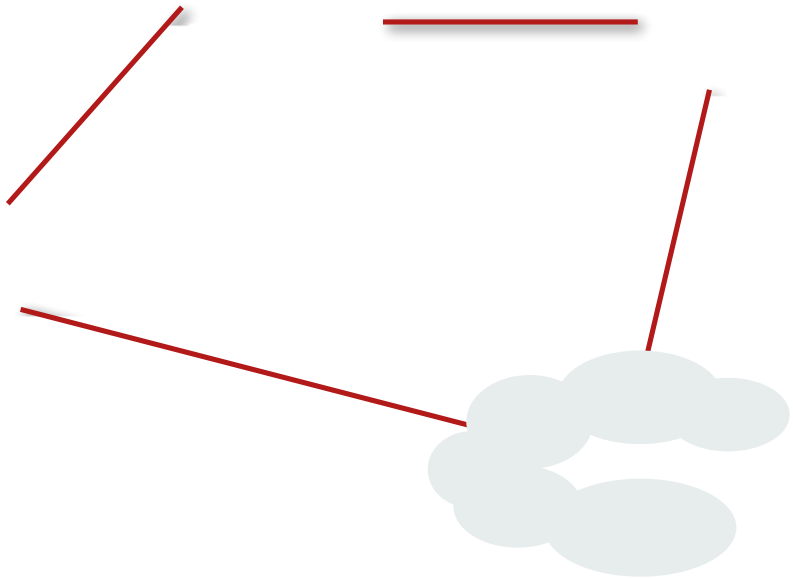

Aggregator

Multi-Exit Discriminator “metric confusion”

MED is non-transitive and optional attribute

BGP Path Selection Al50 1725oAlnd.

BGP Path Selection Algorithm for IOS Part Three

Applying Policy in BGP: Why?

Network operators rarely “plug in routers and go”

External relationships:

- Control who they peer with

BGP Capabilities

- Multiprotocol extensions
  - This is a whole different world, allowing BGP to support more

Scaling iBGP mesh

Avoid ! n(n-1) iBGP mesh

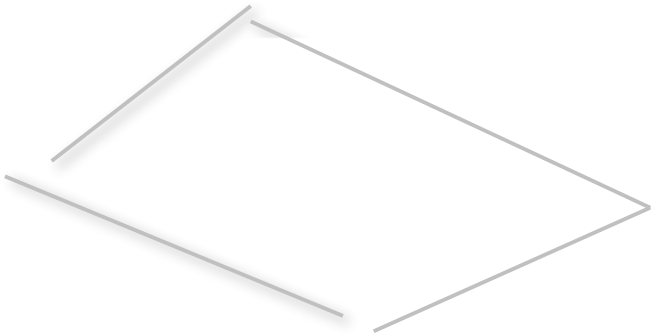

Two solutions

- Route reflector – simpler to deploy and run
- Confederation – more complex, has corner case advantages

Route Reflector: Loop Avoidance

Originator_ID attribute

Route Reflector: Redundancy

32-bit ASNs –

Representation

Representation of 0-4294967295 ASN range

Changes

32-bit ASNs are backward compatible with 16-bit ASNs

There is no flag day

You do NOT need to:

How does it work?

Compatibility Mode:

Local router only supports 16-bit ASN and remote router uses 32-bit ASN

If 32-bit ASN not supported:

Operation

- Add penalty (1000) for each flap
  - Change in attribute gets penalty of 500
- Exponentially decay penalty
  - half life determines decay rate
- Penalty above suppress-limit

Solution:

Do **NOT**

BGP for Internet Service Providers

- BGP Basics
- Scaling BGP
- Using Communities
- Deploying BGP in an ISP network
  - The role of IGPs and iBGP

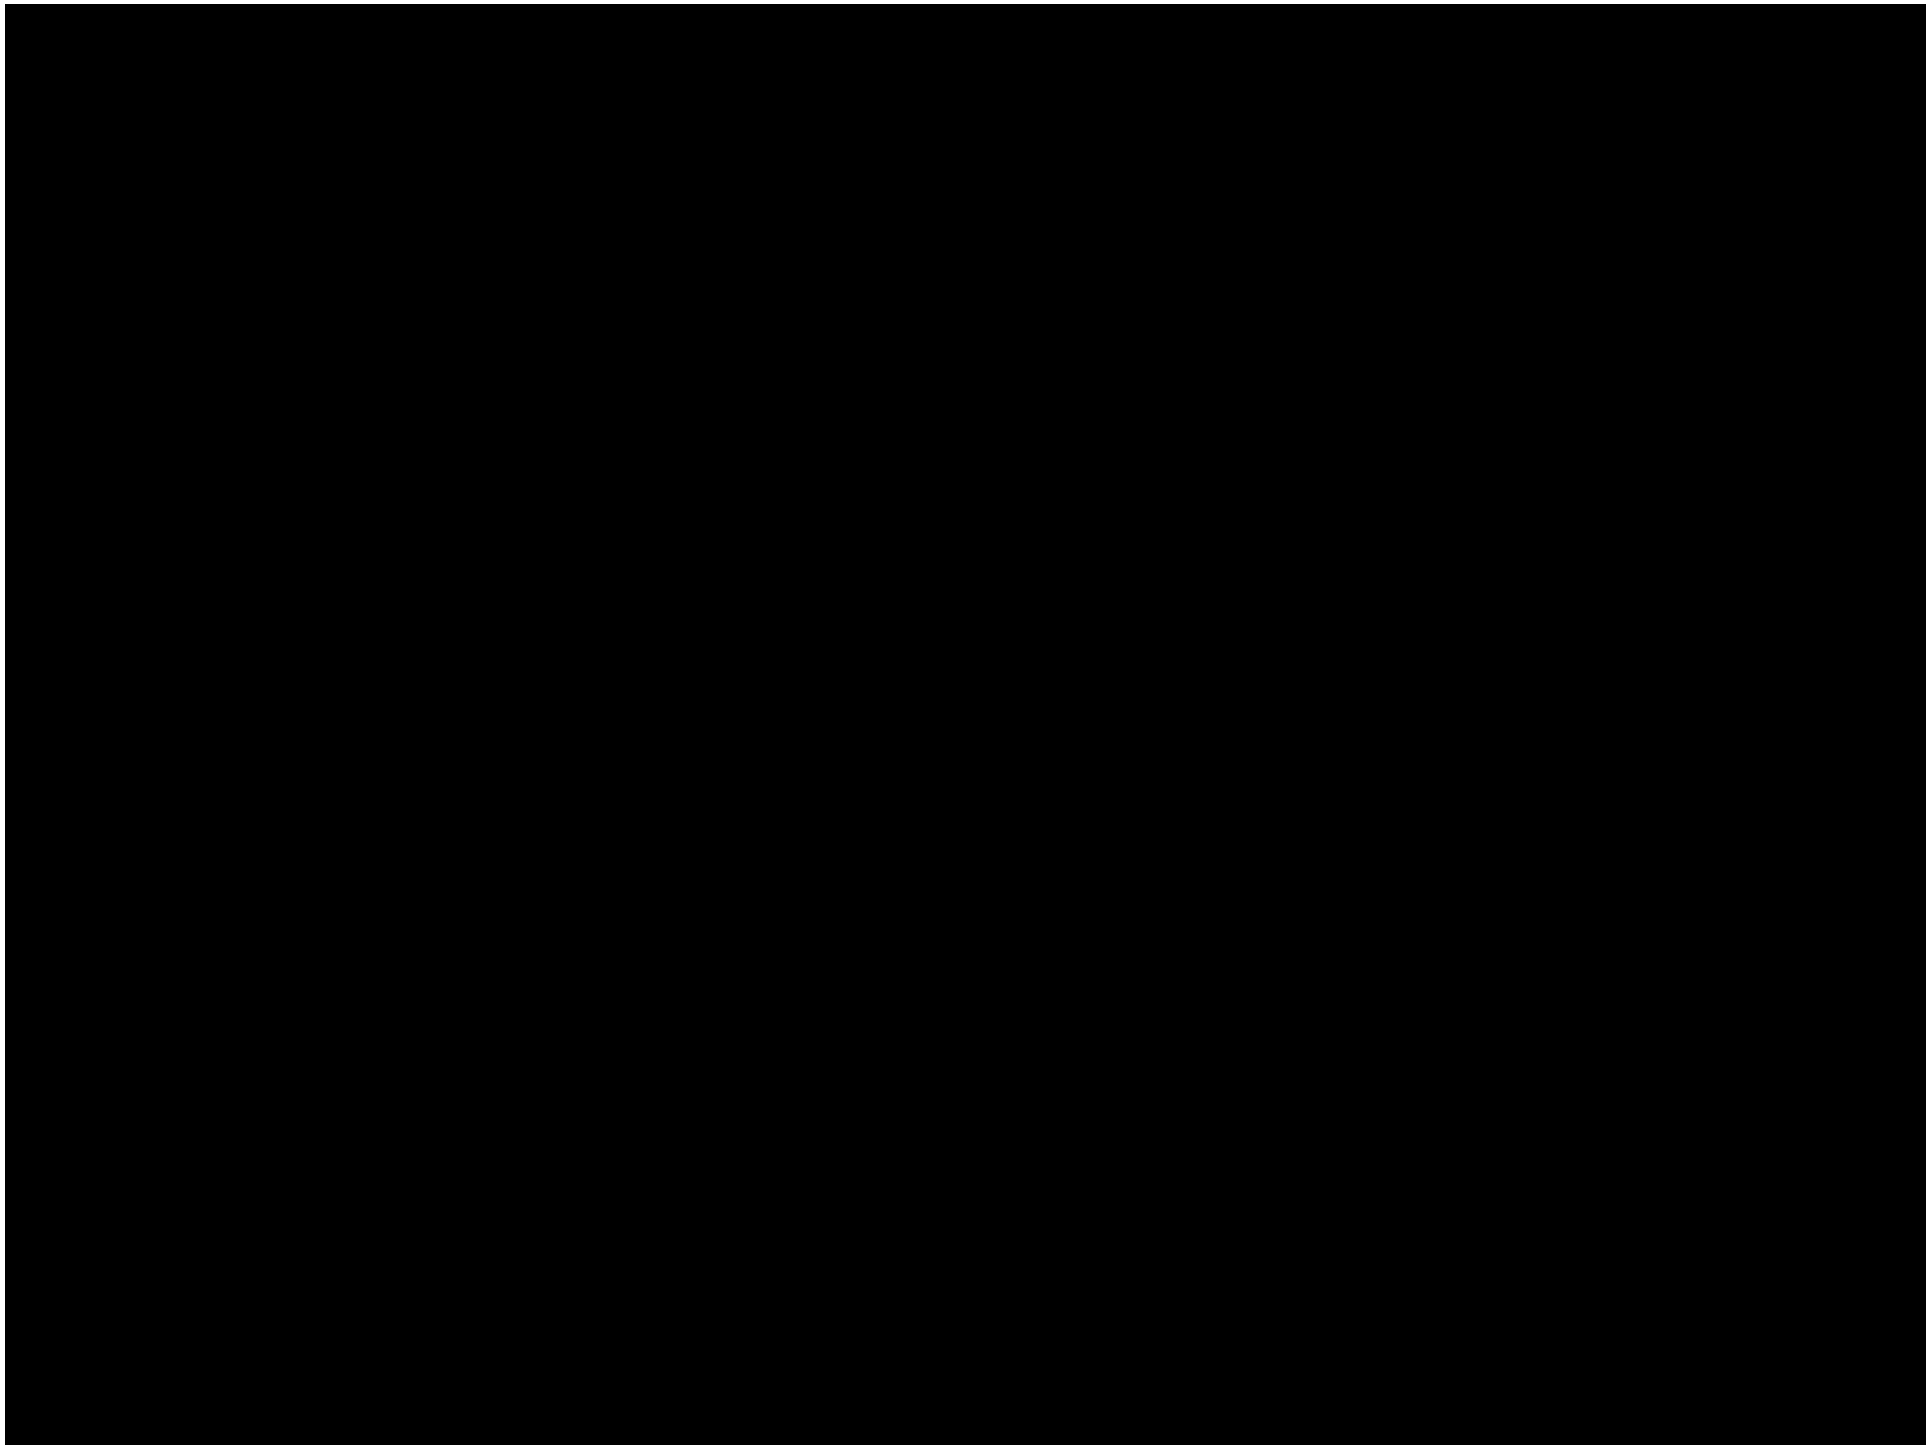

Some ISP Examples
Level 3

I C

BGP/IGP model used in 05999 542 1 1   83t

The Internet Today (31 August 2010)

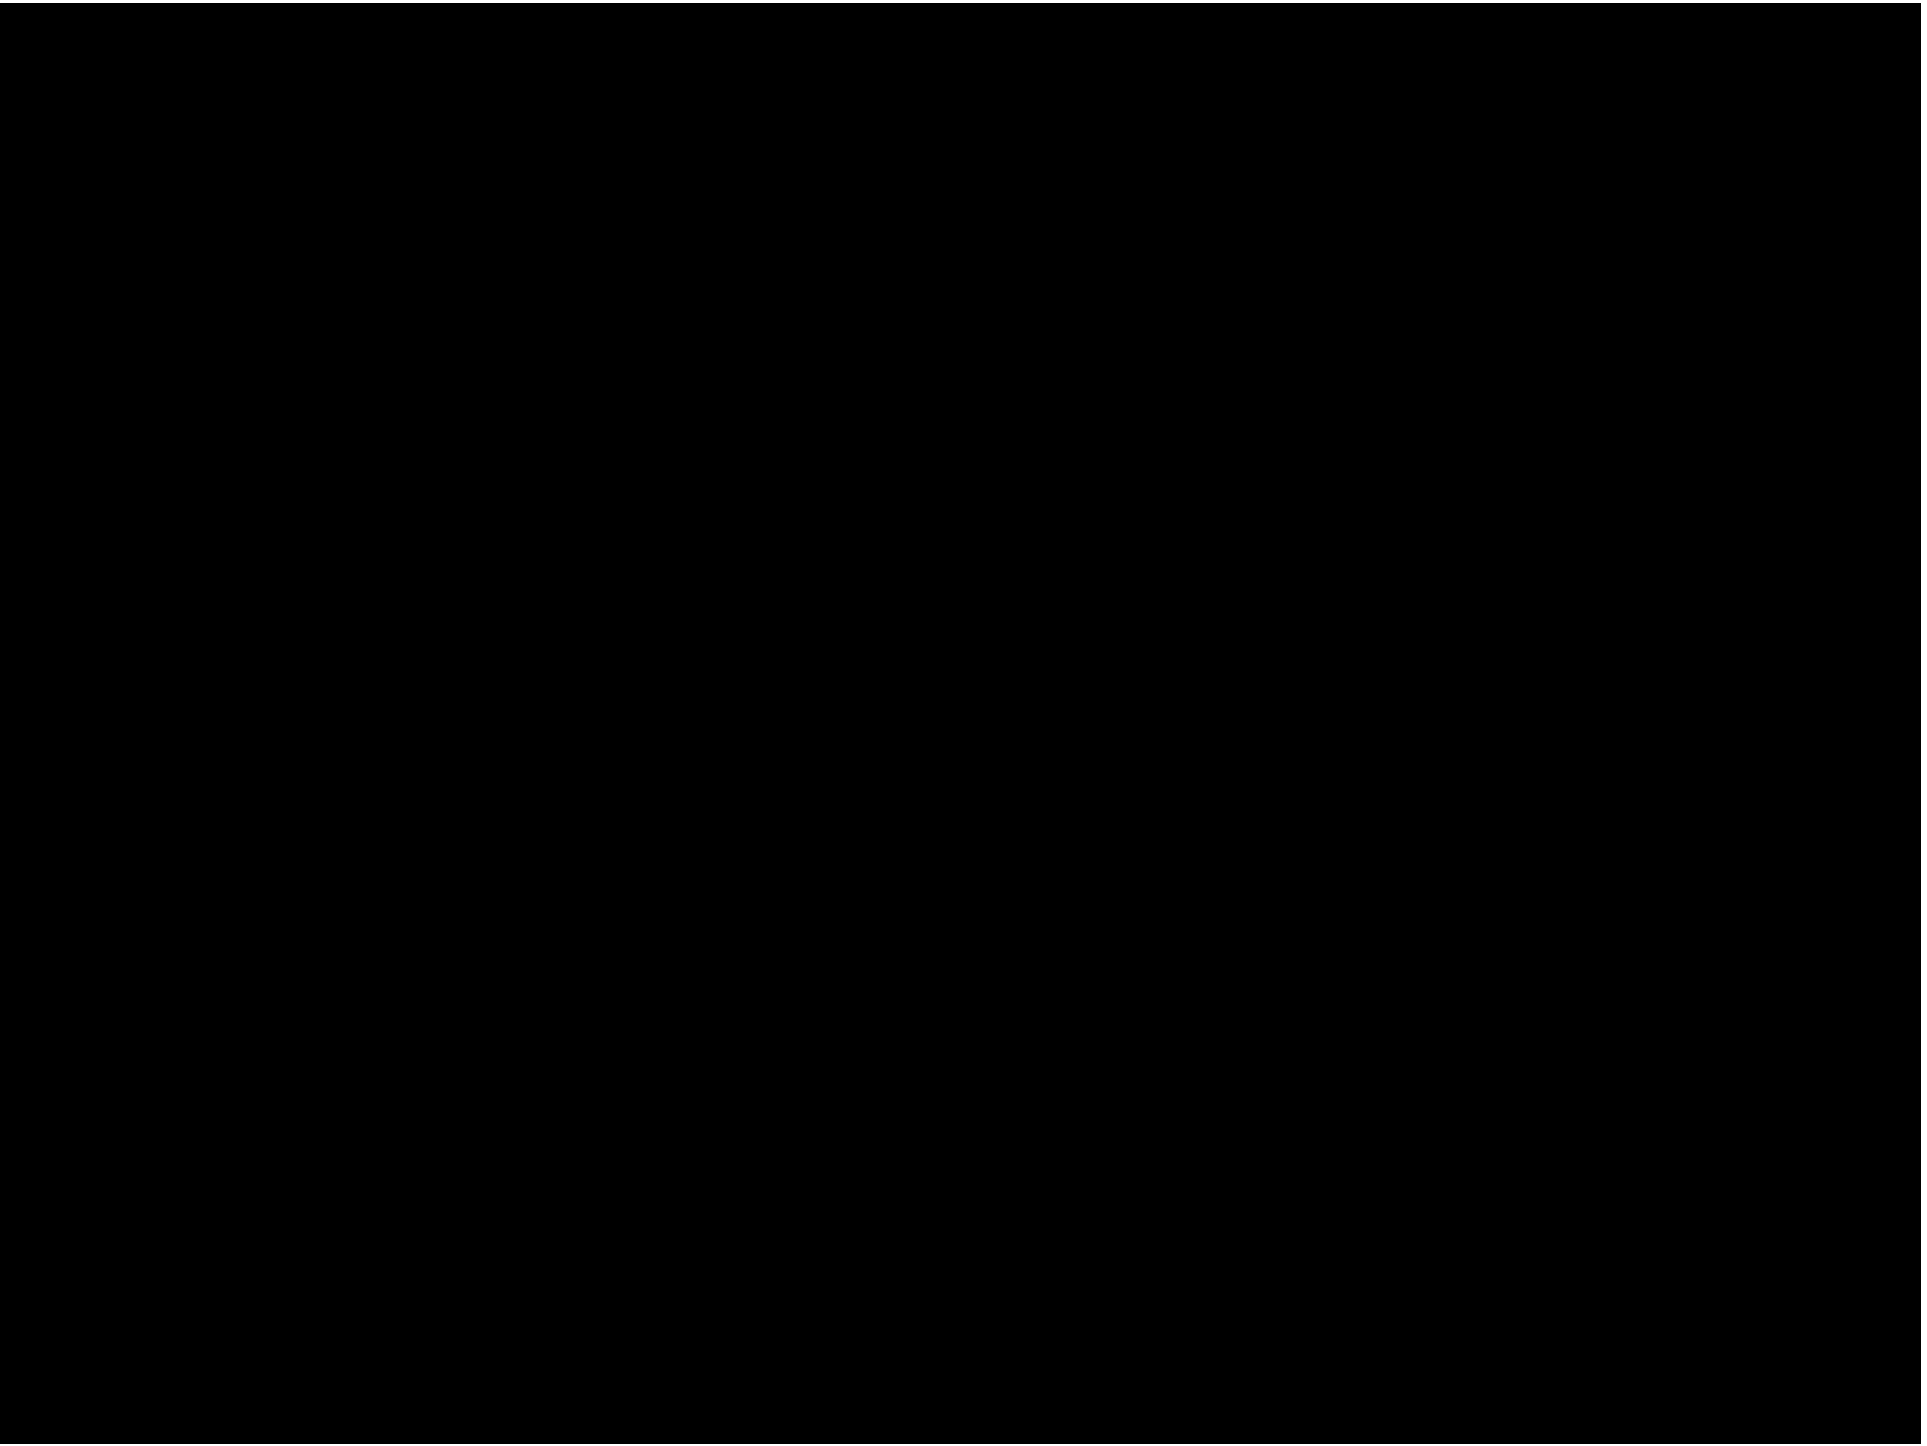

“The New Swamp”
RIR Space

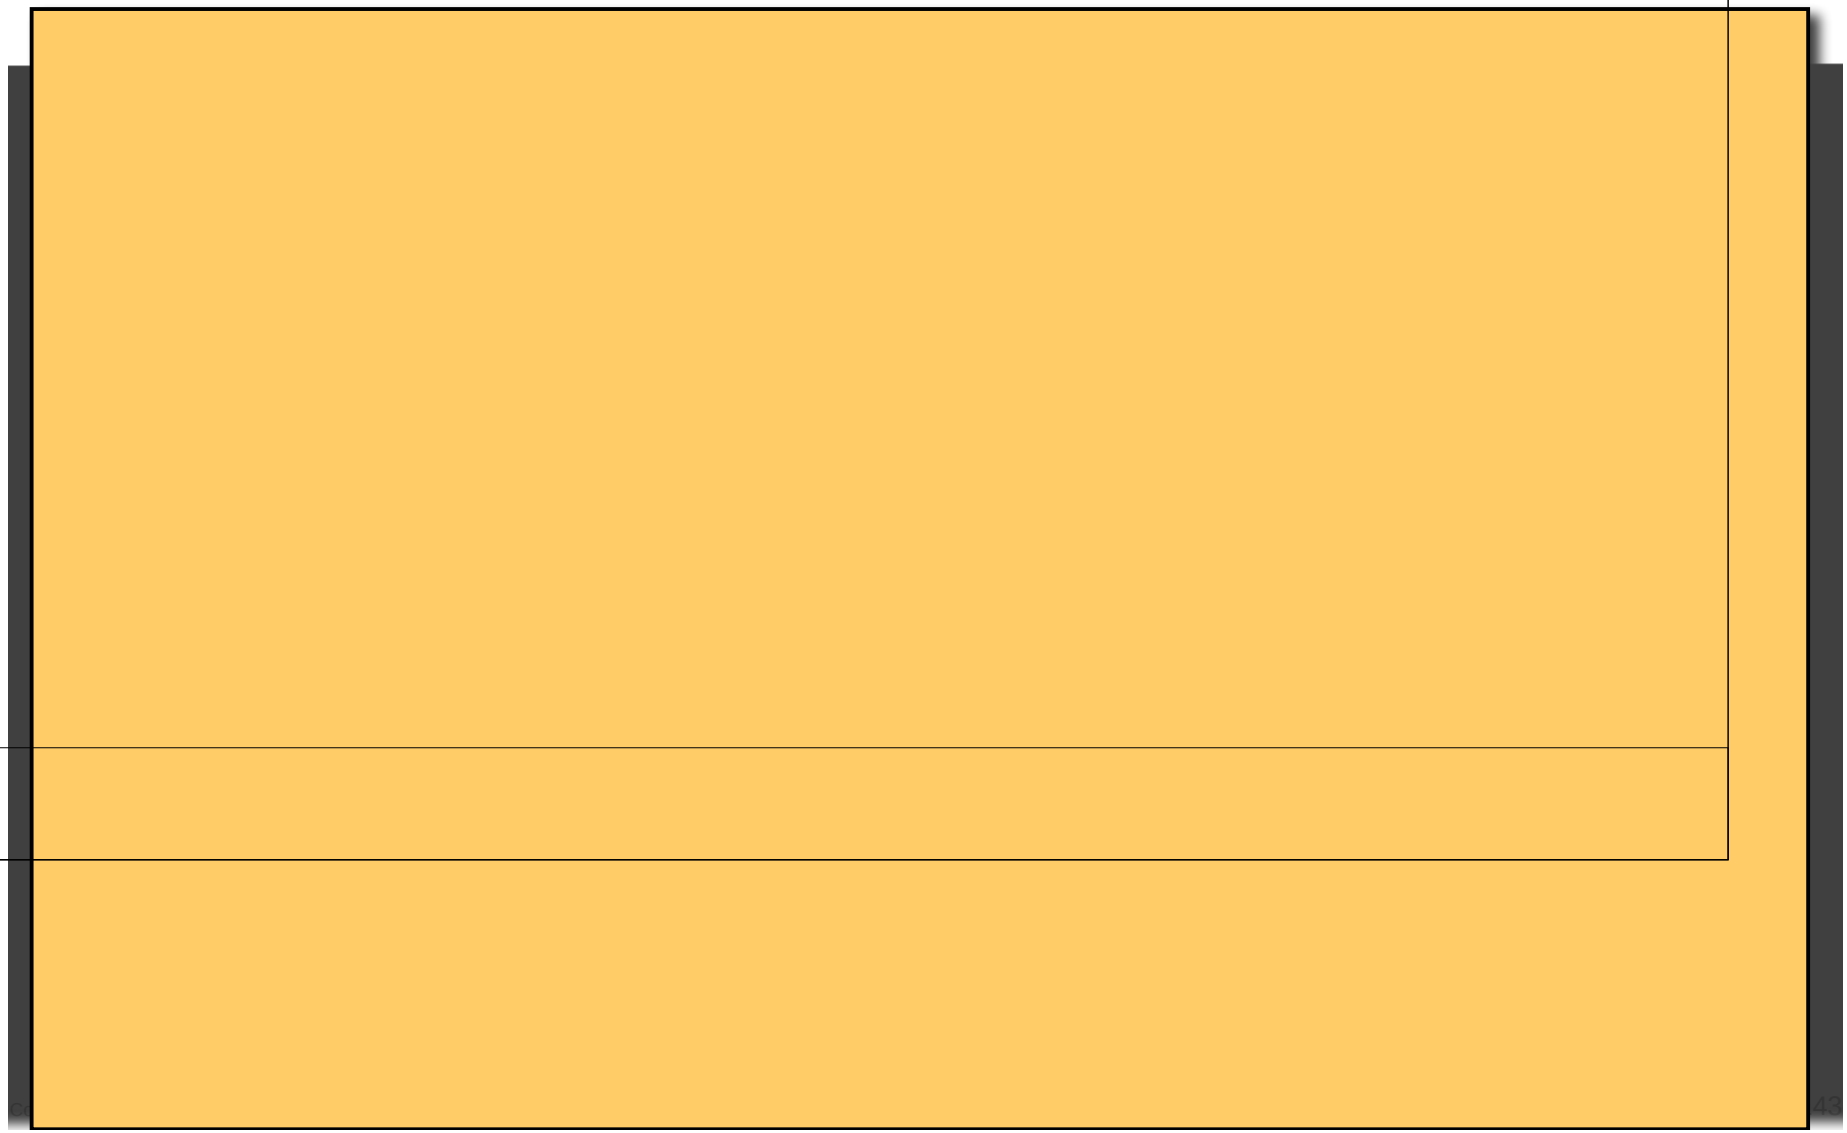

“The New Swamp” Summary

Rest of address space is showing similar deaggregation too

What are the reasons?

- Main justification is traffic engineering

Real reasons are:

- Lack of knowledge
- Laziness
- Deliberate & knowing actions

Importance of Aggregation

Aggregation Summary

Aggregation on the Internet could be MUCH

Receiving Prefixes

Receiving Prefixes:

iBGP and IGPs Reminder!

Limiting AS Path Length

Some announcements have ridiculous lengths of AS-paths:

```
*> 3FFE:1600::/24
```

BGP TTL “hack”

Implement RFC5082 on BGP peerings

(Generalised TTL Security Mechanism)

Neighbour sets TTL to 255

1 0.96 27.64002 719.9999 540 re W n /Cs1 cs 0 0 451 0.827451 0 0