

**The Brothers WISP**



**Route it like it's HOT**

**Greg Sowell Consulting**



# Multihomed BGP Design and Configuration MUM 2012

# What Is BGP

- Border Gateway Protocol is the protocol which makes the core routing decisions on the **Internet**.
- Uses path attributes (PAs) instead of metrics to select best route. Autonomous system path (AS\_PATH) is the default metric.
- Path vector logic – uses AS paths to determine which ASs to hop through to reach destination.
- Uses TCP port 179 to establish neighbor relationships.
- Default Keepalive timer is 60 seconds and default hold timer is 180 seconds.
- iBGP neighbor is an internal neighbor (has the same AS number).
- eBGP neighbor is an external neighbor (has a different AS number).

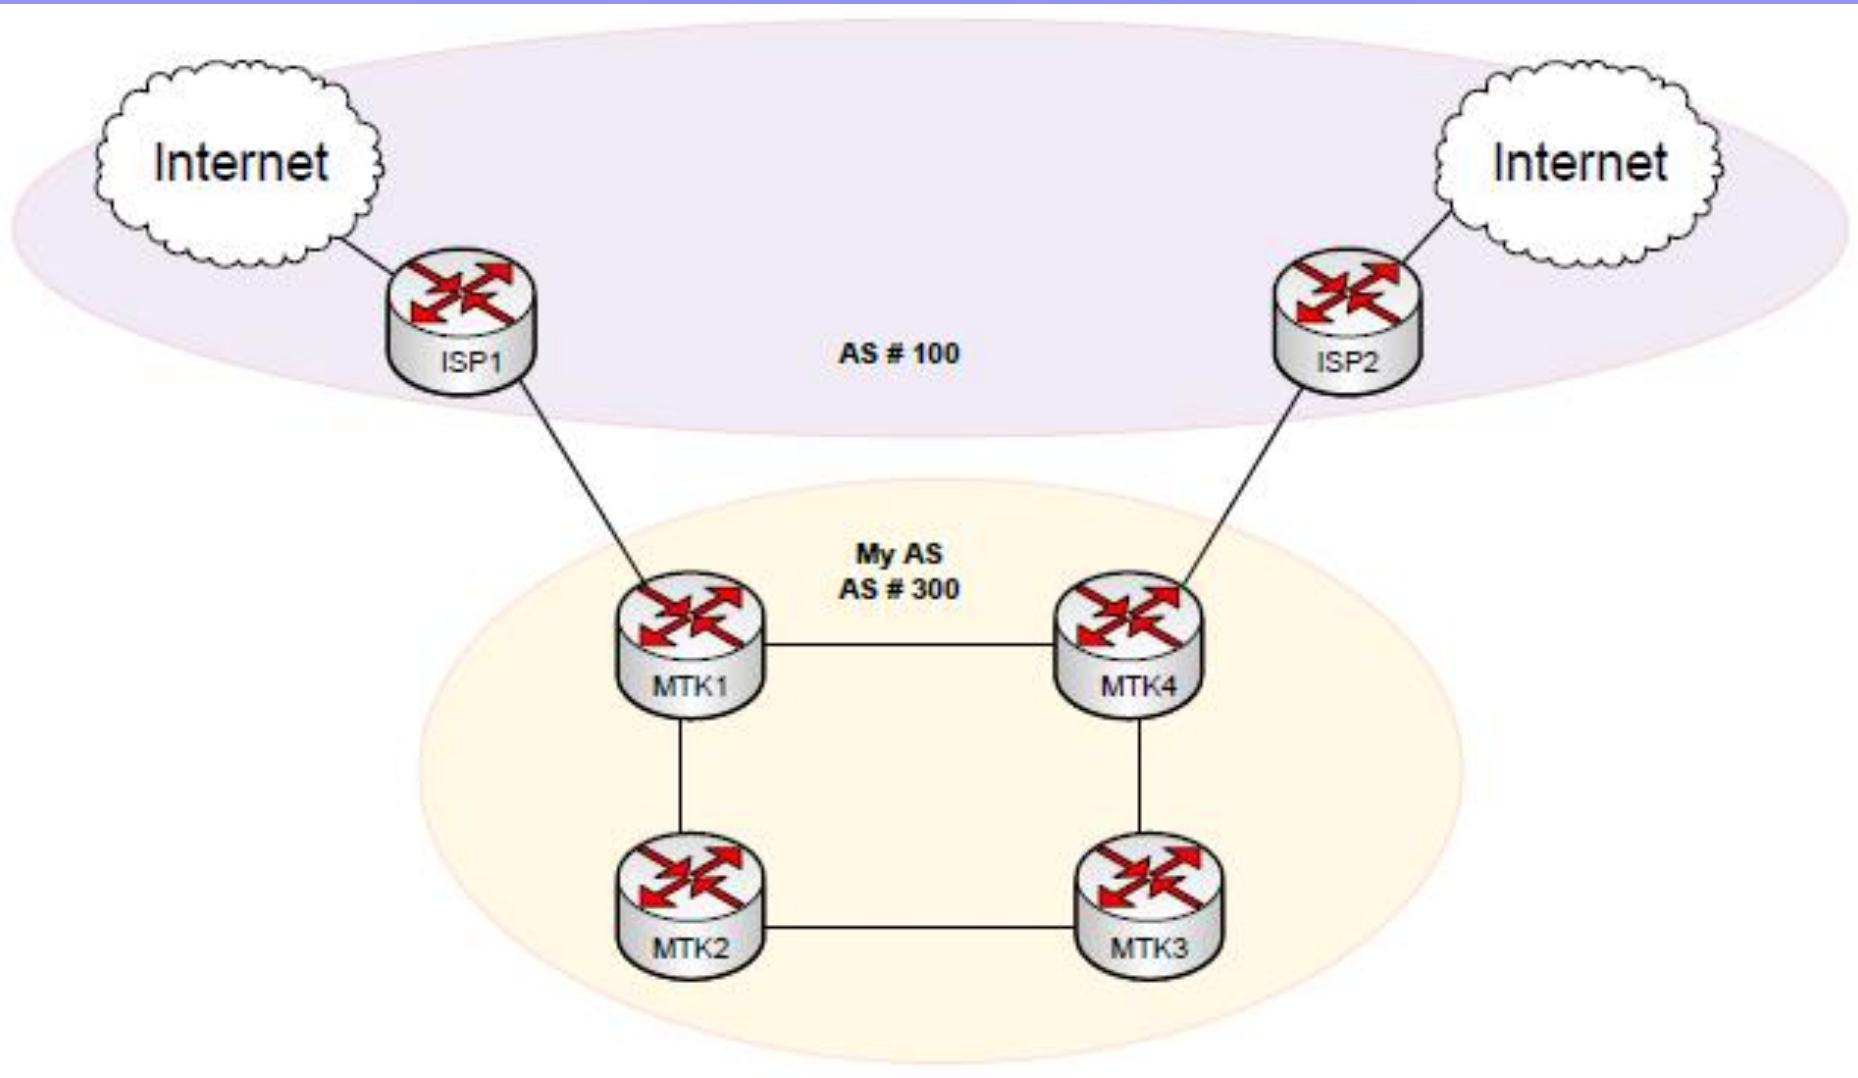
# When To Use BGP

- Do you have public IP addressing assigned to you from your regional provider (ARIN, RIPE)?
- Do you have multiple different upstream ISP connections you want to potentially source the same IP addressing through? “Multi Homed”
- Do you have multiple connections in different areas to the same ISP and want address redundancy?


# BGP VS Internal Routing

- Think of routing as the US postal system.
- BGP is the 18 wheeler trucks that carry mail from city to city. They don't care what street your individual house resides on, just the large city you live in. Equate the city to an autonomous system(AS).
- Once you get inside the city level an internal routing process like OSPF will carry mail to the individual street. In some cases if your city is small enough you could just use static routing.

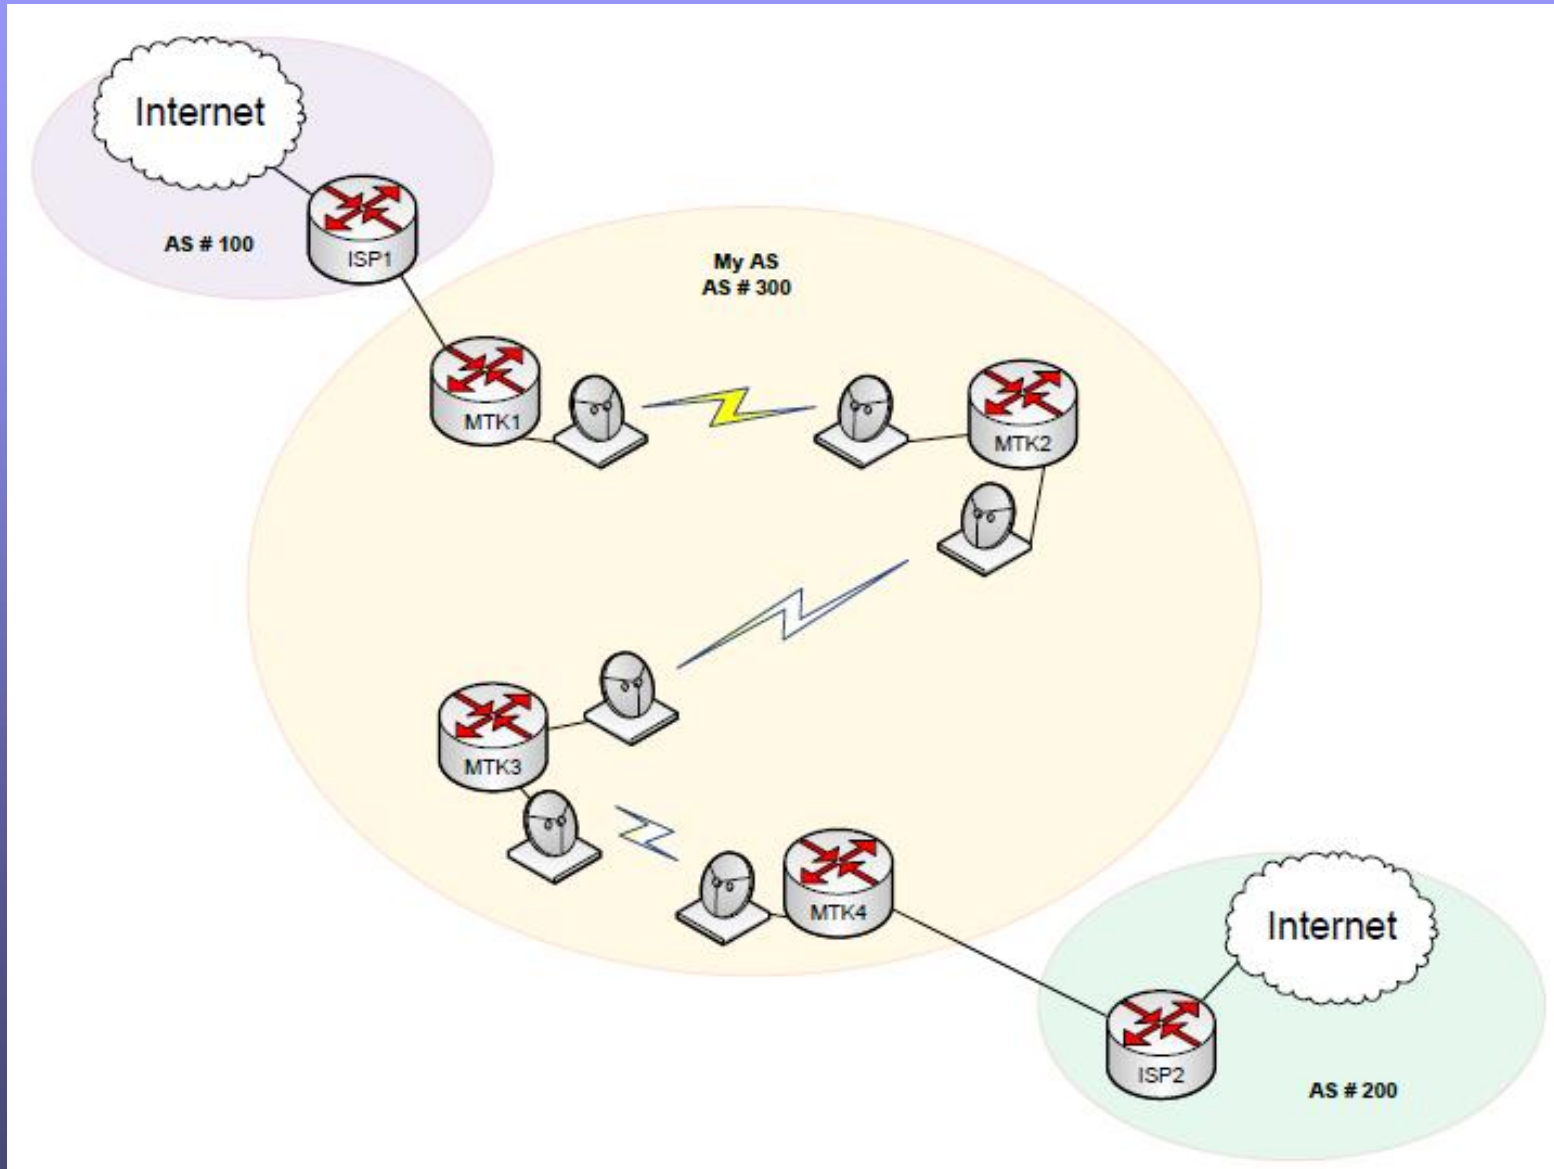
# Multiple Connections To Single ISP



# Multiple ISP Connections



# What Your Network Really Looks Like



# Our Demo Settings

- We have a /22 of addressing. This gives us four /24s to work with. We are breaking our addressing into four /24s because this is the smallest subnet size an upstream ISP will accept.
- We have two upstream providers.
- We have two border routers on opposite ends of our networks.
- We have OSPF as the internal routing process.
- We are accepting only default from our providers.

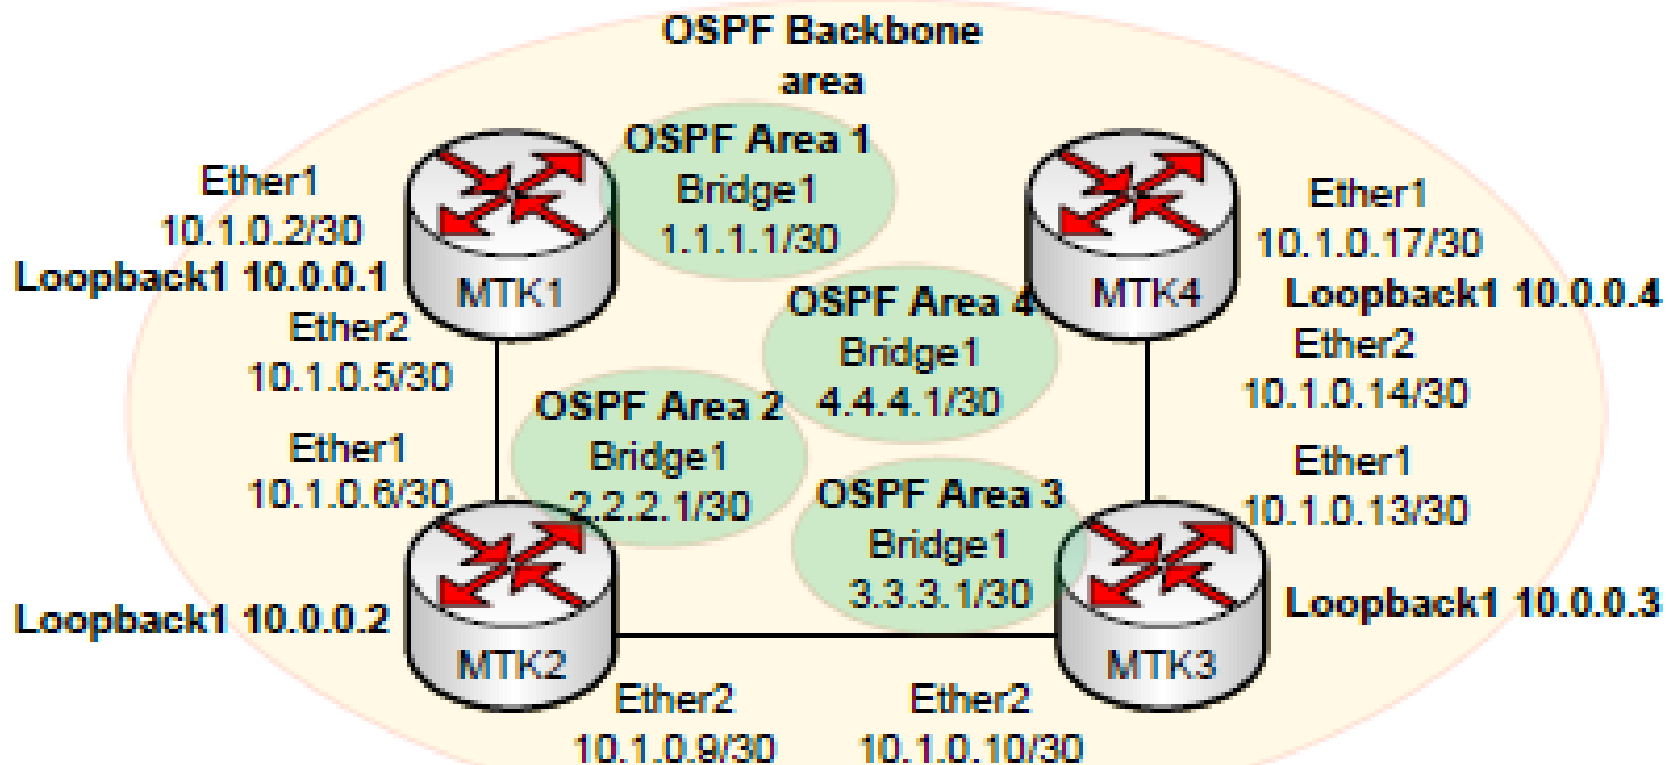


# Start With OSPF

- We use OSPF to carry our interface subnets throughout our internal network.
- We also configure our public addressing in separate OSPF areas to allow for route aggregation and more efficiency throughout our network.
  - Aggregation allows us to summarize all of the smaller subnets into a single /24.
  - Additional efficiency is achieved because when using separate areas any loss of routes causes a partial SPF recalculation within the backbone area so these backbone routers don't take a performance hit.

# OSPF Diagram

- Non-backbone areas are summarized to their respective /24 into the backbone.



# MTK1 Routes

- Note the loopback addresses as well as our public subnets.
- Since we are also summarizing our public addressing a /24 shows up in the route table.
- Create new area

OSPF Area <area1>

Area Name: area1

Instance: default

Area ID: 0.0.0.1

Route List

Routes	Nexthops	Rules	VRF
+	-	✓	✗
	Dst. Address	Gateway	Distance
DAo	1.1.1.0/24		109
DAC	1.1.1.0/30	bridge1 reachable	0
DAo	2.2.2.0/24	10.1.0.6 reachable ether2	110
DAo	3.3.3.0/24	10.1.0.6 reachable ether2	110
DAo	4.4.4.0/24	10.1.0.6 reachable ether2	110
DAC	10.0.0.1	loopback1 reachable	0
DAo	10.0.0.2	10.1.0.6 reachable ether2	110
DAo	10.0.0.3	10.1.0.6 reachable ether2	110
DAo	10.0.0.4	10.1.0.6 reachable ether2	110
DAC	10.1.0.0/30	ether1 reachable	0
DAC	10.1.0.4/30	ether2 reachable	0
DAo	10.1.0.8/30	10.1.0.6 reachable ether2	110
DAo	10.1.0.12/30	10.1.0.6 reachable ether2	110
DAC	192.168.223.0/24	ether4 reachable	0

- Area Range for summarization.
- OSPF network statements.

OSPF Area Range <1.1.1.0/24>

Area: area1

Range: 1.1.1.0/24

Cost: calculated

☒ Advertise

OK

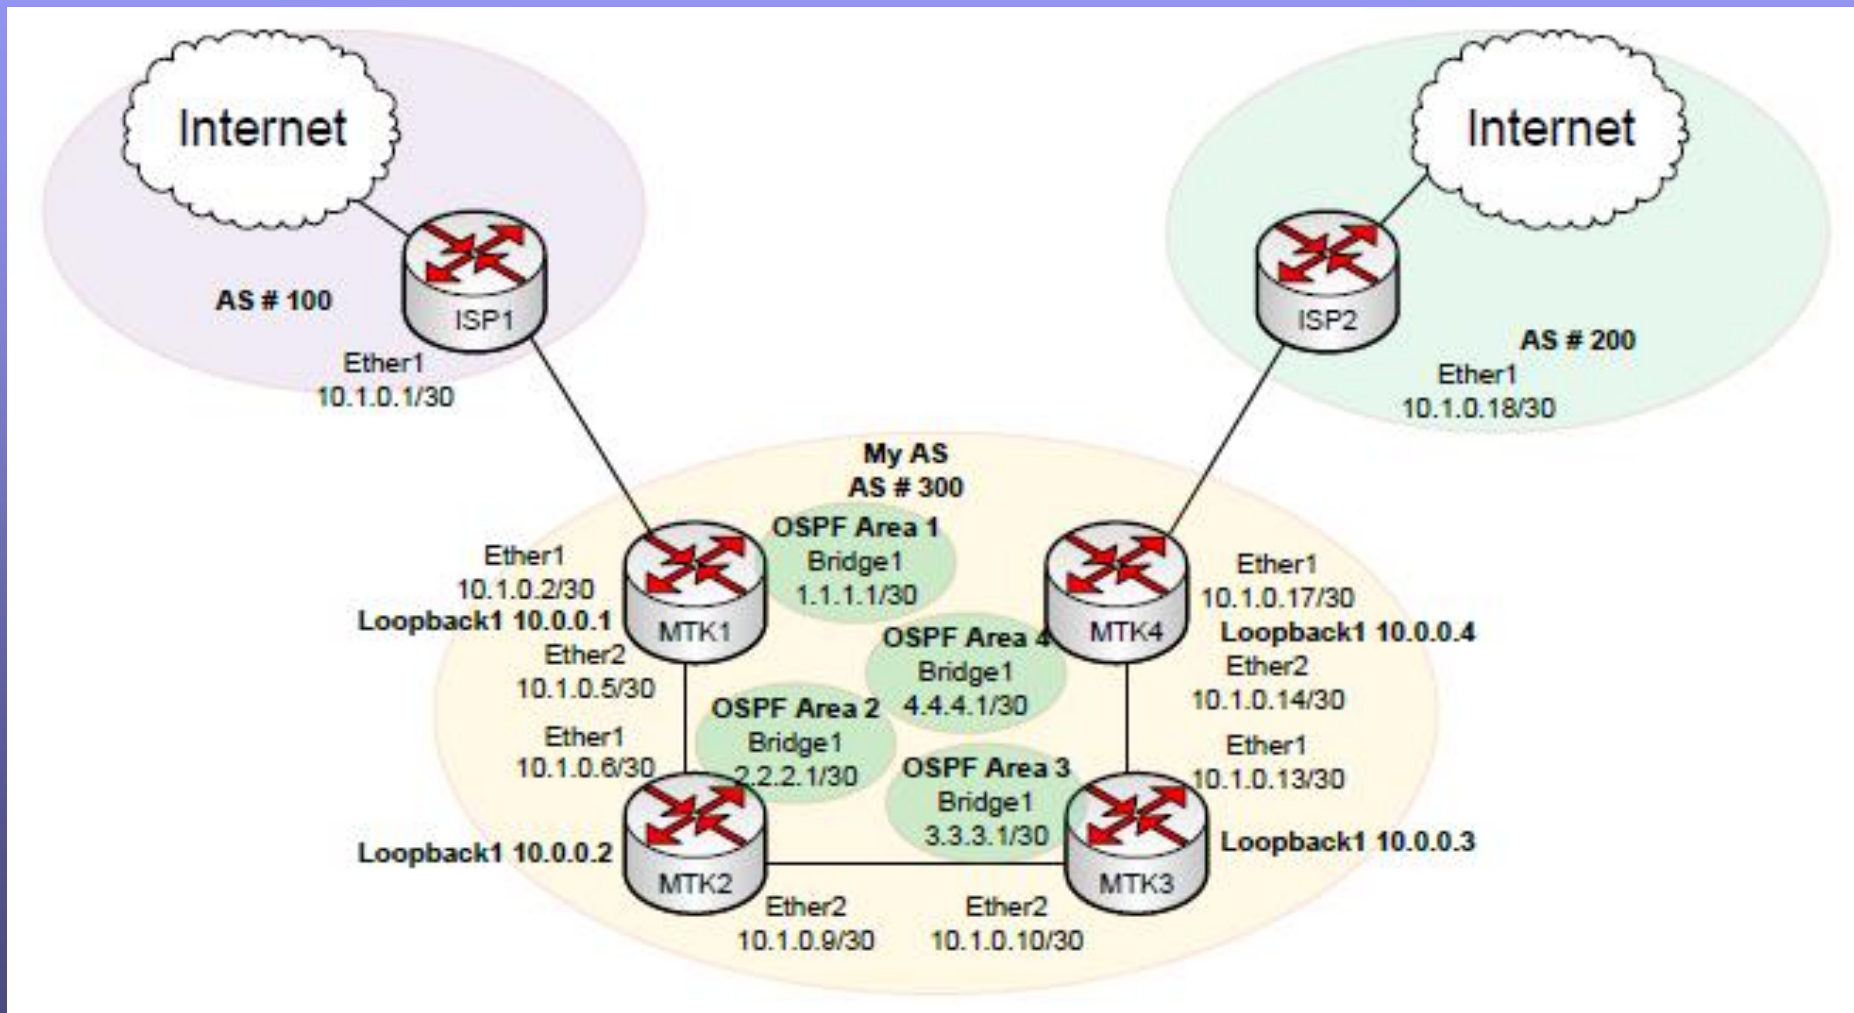
Cancel

Apply

OSPF

Instances	Networks	Areas	Area Range
+	-	✓	✗
	Network	Area	
	1.1.1.0/24	area1	
	10.0.0.0/24	backbone	
	10.1.0.4/30	backbone	

# BGP Diagram

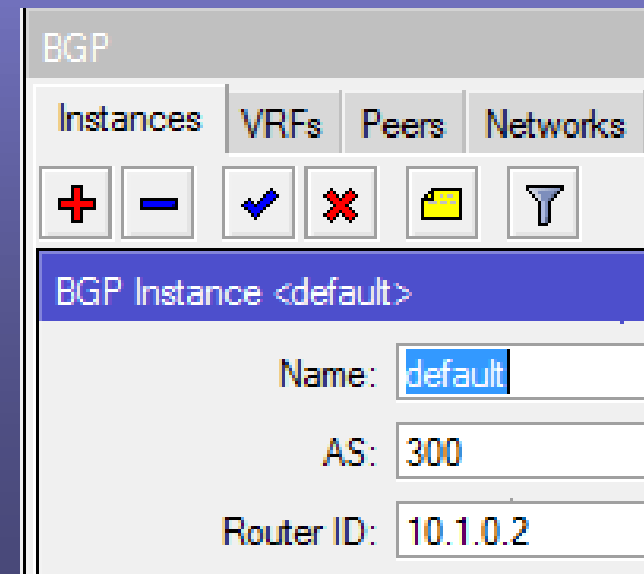


# Configure BGP Peers

- Configure Instance
- Configure Peer
- Configure Networks

# Configure Instance

- Edit the default instance.
- Set the AS# to that of your organization.
  - An Autonomous System (AS) is a collection of connected Internet Protocol (IP) routing prefixes under the control of one or more network operators. A unique AS Number is allocated to each AS for use in BGP routing. AS numbers are important because the ASN uniquely identifies each network on the Internet.
- Set Router ID. This will be Auto configured, but I prefer to Statically set it.

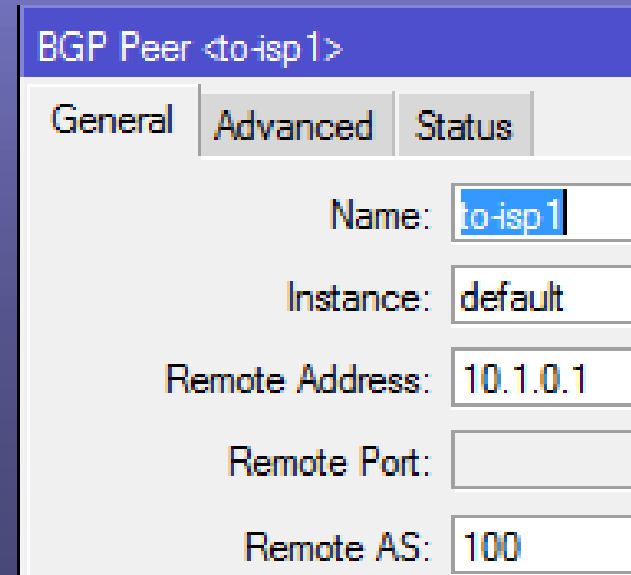


The screenshot shows the Mikrotik WinBox interface for configuring BGP. The 'BGP' window is open, and the 'Instances' tab is selected. Below the tabs are icons for adding (+), removing (-), enabling (checkmark), disabling (X), saving (floppy disk), and applying (funnel). The configuration for the 'BGP Instance <default>' is shown below:

Name:	default
AS:	300
Router ID:	10.1.0.2

# Configure Peer

- Add a peer and give it a descriptive name.
- In standard ISP BGP configurations we peer with the directly connected IP address. We set the Remote Address field to that of the next hop router.
- Set the Remote AS entry as that of our peering ISP.



The screenshot shows the 'BGP Peer' configuration window in Mikrotik WinBox. The window title is 'BGP Peer <to-isp1>'. There are three tabs: 'General' (selected), 'Advanced', and 'Status'. The 'General' tab contains the following fields:


Field	Value
Name	to-isp1
Instance	default
Remote Address	10.1.0.1
Remote Port	
Remote AS	100


# Configure Networks


- The network statements will introduce additional routes into the BGP instance.
- Since we want the ability to advertise all of the networks we add network statements for each.
- Note that we have synchronization enabled for the network statements. This means the exact route must currently exist in the route table if it is to be advertised.


BGP


InstancesVRFsPeersNetworks
















	Network		Synchronize
	 1.1.1.0/24		yes
	 2.2.2.0/24		yes
	 3.3.3.0/24		yes
	 4.4.4.0/24		yes









# Inbound BGP Routes: Full Vs. Default

- The full internet route table is up around 380,000 routes, but do you really need the full route table?
  - The full table is really useful for determining outbound routing path if you have a large volume of ISP connections.
  - If you have a standard WISP network, then your outbound is generally 10% of what your inbound traffic rate is and thus less of a concern.
- Accept just default route. Since we are a standard home user based ISP we are concerned about inbound routing and balancing so we should be able to simply accept the default route from our ISPs. This will allow us to keep our BGP tables small and converge quickly.

# Learned BGP Routes

- We have instructed our upstream providers to only send us the default route, but it seems they are sending us more addressing.
- Now that we are successfully receiving the default route we need to inject this throughout our internal infrastructure.

Route List		
Routes	Nexthops	Rules VRF
		
		
BGP	is	yes
	Dst. Address	Gateway
DAb	0.0.0.0/0	10.1.0.1 reachable ether1
DAb	6.6.6.0/24	10.1.0.1 reachable ether1

# Redistribute Default

- In our border router's OSPF instance, we need to set Redistribute Default Route to if installed (as type 1).
- Setting it as type 1 will take the external cost of the route along with the cost to reach the injecting router.

The screenshot shows the Mikrotik WinBox OSPF configuration window. The 'Instances' tab is selected, displaying a table of OSPF instances. The 'default' instance is highlighted, showing its Router ID as 10.0.0.1 and its status as 'Running'. Below the table, the configuration for the 'default' instance is shown, with the 'Redistribute Default Route' option set to 'if installed (as type 1)'.

Name	Router ID	Running
default	10.0.0.1	yes

OSPF Instance <default>

General Metrics MPLS Status

Name: default

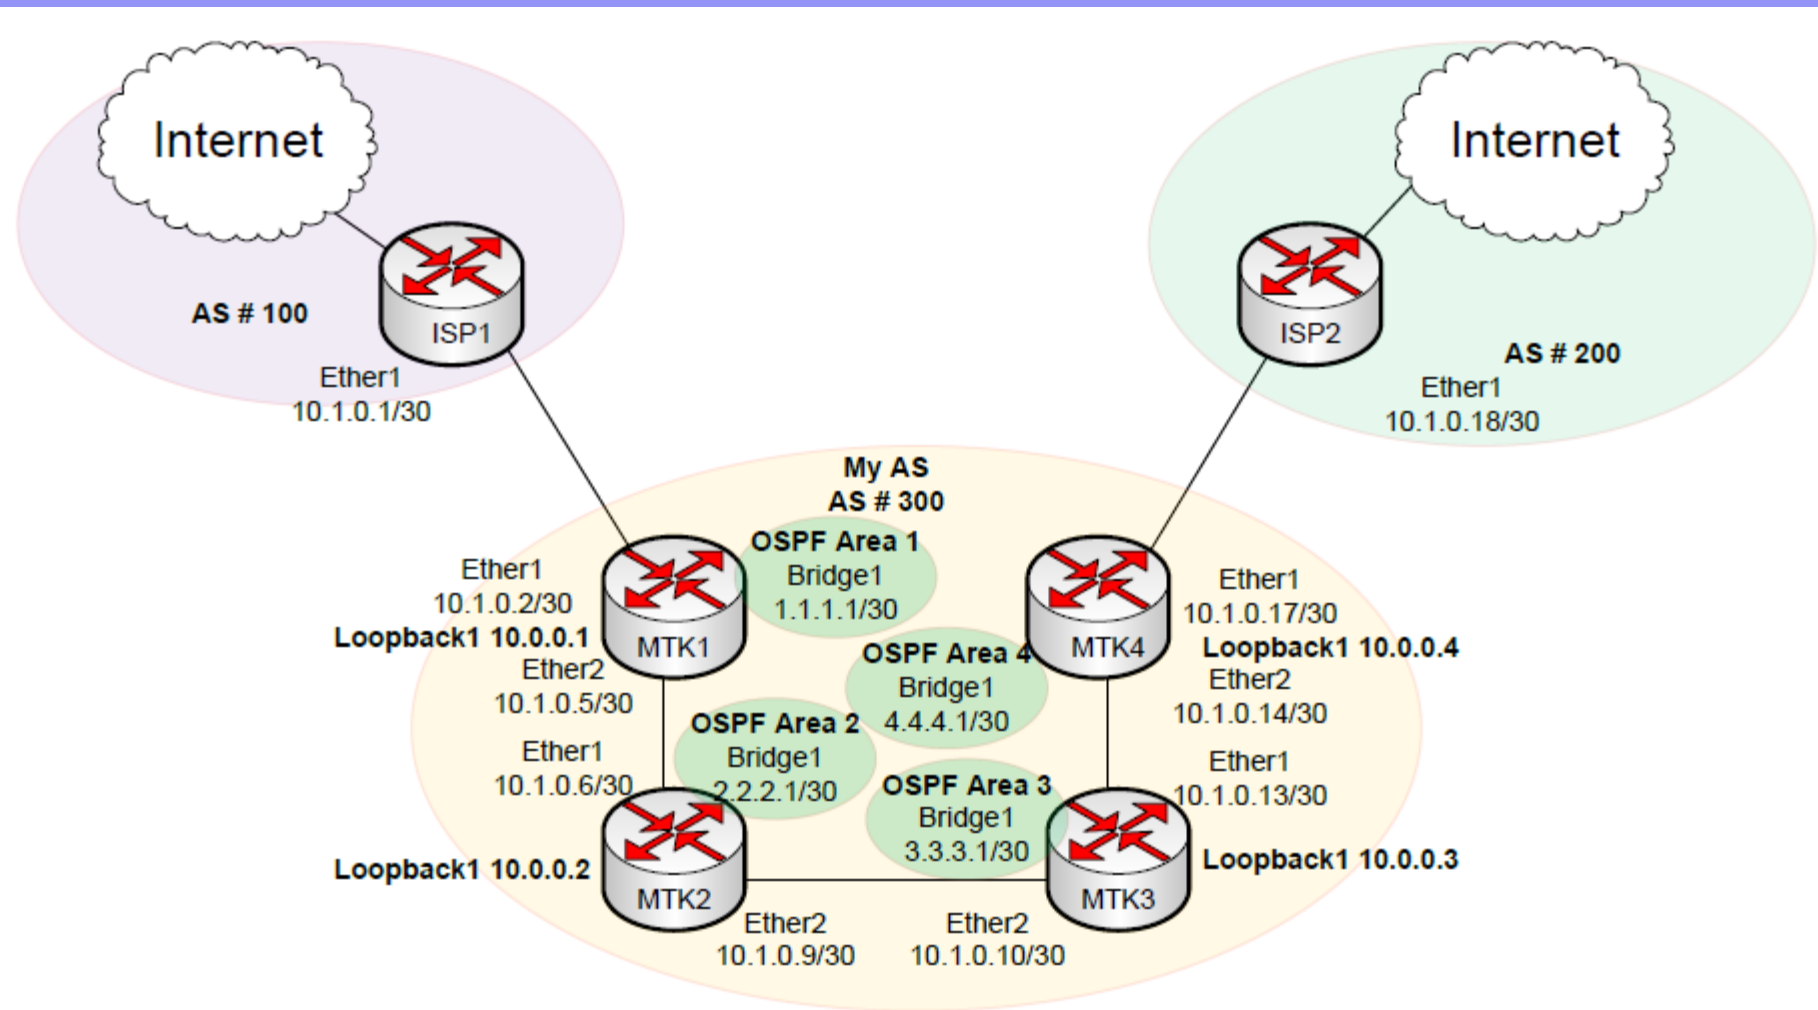
Router ID: 10.0.0.1

Redistribute Default Route: if installed (as type 1)

# Traffic Flow Outbound

- Traffic is currently split for outbound transit. Using standard OSPF costs the traffic is sent out the closest gateway.
  - MTK1 and 2 go out ISP1.
  - MTK3 and 4 go out ISP2.
- If there is a failure to either ISP the default route from that BGP session is removed and thus is no longer redistributed into OSPF. The remote default route is then accepted into all remaining routers ensuring traffic is still able to flow to the internet.

# Current Configuration

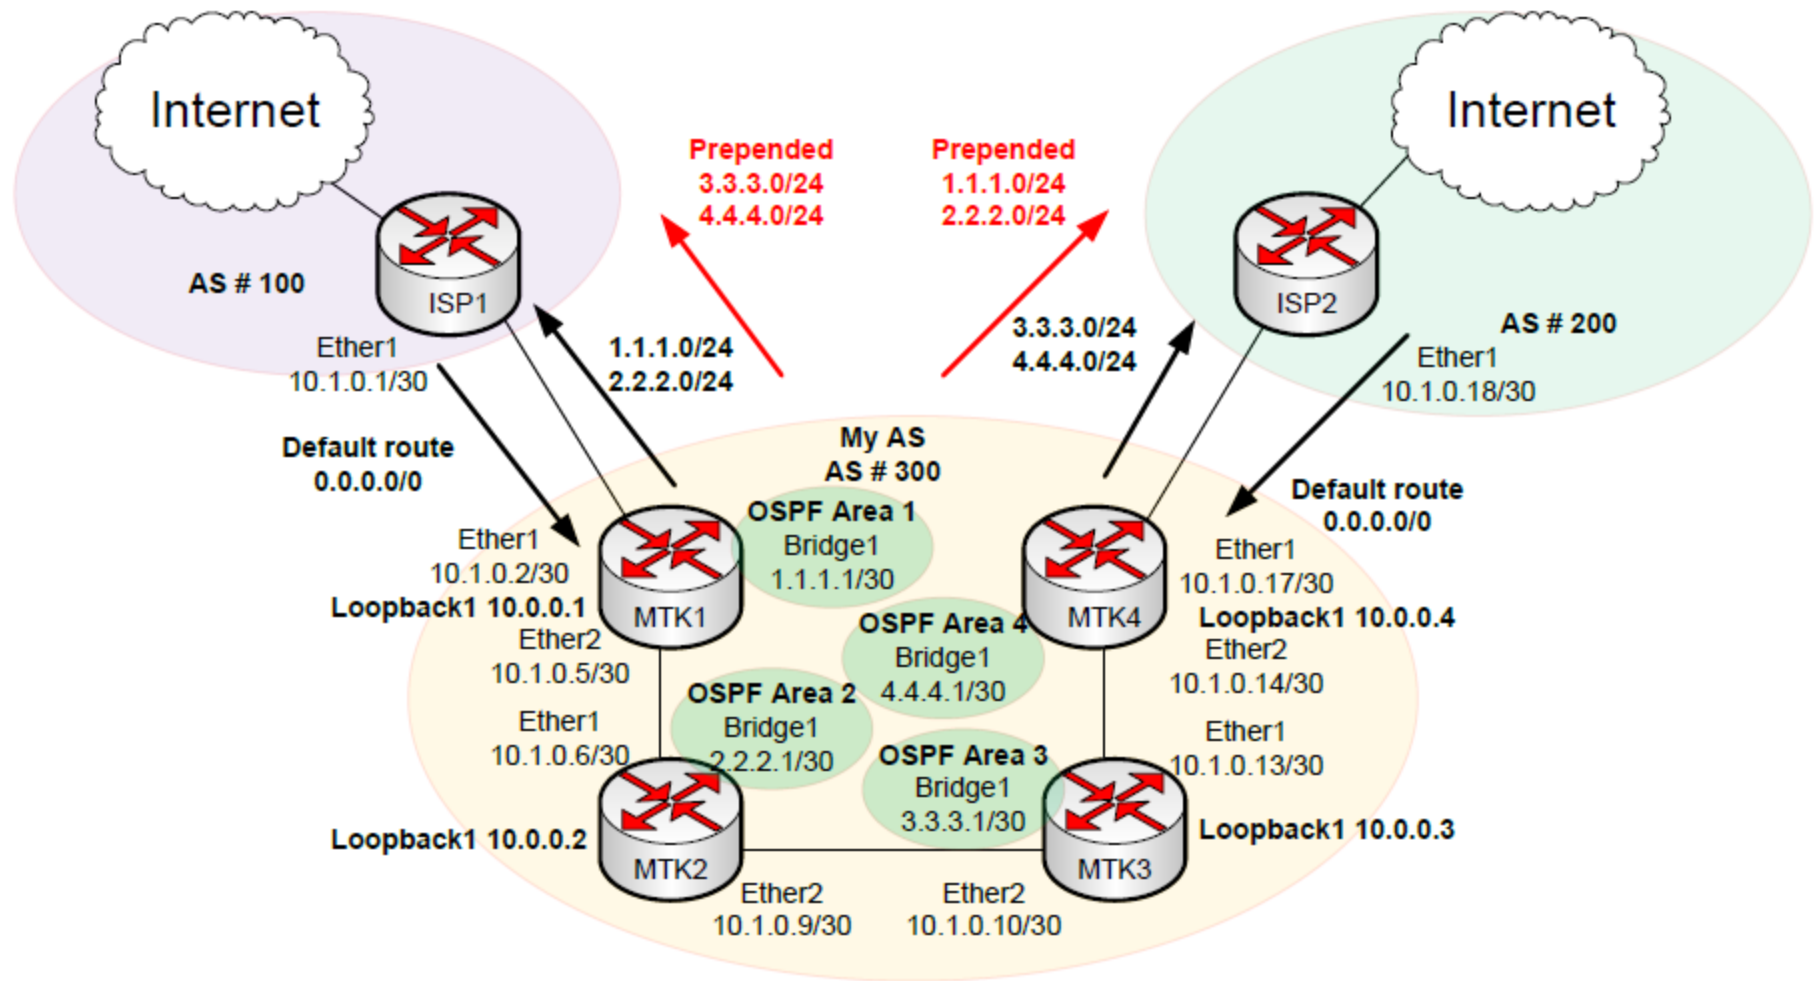


# Verify OSPF Default

# Route Flow and Filtering

- Currently we are not doing any filtering or manipulation of our outgoing routes. This means that whichever ISP happens to have the shortest AS path to reach these subnets will carry the traffic.
- We want subnet 1.1.1.0/24 and 2.2.2.0/24 to be homed out of ISP1 with failover to ISP2.
- We also want subnet 3.3.3.0/24 and 4.4.4.0/24 to be homed out of ISP2 with failover to ISP1.
- We also want to ensure that we are only advertising out our 4 public subnets and nothing else.
- Currently we are also accepting all routes from our ISPs. We want to limit these routes to just the default.

# Current Configuration





# Manipulating BGP Advertisements

- We use filters to manipulate our BGP advertisements. This will allow us to adjust what we send out and in.
- In these filters we will also use the concept of prepending. In essence we will artificially add additional AS hops to the AS\_path attribute. This will make a specific route seem farther away and thus less desirable.

# Inbound Filter

- We will allow the default route in and deny all others.

Route Filter <0.0.0.0/0>

Matchers BGP Actions BGP

Chain:

Prefix: ☐ 0.0.0.0/0

Route Filter <0.0.0.0/0>

Matchers BGP Actions BGP

Action:

Route Filter <0.0.0.0/0>

Matchers BGP Actions BGP

Chain:

Prefix: ☐ 0.0.0.0/0

Prefix Length: ☐ 0-32

Route Filter <0.0.0.0/0>

Matchers BGP Actions BGP

Action:

Route Filters

#	Chain	Prefix	Prefix Length	Protocol	BGP AS Path	Action
0	isp 1-in	0.0.0.0/0				accept
1	isp 1-in	0.0.0.0/0	0-32			discard

BGP Peer <to-isp1>

General Advanced Status

Name:

Instance:

Remote Address:

Remote Port:

Remote AS:

TCP MD5 Key:

Nexthop Choice:

☐ Multihop

☐ Route Reflect

Hold Time:

Keepalive Time:

TTL:

Max Prefix Limit:

Max Prefix Restart Time:

In Filter:

# Verify Inbound BGP Filter

# Outbound Filter

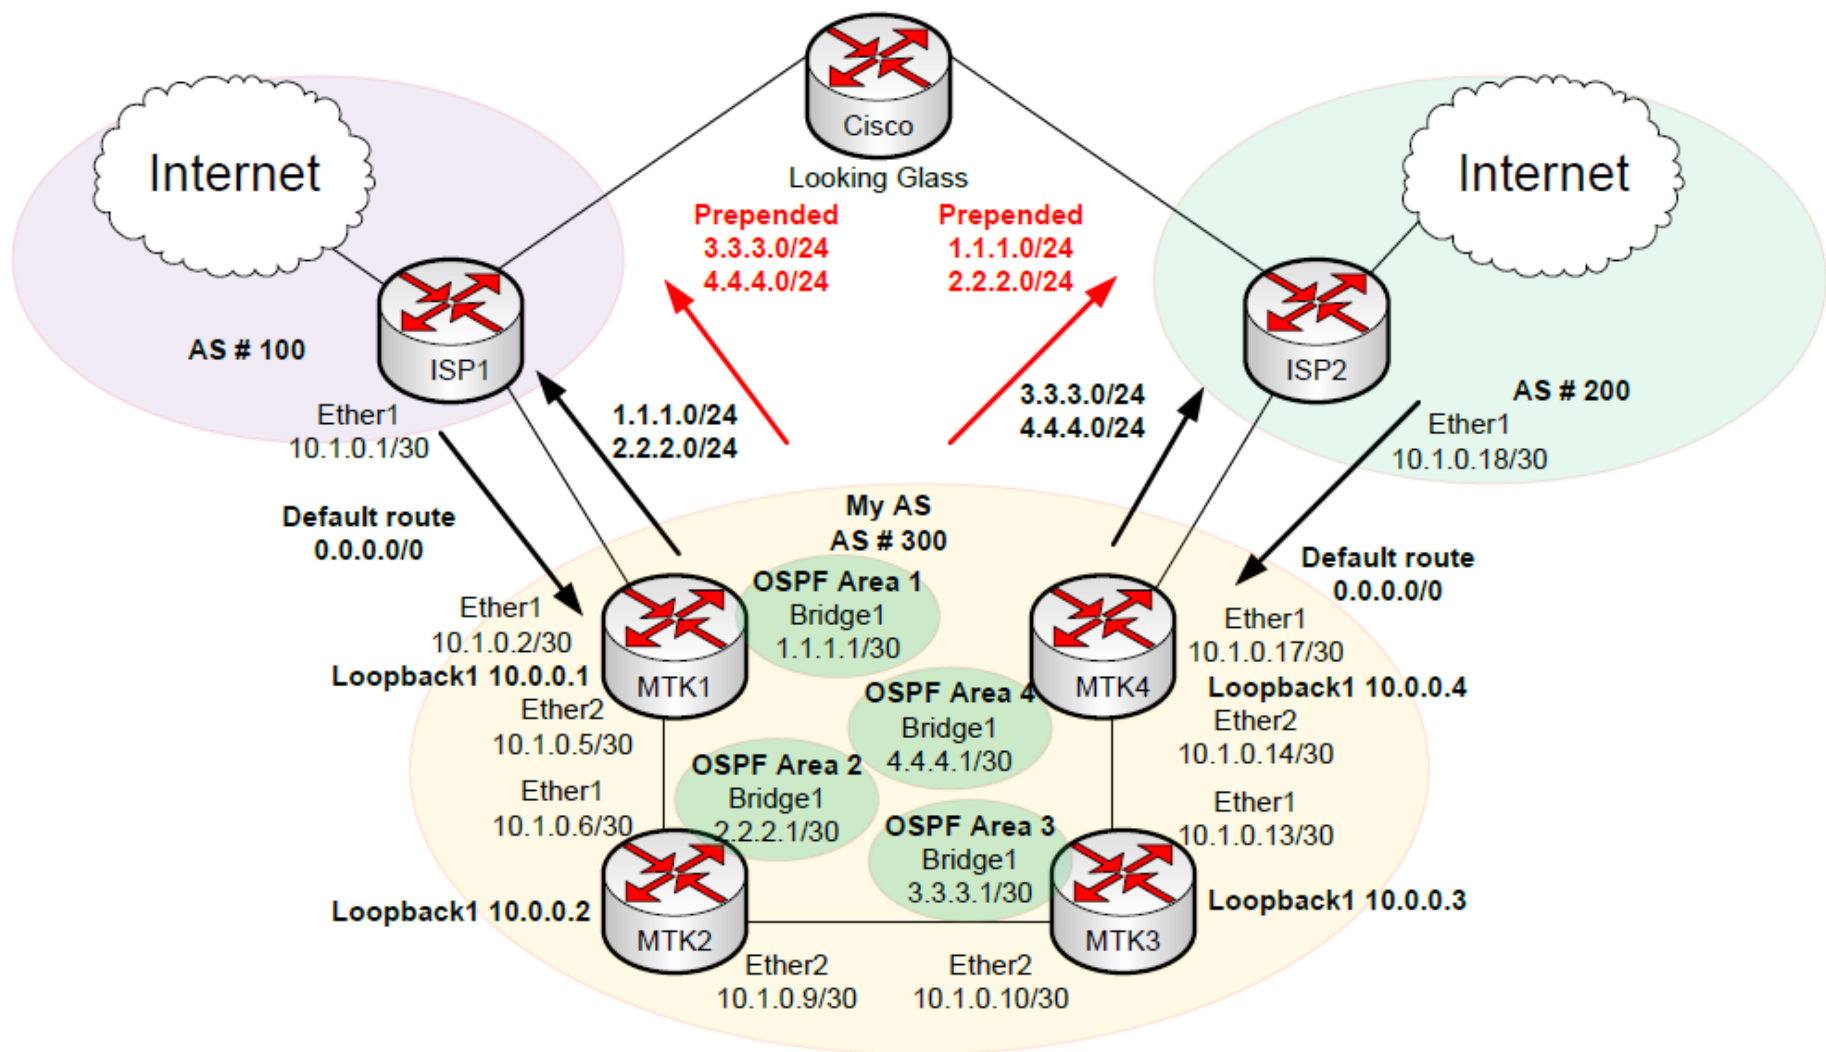
- We will allow 1.1.1.0 and 2.2.2.0/24 out.
- We will allow 3.3.3.0/24 and 4.4.4.0/24 out and also prepend our AS number 3 times.
- All other routes will be denied.

The image displays six Mikrotik WinBox windows for configuring Route Filters. Each window has tabs for 'Matchers', 'BGP', 'Actions', and 'BGP Actions'.

- Route Filter <1.1.1.0/24>**: Chain: `isp1-out`, Prefix: `1.1.1.0/24`.
- Route Filter <2.2.2.0/24>**: Chain: `isp1-out`, Prefix: `2.2.2.0/24`.
- Route Filter <3.3.3.0/24>**: Chain: `isp1-out`, Prefix: `3.3.3.0/24`. BGP Actions: Set BGP Weight: (empty), Set BGP Local Pref.: (empty), Set BGP Prepend: `3`.
- Route Filter <4.4.4.0/24>**: Chain: `isp1-out`, Prefix: `4.4.4.0/24`. BGP Actions: Set BGP Weight: (empty), Set BGP Local Pref.: (empty), Set BGP Prepend: `3`.
- Route Filter <0.0.0.0/0>**: Chain: `isp1-out`, Prefix: `0.0.0.0/0`, Prefix Length: `0-32`. Action: `discard`.

Below the main windows, there are smaller windows showing the 'Action' for each filter: `accept` for 1.1.1.0/24, 2.2.2.0/24, 3.3.3.0/24, and 4.4.4.0/24, and `discard` for 0.0.0.0/0.

# Current Configuration

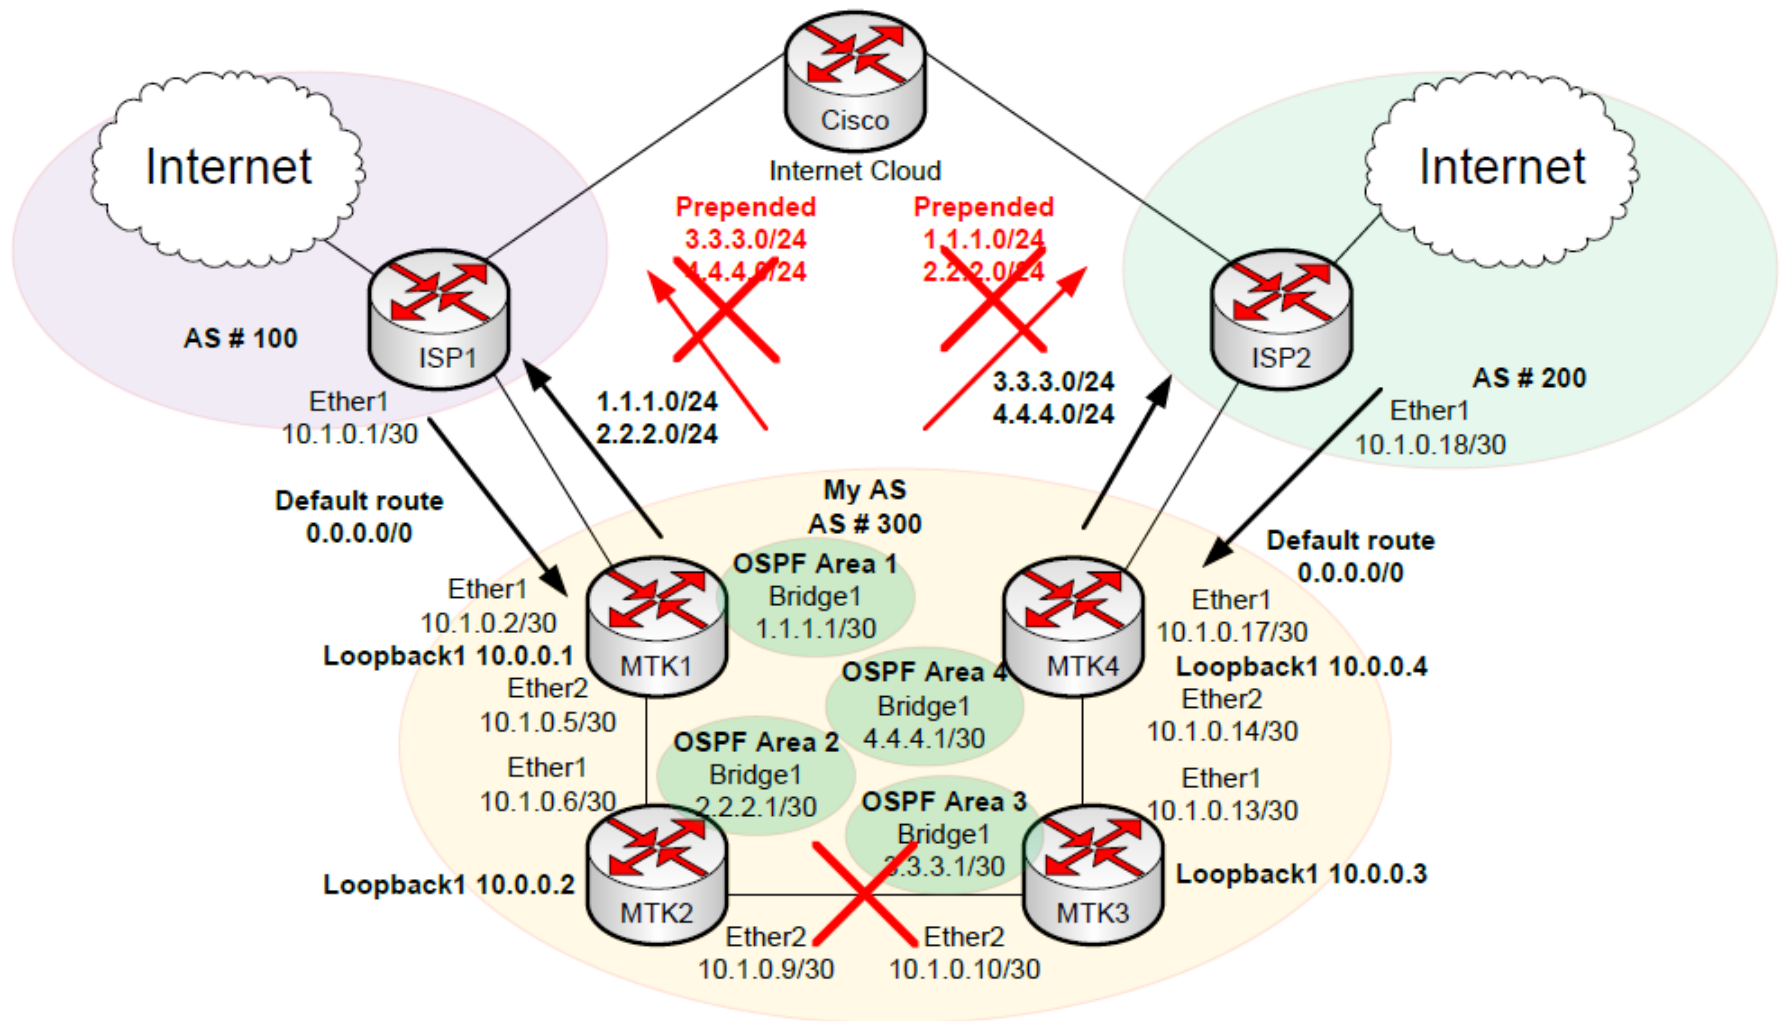


# Verify Outbound BGP Filter

Network	Next Hop	Metric	LocPrf	Weight	Path
* 1.1.1.0/24	192.168.223.212	0	200	300	300 300 i
*>	192.168.223.211	0	100	300	i
* 2.2.2.0/24	192.168.223.212	0	200	300	300 300 i
*>	192.168.223.211	0	100	300	i
*> 3.3.3.0/24	192.168.223.212	0	200	300	i
*	192.168.223.211	0	100	300	300 300 i
*> 4.4.4.0/24	192.168.223.212	0	200	300	i
*	192.168.223.211	0	100	300	300 300 i

# Verify Failover Routing

# Reaching Disconnected Subnets





# Reaching Disconnected Subnets

