



PEERING STRATEGIES, TOOLS & BEST PRACTICES

Dr Aeffendi Hashim

On behalf of Global IP, Core IP Development, TM

MyIX Peering Forum,

10 October 2016, Aloft Hotel

Content

- Overview of TM's network
- Peering Strategies & Tools
 - Top down
 - Bottom up
- Best Practices
 - Inbound filters
 - Outbound filters
 - Aggregation
- MyIX Peering Snapshot



TM Overview

TM's a Converged Communications Service Provider

Others	<ul style="list-style-type: none"> Value Added Service
Layer 3	<ul style="list-style-type: none"> Virtual Private Networks (VPNs) Fixed-Line Voice Services Fixed-Line Video Services (IPTV) Fixed Internet Access Broadband (2.3 million subscribers) WiFi internet Wholesale internet (transit) Mobile Voice (Webe) Mobile Internet (Webe) Internet security
Layer 2	<ul style="list-style-type: none"> Ethernet Access ELAN/ELINE circuits Managed Data Centres
Layer 1	<ul style="list-style-type: none"> Transmission Links (leased circuits) Copper Access Networks Optical Access Networks Data Centres Infrastructure

Internet connectivity is fundamental to TM's business

Backhaul


Access

Infra

TM Overview

CAIDA's IPv4 AS Core AS-level INTERNET GRAPH

Skitter January 2000




A rapidly expanding infrastructure

copyright © 2000 UC Regents. all rights reserved.


CAIDA's IPv4 AS Core AS-level INTERNET GRAPH

Archipelago January 2015




copyright © 2015 UC Regents. All rights reserved.

TM Overview

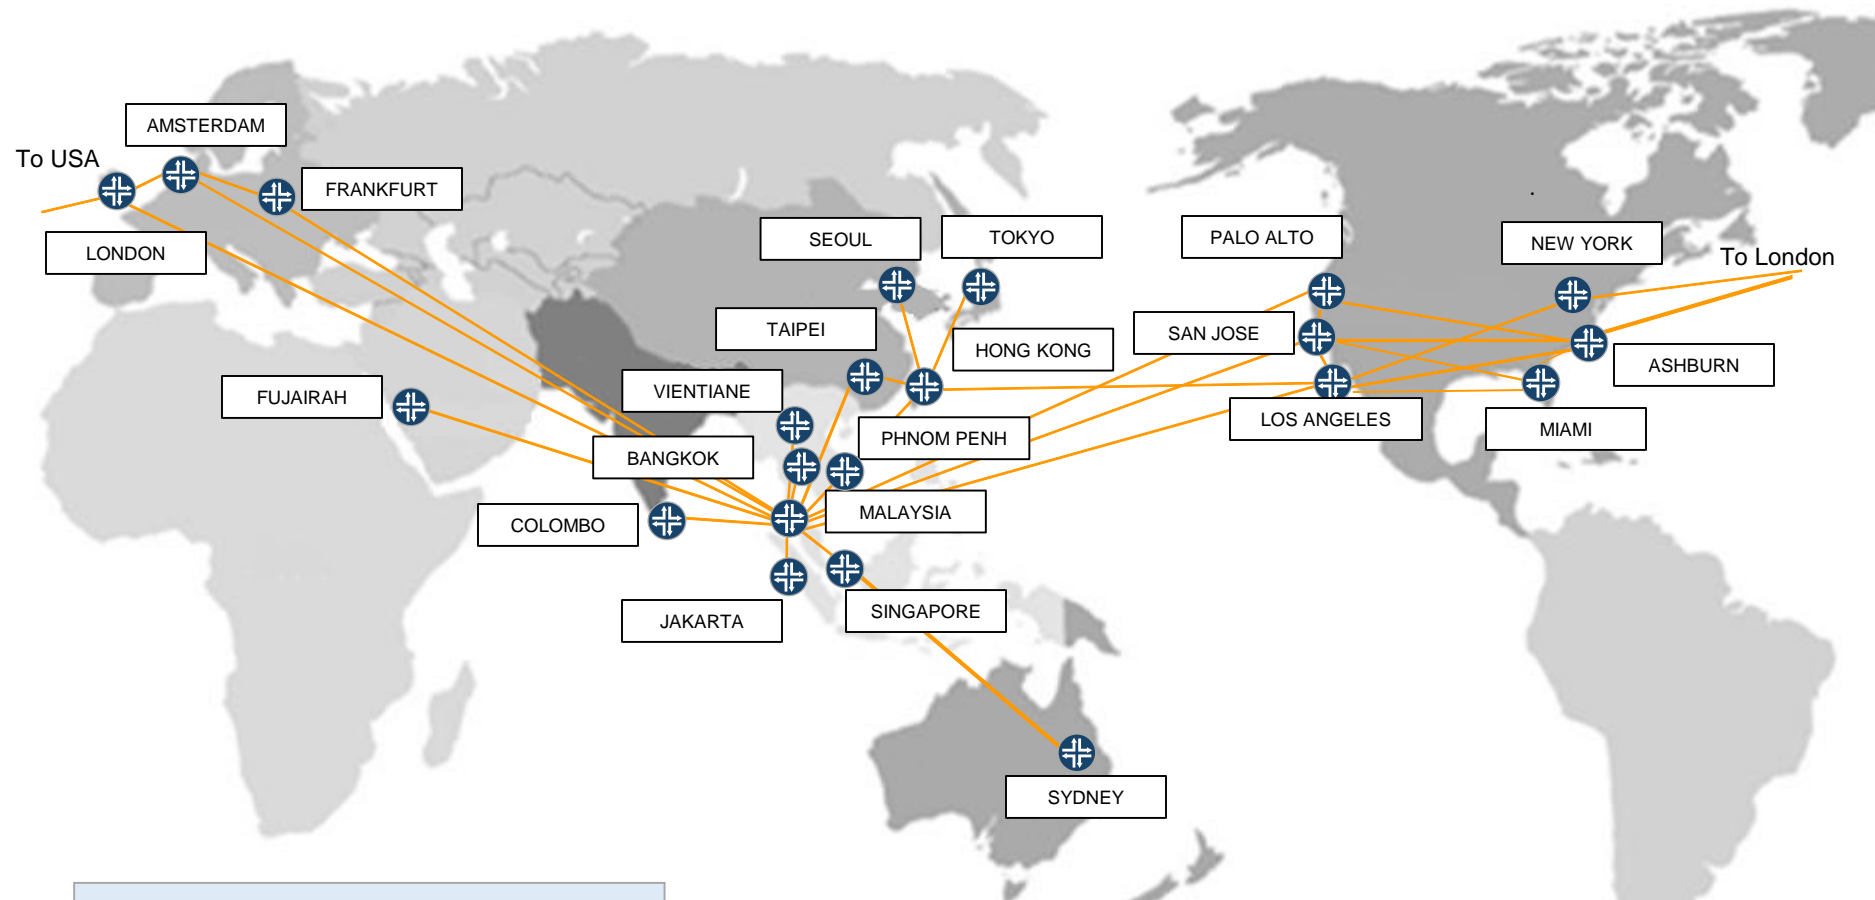


- Basically BGP peerings between ISPs
- Everyone peers!
- Market price of transit service is expensive
- Good to peer (market price: a third of transit cost)
- When peering locally
= better latency, less cost for international nodes & transmission links

	Peers	Transit Provider	Transit Subscriber
TIER 1 ISPs	✓	✓	✗
TIER 2 ISPs	✓	✓	✓
TIER 3 ISPs	✓	✗	✓



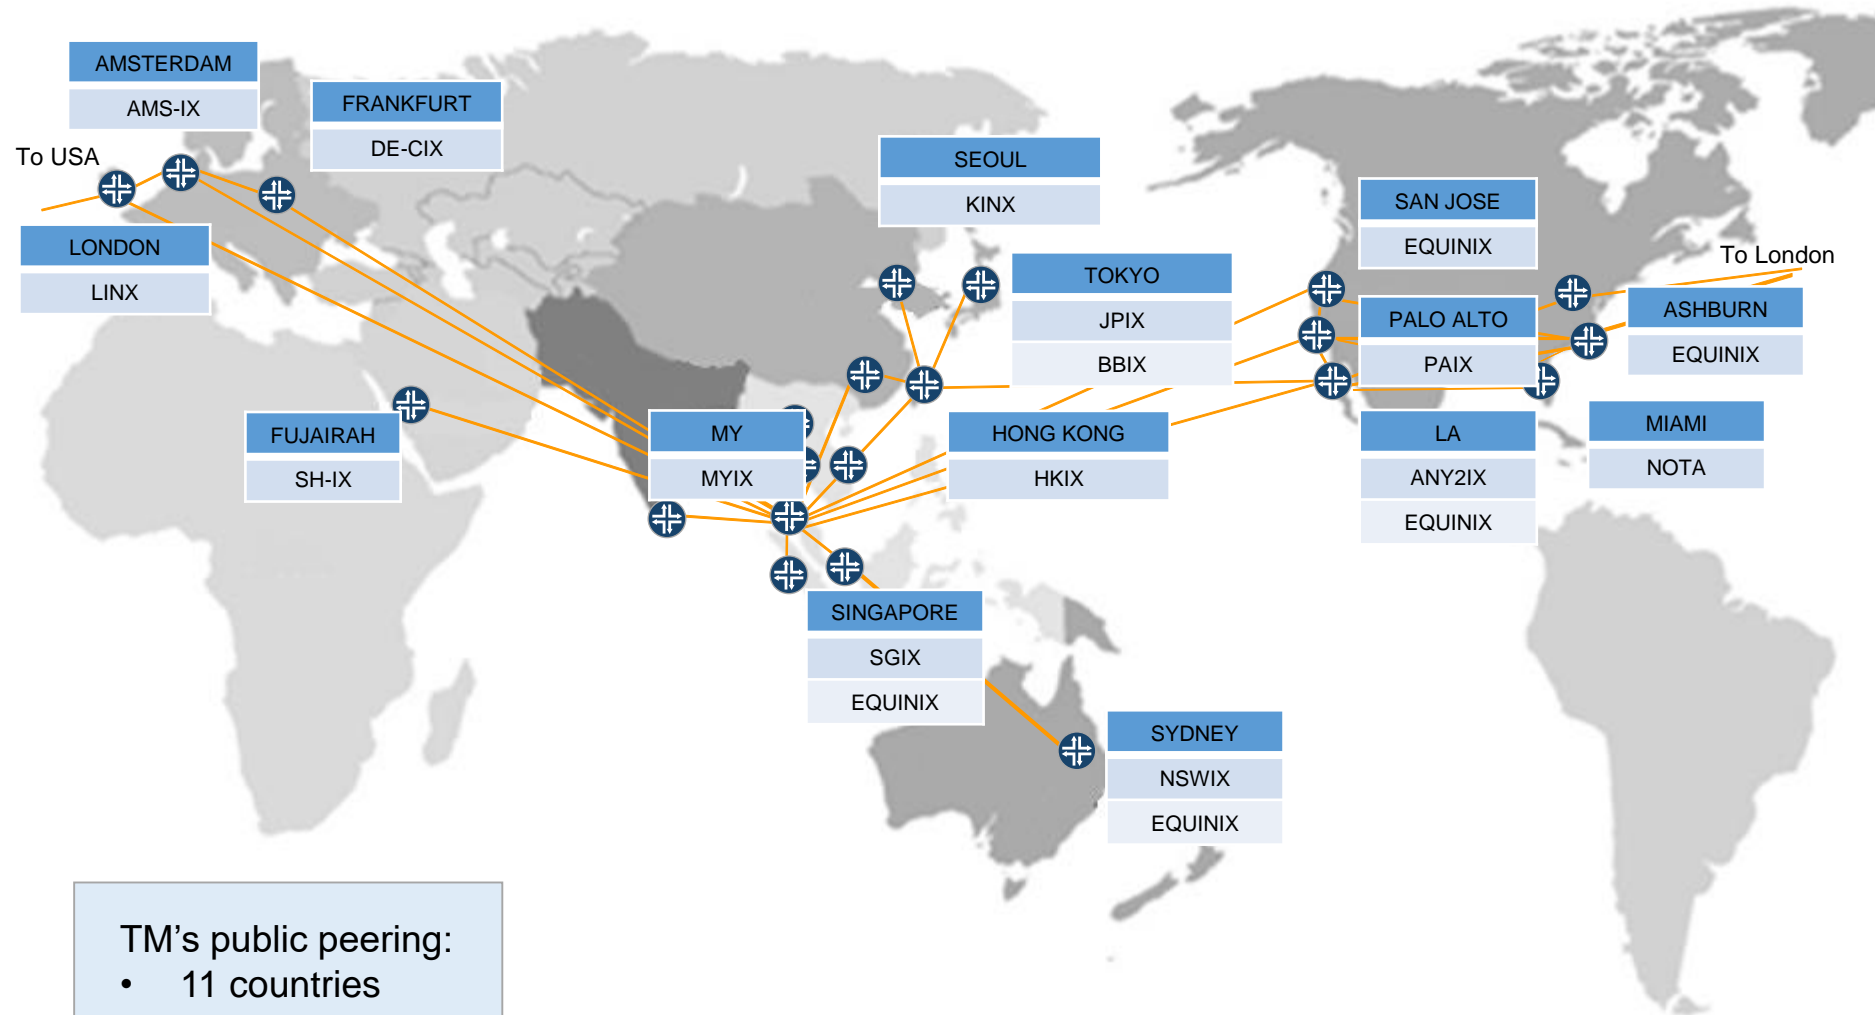
TM Overview



~2 Tb/s internet traffic capacity

- Upstream
- Peering
- Cache

TM Overview



TM's public peering:

- 11 countries
- 19 locations

Content

- Overview of TM's network
- Peering Strategies & Tools
 - Top down
 - Bottom up
- Best Practices
 - Inbound filters
 - Outbound filters
 - Aggregation
- MyIX Peering Snapshot

Peering Strategies

- 1 **Name** Top Down
- 2 **Goal** To get as much direct connectivity to internet prefixes as possible with the minimal number of peerings
- 3 **Why?**
 - Provides low-latency connectivity to as many destination prefixes as possible
 - Satisfies the demand of the majority of our customers
- 4 **How?**
 - Step 1: Identify peer
 - Step 2: Identify peering location
 - Step 3: Peer

Peering Tools

[AS Ranking](#)
[Org Ranking](#)
[Information for a single AS](#)
[Information for a single Org](#)
[Background](#)
[Data Sources](#)
[Help](#)
[AS Ranking Help](#)

The top ASes ranked by customer cone size are displayed below.

For information about a specific AS, enter its AS name, its AS number, or the name of the Org of which the AS is a member.

Dataset: 2016-06-01 IPv4

Table shows of 54772 ASes, sorted by

AS rank	AS number	AS name	Org name	AS Type(s)	customer cone						AS transit degree
					Number of			Percentages of all			
					ASes	IPv4 Prefixes	IPv4 Addresses	ASes	IPv4 Prefixes	IPv4 Addresses	
1	3356	LEVEL3	Level 3 Communications, Inc.	Tr/Ac	29,494	224,970	783,401,728	33%	34%	36%	4138
2	174	COGENT-174	Cogent Communications	Tr/Ac	23,299	172,963	616,423,936	42%	26%	28%	4567
3	1299	TELIANET	TeliaSonera AB	Tr/Ac	21,954	191,391	667,346,176	40%	29%	31%	1272
4	2914	NTT-COMMUN...	NTT America, Inc.	Tr/Ac	18,991	174,304	642,432,768	34%	26%	29%	1352
5	3257	GTT-BACKBONE	Tinet Spa	Tr/Ac	18,140	161,377	565,089,024	33%	24%	26%	1282
6	6762	SEABONE-NET	TELECOM ITALIA SPARKLE S.p.A.	Tr/Ac	14,394	123,771	329,530,624	26%	18%	15%	534
7	6453	AS6453	TATA COMMUNICATIONS (AMERICA) INC	Tr/Ac	12,300	135,127	533,133,824	22%	20%	24%	685
8	6939	HURRICANE	Hurricane Electric, Inc.	Tr/Ac	8,088	79,800	278,942,720	14%	12%	12%	4809
9	2828	XO-AS15	XO Communications	Tr/Ac	6,251	60,271	250,568,448	11%	9.2%	11%	1089
10	1273	CW	Cable and Wireless Worldwide plc	Tr/Ac	5,878	42,258	173,223,936	10%	6.4%	8.1%	296

Step 1

CAIDA
<http://as-rank.caida.org/>
 can suggest ISPs with a lot of prefixes (ie customers & content)

Peering Tools

The relationship table below displays the neighbors of **AS 4788**, and each neighbor's inferred relationship type with AS 4788.

Table shows of 178 neighbor ASes, sorted by relationship type and AS rank, with details.

neighbor					type
AS rank	AS	AS name	AS type(s)	Org name	
1	3356	LEVEL3	Tr/Ac	Level 3 Communications, Inc.	↑ provider
2	174	COGENT-174	Tr/Ac	Cogent Communications	↑ provider
3	1299	TELIANET	Tr/Ac	TeliaSonera AB	↑ provider
4	2914	NTT-COMMUN...	Tr/Ac	NTT America, Inc.	↑ provider
5	3257	GTT-BACKBONE	Tr/Ac	Tinet Spa	↑ provider
7	6453	AS6453	Tr/Ac	TATA COMMUNICATIONS (AMERICA) INC	↑ provider
8	6939	HURRICANE	Tr/Ac	Hurricane Electric, Inc.	↑ provider
10	1273	CW	Tr/Ac	Cable and Wireless Worldwide plc	↑ provider
14	4436	AS-GTT-4436	Tr/Ac	nLayer Communications, Inc.	↑ provider
17	3491	BTN-ASN	Tr/Ac	Beyond The Network America, Inc.	↑ provider
20	3320	DTAG	Tr/Ac	Deutsche Telekom AG	↑ provider
38	43531	IXREACH	Tr/Ac	IX Reach Ltd	↑ provider
16	9002	RETN-AS	Tr/Ac	RETN Limited	↔ peer
18	20485	TRANSTELECOM	Tr/Ac	Closed Joint Stock Company TransTeleCom	↔ peer
22	12389	ROSTELECOM-AS	Tr/Ac	PJSC Rostelecom	↔ peer
24	7473	SINGTEL-AS-AP	Tr/Ac	Singapore Telecommunications Ltd	↔ peer
26	3216	SOVAM-AS	Tr/Ac	OJSC "Vimpelcom"	↔ peer
30	5580	HIBERNIA	Tr/Ac	Hibernia Networks (Netherlands) BV	↔ peer
31	8359	MTS	Tr/Ac	MTS PJSC	↔ peer
32	20764	RASCOM-AS	Tr/Ac	CJSC RASCOM	↔ peer


Step 1

CAIDA
<http://as-rank.caida.org/>


or with diverse
connectivity to other
networks

Peering Tools

AS4788 IPv4 Peers



AS4788 IPv6 Peers



ASN	Name
AS6453	TATA COMMUNICATIONS (AMERICA) INC
AS3257	Tinet Spa
AS1299	Telia Company AB
AS6939	Hurricane Electric, Inc.
AS2914	NTT America, Inc.
AS3320	Deutsche Telekom AG
AS1273	Cable and Wireless Worldwide plc
AS9381	Wharf T&T Ltd.
AS23678	MyKRIS Asia Sdn Bhd
AS174	Cogent Communications

ASN	Name
AS1299	Telia Company AB
AS6939	Hurricane Electric, Inc.
AS174	Cogent Communications
AS3257	Tinet Spa
AS2914	NTT America, Inc.
AS6453	TATA COMMUNICATIONS (AMERICA) INC
AS3356	Level 3 Communications, Inc.
AS24514	Malaysian Research & Education Network
AS4635	Hong Kong Internet Exchange--Route Server 1

Step 1

Also Hurricane Electric's
<http://bgp.he.net/>

Shows an ISP's
connectivity to other
networks

Peering Tools

Organization	Telekom Malaysia Berhad (TM)
Also Known As	TM Global IP Network
Company Website	http://www.tm.com.my
Primary ASN	4788
IRR Record	AS-4788
Route Server URL	
Looking Glass URL	
Network Type	Cable/DSL/ISP
IPv4 Prefixes	10000
IPv6 Prefixes	500
Traffic Levels	500-1000 Gbps
Traffic Ratios	Mostly Inbound
Geographic Scope	Global
Protocols Supported	<input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> IPv6
Last Updated	2016-06-28T07:33:07Z
Notes	TM Global Network is MPLS enabled and IPV6 ready
Peering Policy Information	
Peering Policy	http://www.tm.com.my
General Policy	Selective
Multiple Locations	Required - International
Ratio Requirement	No
Contract Requirement	Private Only

Public Peering Exchange Points			Filter
Exchange ▼	IPv4	Speed	
ASN	IPv6	RS Peer	
AMS-IX	80.249.209.91	10G	
4788	2001:7f8:1::a500:4788:1		
BBIX Tokyo	218.100.6.118	10G	
4788	2001:de8:c::4788:1		
CoreSite - Any2 California	206.72.211.84	10G	
4788	2001:504:13::211:84		
DE-CIX Frankfurt Main	80.81.194.24	10G	
4788	2001:7f8::12b4:0:1		
Equinix Ashburn	206.126.236.176	10G	
4788	2001:504:0:2::4788:1		
Equinix Los Angeles	206.223.123.94	10G	
4788	2001:504:0:3::4788:1		
Equinix Palo Alto	198.32.176.26	10G	
4788	2001:504:d::1a		
Equinix San Jose	206.223.116.120	10G	
4788	2001:504:0:1::4788:1		
Equinix Singapore	27.111.228.15	30G	
4788	2001:de8:4::4788:1		
Equinix Sydney	202.167.228.172	1G	
4788	2001:de8:6::4788:1		
HKIX	123.255.91.222	10G	
4788	2001:7fa:0:1::ca28:a1de		
IX Australia NSW	218.100.52.173	1G	
4788	2001:7fa:11:4:0:12b4:0:1		
JPIX	210.171.224.54	2G	
4788	2001:de8:8::4788:2		
KINX	192.145.251.44	2G	
4788	2001:7fa:8::16		
LINX LON1	195.66.224.47	10G	
Private Peering Facilities			Filter
Facility ▼	Country		

Step 2 & 3

We can then find out where to peer using Peering DB

<https://www.peeringdb.com>

Shows the ISP's peering presence


Peering Strategies

- 1 **Name** Bottom Up
- 2 **Goal** To get direct connectivity to specifically identified internet prefixes
- 3 **Why?**
 - A lot of customer traffic goes to or comes from the prefixes
 - Requested by a high-priority customer
 - A tender requirement
- 4 **How?**
 - Step 1: Identify desired prefix
 - Step 2: Identify peer
 - Step 3: Identify peering location
 - Step 4: Peer

Peering Tools


























































Step 1 & 2

- Example of a traffic analysis tool
- Shows traffic volume to/from an Autonomous System
- Recommends new peerings

Click ASN to drill-down into traffic through that ASN 

DETAILS Period: Today Hide Managed Objects: ☐  Update

Current | [Average](#) | [Max](#) | [PCT95](#)

Rank	ASN Name	Relationship	Into Network		Out of Network		Total ▼
1	32934 FACEBOOK	Peer		Gbps		Gbps	 Gbps
2	15169 GOOGLE	Peer		Gbps		Gbps	 Gbps
3	3491 BTN	Peer		Gbps		Gbps	 Gbps
4	1273 CW	Peer		Gbps		Gbps	 Gbps
5	1299 TELIANET	Peer		Gbps		Gbps	 Gbps
6	3257 TINET-BACKBONE	Peer		Gbps		Gbps	 Gbps
7	2914 NTT-COMMUNICATIONS-2914	Peer		Gbps		Gbps	 Gbps
8	6453 GLOBEINTERNET	Peer		Gbps		Gbps	 Gbps
9	7473 SINGTEL-AS-AP	Peer		Gbps		Gbps	 Gbps
10	4134 CHINANET-BACKBONE	Peer		Gbps		Gbps	 Gbps
11	4837 CHINA169-BACKBONE	Peer		Gbps		Gbps	 Gbps
12	174 COGENT	Peer		Gbps		Gbps	 Gbps
13	714 APPLE	Peer		Gbps		Mbps	 Gbps
14	6185 APPLE-AUSTIN			Gbps		Mbps	 Gbps
15	3356 LEVEL3	Peer		Gbps		Gbps	 Gbps
16	41690 DAILYMOTION	Peer		Gbps		Mbps	 Gbps
17	8075 MICROSOFT-CORP-MSN-AS-BLOCK	Peer		Gbps		Gbps	 Gbps
18	1239 AS1239	Peer		Gbps		Gbps	 Gbps
19	20940 AKAMAI-ASN1	Peer		Gbps		Gbps	 Gbps
20	38895 AMAZON-AS-AP	Peer		Gbps		Gbps	Gbps
21	64049 IANA-RSVD			Mbps		Gbps	Gbps
22	55836 RELIANCE-INFOTEL-IN			Mbps		Gbps	Gbps

Peering Tools

Team Cymru IP to ASN Lookup v1.0



[CYMRU] [ASN LOOKUP] [HTTP(S) ASN LOOKUP]

Family: ☒ IPv4 ☐ IPv6 Methods: ☒ whois ☐ peer-whois
 Flags: ☐ prefix ☐ cc ☐ registry ☐ allocated ☐ nottruncate ☐ verbose

58.27.1.1

Insert your IP or ASN in the textbox above.

IPv4 [OPTIONAL COMMENT]

Eg. '4.2.2.2 2004-12-10 11:33:21 GMT'

AS#

Eg. 'AS23028'

Step 1 & 2

Team Cymru's

<https://asn.cymru.com/cgi-bin/whois.cgi>

Gives the owner or AS number of the owner of the identified/requested prefix

Executing commands. Please be patient!

v4.whois.cymru.com

The server returned 4 line(s).

[Querying v4.whois.cymru.com]

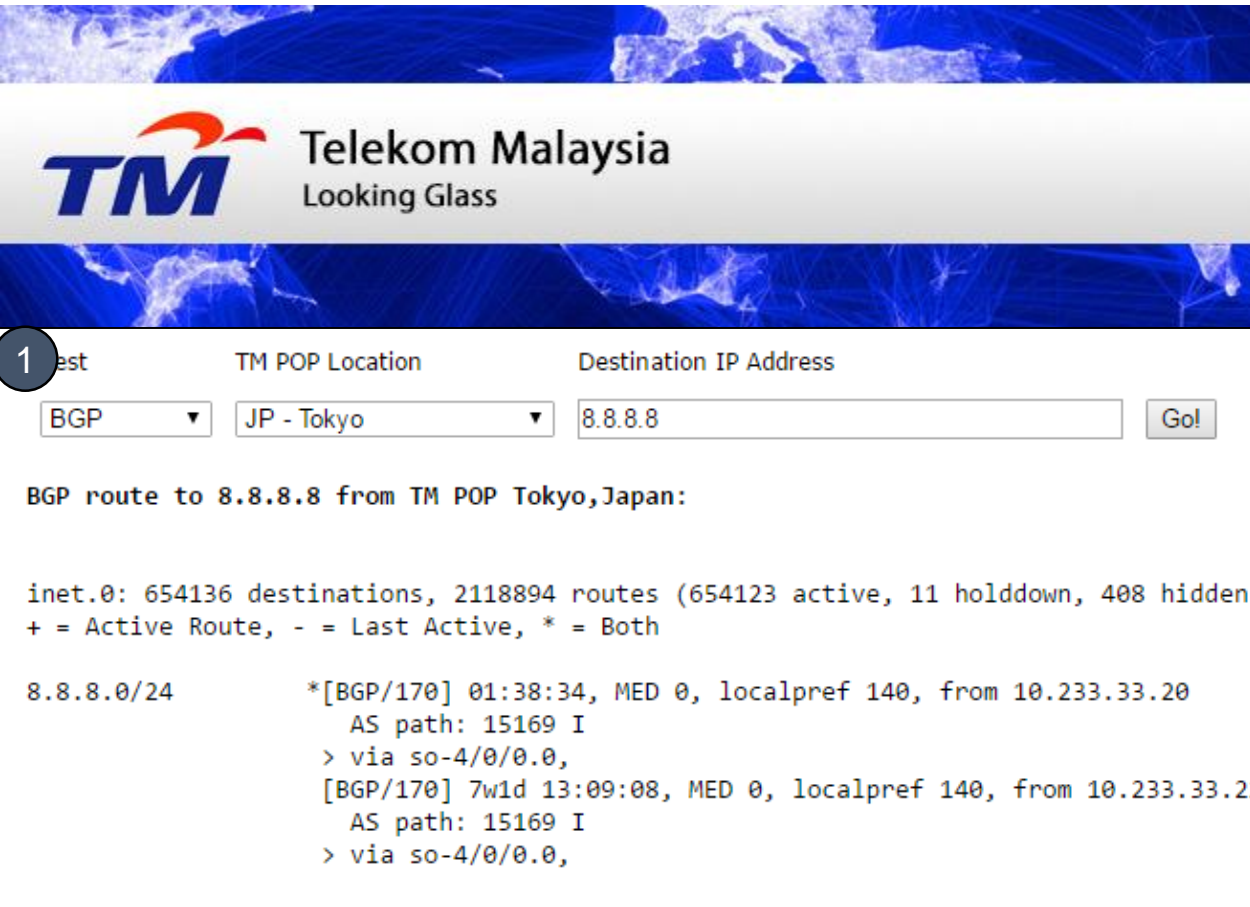
[v4.whois.cymru.com]

AS	IP	AS Name
4788	58.27.1.1	TMNET-AS-AP TM Net, Internet Service Provider, MY

Peering Tools

Step 2

- If direct peering not possible, peer with a partner with good connectivity to the desired prefix and negotiate passage to the prefix
- Looking Glass:** a free, publicly accessible tool to understand how a peer routes traffic



Telekom Malaysia
Looking Glass

1 Test TM POP Location Destination IP Address
BGP JP - Tokyo 8.8.8.8 Go!

BGP route to 8.8.8.8 from TM POP Tokyo, Japan:

```
inet.0: 654136 destinations, 2118894 routes (654123 active, 11 holddown, 408 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24      *[BGP/170] 01:38:34, MED 0, localpref 140, from 10.233.33.20
                AS path: 15169 I
                > via so-4/0/0.0,
                [BGP/170] 7w1d 13:09:08, MED 0, localpref 140, from 10.233.33.22
                AS path: 15169 I
                > via so-4/0/0.0,
```

2 Test TM POP Location Destination IP Address
IPv4 Ping DE - Frankfurt 8.8.8.8 Go!

Executing 5 time ping to 8.8.8.8 from TM POP Frankfurt, Germany:

```
PING 8.8.8.8 (8.8.8.8): 56 data bytes
64 bytes from 8.8.8.8: icmp_seq=0 ttl=61 time=203.224 ms
64 bytes from 8.8.8.8: icmp_seq=1 ttl=61 time=181.604 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=61 time=181.550 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=61 time=181.550 ms
64 bytes from 8.8.8.8: icmp_seq=4 ttl=61 time=187.122 ms

--- 8.8.8.8 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 181.550/187.010/203.224/8.388 ms
```

3 Test TM POP Location Destination IP Address
IPv4 Trace HK - Hong Kong 8.8.8.8 Go!

Executing traceroute to 8.8.8.8 from TM POP Hong Kong 1, Hong Kong SAR:

```
1  10.55.192.101 (10.55.192.101)  45.616 ms  45.670 ms  46.069 ms
2  72.14.197.66 (72.14.197.66)   42.757 ms  42.767 ms  72.14.213.128 (72.14.213.128)
3  216.239.59.245 (216.239.59.245) 46.211 ms  216.239.59.229 (216.239.59.229) 46.
4  209.85.249.105 (209.85.249.105) 55.217 ms  72.14.234.19 (72.14.234.19)  51.788
5  8.8.8.8 (8.8.8.8)             46.084 ms  46.086 ms  46.064 ms
```

Peering Tools

Step 2 Looking Glass: a common free tool offered by major ISPs




Router:
 Kuala Lumpur - MY ▼

Query:
 BGP ▼

IP Address

☐ Your current IP Address: 202.188.60.37
☒ Specify an IP Address (IPv4 or IPv6) 58.27.1.1

Only IP addresses are allowed parameters for BGP Queries. FQDN can not be used.


Query Results:
Router: Kuala Lumpur - MY
Command: show route protocol bgp 58.27.1.1 terse

```

inet.0: 621058 destinations, 7460580 routes (620476 active, 32 holddown, 482077 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination      P Prf  Metric 1  Metric 2  Next hop          AS path
* ? 58.27.0.0/18      B 170    120      >203.115.193.178  4788 I
   unverified
  
```

{master}


SingTel Internet Exchange(STIX) | Looking Glass

Query Type :
 ☒ Ping To
 ☐ Traceroute To
 ☐ BGP

Router : Chaiwan, HK (IPv4, IPv6) ▼

Ping To : 58.27.1.1

Query Results [Print](#)

202.188.60.37

Taipei, Taiwan Router ping to 58.27.1.1:

Query Request Start Time: 2016-08-08 15:57:02
 Query Request End Time: 2016-08-08 15:56:27

ping 58.27.1.1
 Type escape sequence to abort.
 Sending 5, 100-byte ICMP Echos to 58.27.1.1, timeout is 2 seconds:

Peering Tools

Organization	Telekom Malaysia Berhad (TM)
Also Known As	TM Global IP Network
Company Website	http://www.tm.com.my
Primary ASN	4788
IRR Record	AS-4788
Route Server URL	
Looking Glass URL	
Network Type	Cable/DSL/ISP
IPv4 Prefixes	10000
IPv6 Prefixes	500
Traffic Levels	500-1000 Gbps
Traffic Ratios	Mostly Inbound
Geographic Scope	Global
Protocols Supported	<input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> IPv6
Last Updated	2016-06-28T07:33:07Z
Notes	TM Global Network is MPLS enabled and IPV6 ready
Peering Policy Information	
Peering Policy	http://www.tm.com.my
General Policy	Selective
Multiple Locations	Required - International
Ratio Requirement	No
Contract Requirement	Private Only

Public Peering Exchange Points			Filter
Exchange ▼	IPv4	Speed	
ASN	IPv6	RS Peer	
AMS-IX	80.249.209.91	10G	
4788	2001:7f8:1::a500:4788:1	○	
BBIX Tokyo	218.100.6.118	10G	
4788	2001:de8:c::4788:1	○	
CoreSite - Any2 California	206.72.211.84	10G	
4788	2001:504:13::211:84	○	
DE-CIX Frankfurt Main	80.81.194.24	10G	
4788	2001:7f8::12b4:0:1	○	
Equinix Ashburn	206.126.236.176	10G	
4788	2001:504:0:2::4788:1	○	
Equinix Los Angeles	206.223.123.94	10G	
4788	2001:504:0:3::4788:1	○	
Equinix Palo Alto	198.32.176.26	10G	
4788	2001:504:d::1a	○	
Equinix San Jose	206.223.116.120	10G	
4788	2001:504:0:1::4788:1	○	
Equinix Singapore	27.111.228.15	30G	
4788	2001:de8:4::4788:1	○	
Equinix Sydney	202.167.228.172	1G	
4788	2001:de8:6::4788:1	○	
HKIX	123.255.91.222	10G	
4788	2001:7fa:0:1::ca28:a1de	○	
IX Australia NSW	218.100.52.173	1G	
4788	2001:7fa:11:4:0:12b4:0:1	○	
JPIX	210.171.224.54	2G	
4788	2001:de8:8::4788:2	○	
KINX	192.145.251.44	2G	
4788	2001:7fa:8::16	○	
LINX LON1	195.66.224.47	10G	
4788	2001:7fa:1::16	○	
Private Peering Facilities			Filter
Facility ▼	Country		

Step 3 & 4

We can then find out where to peer using Peering DB

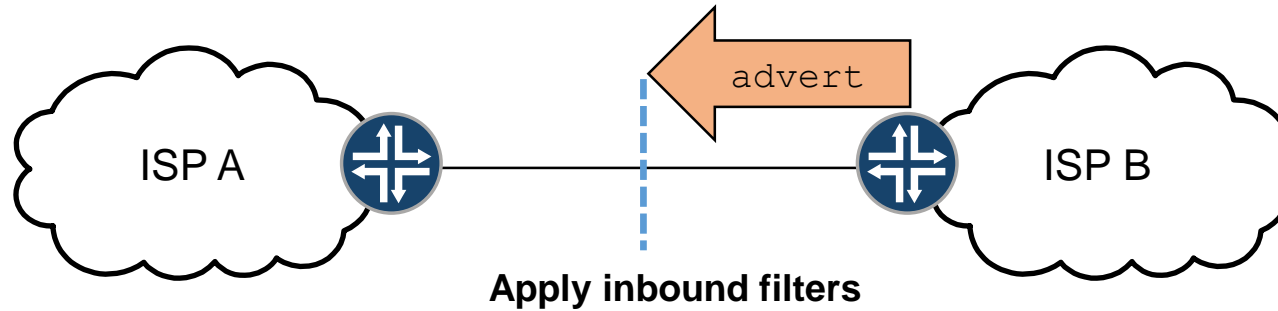
<https://www.peeringdb.com>

Shows the ISP's peering presence

Content

- Overview of TM's network
- Peering Strategies & Tools
 - Top down
 - Bottom up
- Best Practices
 - Inbound filters
 - Outbound filters
 - Aggregation
- MyIX Peering Snapshot

Peering Best Practices



1

Aim:

- Prevents the receipt of prefixes that can interfere with our network or cause customer traffic to be wrongly routed

2

General filter:

- Accept only prefixes agreed in the peering agreement
- Accept only prefixes with the ISP's ASN as the most recent ASN in AS path
- Set a prefix-limit
- Reject default routes
- Reject own prefixes
- Reject multicast prefixes

3

V4 filter:


- Reject RFC1918 prefixes (private IPs)
- Reject RFC5735 prefixes (loopback, test, reserved etc prefixes)
- Reject longer than /24 prefixes (accept only summarised prefixes)
- Reject other bogon prefixes: unallocated prefixes (but RFC6441 recommends the filter to be removed for IPv4)
→ <http://www.team-cymru.org/bogon-reference-bgp.html>

4

V6 filter:

- Reject RFC5156 prefixes (link-scoped, ipv4-compatible, ipv4-mapped etc prefixes)
- Reject prefixes longer than /64 (accept only summarised prefixes)
- Reject other bogons prefixes: unallocated prefixes
→ <http://www.team-cymru.org/bogon-reference-bgp.html>
- RIPE IPV6 filter guide:
→ <http://www.space.net/~gert/RIPE/ipv6-filters.html>

Peering Best Practices



1

Aim:

- Prevents the advertisement of prefixes that can interfere with other networks or cause their customer traffic to be wrongly routed

2

General filter:

- Send only prefixes agreed in the peering agreement
- Reject default routes
- Reject own prefixes
- Reject multicast

3

V4 filter:


- Reject RFC1918 prefixes (private IPs)
- Reject RFC5735 prefixes (loopback, test, reserved etc prefixes)
- Reject longer than /24 prefixes (accept only summarised prefixes)
- Reject other bogon prefixes: unallocated prefixes (but RFC6441 recommends the filter to be removed for IPv4)
→ <http://www.team-cymru.org/bogon-reference-bgp.html>

4

V6 filter:

- Reject RFC5156 prefixes (link-scoped, ipv4-compatible, ipv4-mapped etc prefixes)
- Reject prefixes longer than /64 (accept only summarised prefixes)
- Reject other bogons prefixes: unallocated prefixes
→ <http://www.team-cymru.org/bogon-reference-bgp.html>
- RIPE IPV6 filter guide:
→ <http://www.space.net/~gert/RIPE/ipv6-filters.html>

Peering Best Practices



1

- Aggregate when advertising.
- Perfect case: announce the address block assigned by RIRs
 - APNIC : IPv4:/24 IPv6:/48
 - LACNIC : IPv4:/22 IPv6:/32
 - AFRINIC: IPv4:/22 IPv6:/32
 - RIPE : IPv4:/24 IPv6:/32
 - ARIN : IPv4:/24 IPv6:/32

2

- Aggregates should be generated internally
 - for example RRs, not at the borders
 - easier to manage, few RRs, many border routers

3


- Reduces the number of prefixes in the BGP routing table
 - reduces memory requirement of routers
 - more stable routing table
 - reduces recovery time
 - speeds up convergence – bigger tables take longer for CPUs to process
 - reduces size of BGP updates. Bigger updates take longer to process
- Examples later

Aggregation Example

1

Steady-state:

2 × /24 sent to the internet instead of one /23
→memory consumption



2

Link to customer 1 fails:

- 1 × /24 withdrawn by ISP A
- Update propagated to the rest of the internet
- All internet routers need to process the update
- New converged BGP table calculated by all
→CPU consumption

3

Link to customer 1 restored:


- 1 × /24 re-advertised by ISP A
- Update propagated to the rest of the internet
- All internet routers need to process the update
- New converged BGP table calculated by all
→CPU utilisation
- Route damping (suppression of flapping routes) implemented by some ISPs
→traffic from some portion of the Internet does not immediately reach Customer 1
→extended downtime

Aggregation Example

1

Steady-state:

1 × /23 sent to the internet instead of 2 × /24
→ reduced memory consumption



2

Link to customer 1 fails:

- 1 × /23 remains advertised by ISP A
- Not updated to the rest of the internet
→ Internet BGP table remains as steady-state
→ reduced CPU consumption
- Traffic sent to ISP A but dropped

3

Link to customer 1 restored:

- 1 × /24 re-advertised to ISP A
- Not updated to the rest of the internet
→ Internet BGP table remains as steady-state
→ reduced CPU consumption
- All traffic from the Internet immediately reaches Customer 1
→ less downtime

Aggregation

CIDR <http://www.cidr-report.org/as2.0/> provides a good indication of how much more aggregation can be done

IPv4

--- 04Aug16 ---					
ASnum	NetsNow	NetsAggr	NetGain	% Gain	Description
Table	621994	345834	276160	44.4%	All ASes
AS39891	3329	16	3313	99.5%	ALJAWWALSTC-AS , SA
AS7545	3484	429	3055	87.7%	TPG-INTERNET-AP TPG Telecom Limited, AU
AS4538	5552	2647	2905	52.3%	ERX-CERNET-BKB China Education and Research Network Center, CN
AS17974	2929	77	2852	97.4%	TELKOMNET-AS2-AP PT Telekomunikasi Indonesia, ID
AS6389	2092	44	2048	97.9%	BELLSOUTH-NET-BLK - BellSouth.net Inc., US

IPv6

--- 04Aug16 ---					
ASnum	NetsNow	NetsAggr	NetGain	% Gain	Description
Table	31468	22414	9054	28.8%	All ASes
AS3651	573	65	508	88.7%	SPRINT-BB6 - Sprint, US
AS22773	460	84	376	81.7%	ASN-CXA-ALL-CCI-22773-RDC - Cox Communications Inc., US
AS13206	258	1	257	99.6%	REALMOVE-AS-AP Realmove Company Limited, TH
AS28573	431	183	248	57.5%	CLARO S.A., BR
AS20277	255	8	247	96.9%	OSP-EU-WEST-001 , GB

Aggregation

CIDR <http://www.cidr-report.org/as2.0/> same tool can be used to help plan route aggregation

Selected AS Report

Enter an AS here to generate an aggregation report for the AS.


Enter AS (e.g. "AS1221")

Prefix	AS Path	Aggregation Suggestion
1.9.0.0/16	4777 2516 4788	
1.9.34.0/24	4608 24130 7545 6939 4788	
1.9.37.0/24	4608 24130 7545 6939 4788	
1.9.40.0/23	4608 24130 7545 6939 4788	+ Announce - aggregate of 1.9.40.0/24 (4608 24130 7545 6939 4788) and 1.9.41.0/24 (4608 24130 7545 6939 4788)
1.9.40.0/24	4608 24130 7545 6939 4788	- Withdrawn - aggregated with 1.9.41.0/24 (4608 24130 7545 6939 4788)
1.9.41.0/24	4608 24130 7545 6939 4788	- Withdrawn - aggregated with 1.9.40.0/24 (4608 24130 7545 6939 4788)
1.9.65.0/24	4608 24130 7545 6939 4788	
1.9.79.0/24	4608 24130 7545 6939 4788	
1.32.0.0/17	4777 2516 4788	
1.32.0.0/17	4608 1221 4637 4788	+ Announce - aggregate of 1.32.0.0/18 (4608 1221 4637 4788) and 1.32.64.0/18 (4608 1221 4637 4788)
1.32.0.0/19	4608 1221 4637 4788	- Withdrawn - aggregated with 1.32.32.0/19 (4608 1221 4637 4788)
1.32.32.0/19	4608 1221 4637 4788	- Withdrawn - aggregated with 1.32.0.0/19 (4608 1221 4637 4788)
1.32.64.0/18	4608 1221 4637 4788	- Withdrawn - aggregated with 1.32.0.0/18 (4608 1221 4637 4788)

Content

- Overview of TM's network
- Peering Strategies & Tools
 - Top down
 - Bottom up
- Best Practices
 - Inbound filters
 - Outbound filters
 - Aggregation
- MyIX Peering Snapshot

MyIX Peering Snapshot



Filtering

Prefix filtering v4

of peers = 66
 # of received prefixes = 2748
 # of accepted prefixes = 2748
 # of filtered prefixes = 0
 → good

Prefix filtering v6

of peers = 13
 # of received prefixes = 47
 # of accepted prefixes = 47
 # of filtered prefixes = 0
 → good

Aggregation

Prefix aggregation v4

of prefixes = 2748
 # if perfectly aggregated = 2032
 % potential improvement = 26%
 → significant

Prefix aggregation v6

of prefixes = 47
 # if perfectly aggregated = 44
 % potential improvement = 6%
 → OK

Conclusion

- Overview of TM's network
 - we are available at many peering locations
- Peering Strategy & Tools
 - Top down approach → to get as much routes as possible
 - Bottom up approach → to get targeted routes
- Best Practices
 - Inbound filters → a must
 - Outbound filters → a courtesy
 - Aggregation → is good
- MyIX Peering Snapshot:
 - OK. v4 aggregation can be better



Thank You
aeffendi@tm.com.my