

Mikrotik in real life, full scale and low budget ISP

wima.sy@gmail.com

additional presentation

How do the youngest country in the world ISPs run their bussiness more efficient, and more reliable with Mikrotik

about me

- Working in ISP industries since 1994
- Currently working as consultant engineer for asia pacific oceania countries company & organization

Definition


- Real life
 - Not in simulation, or lab scale
 - In bussiness, operational, still operational
- Full Scale
 - Have ASN, buy transit, peering
 - Connect to IX
- Low Budget
 - Is low budget

Disclaimer

- This presentation will not talk in depth about BGP, OSPF & Traffic Engineering
- I just share simple example, and how to do it with Mikrotik
- It is real case, some IP/AS is fake, for security

Before

- Cisco 7200 VXR
 - Border Router
 - BGP Peering to Transit Provider
 - BGP Peering to Local IXP
 - Customer Access Router
- IBM e Series
 - FreeBSD / Quagga
 - BGP Peering to Local IXP
 - BGP Peering to IP Transit Customer
- Problem
 - Expensive Router
 - Difficult to Maintain



Reason to Upgrade


- Efficiency
- Performance
- Maintenance
- Cost
- Growth

After

- CCR 1036-8 – Transit
 - OSPF
 - BGP Peering to Transit Provider
- CCR 1036-8 – IXP
 - OSPF
 - BGP Peering to IXP
- CCR 1036-12 – Core
 - OSPF
 - BGP Route collector
- CCR 1009 – Access
 - Static Routing, VLAN, Trunk
 - Management Router
- RB750 – OOB
 - VPN



Physical Network Diagram



Configuration

Pre-config

- Turn off unused service features
 - Web, telnet, ftp, etc
- Winbox / SSH only available from Remote Access IP
- Change default port

Name	Port	Available From	Certificate
X <input checked="" type="checkbox"/> api	8728		
X <input type="checkbox"/> api-ssl	8729		none
X <input type="checkbox"/> ftp	21		
<input checked="" type="checkbox"/> ssh	234	172.16.16.xx	
X <input type="checkbox"/> telnet	23		
<input checked="" type="checkbox"/> winbox	8291	172.16.16.xx	
X <input type="checkbox"/> www	80		
X <input type="checkbox"/> www-ssl	443		none

8 items (1 selected)

Configuration

Turn off unused packages features

- Disable features/packages

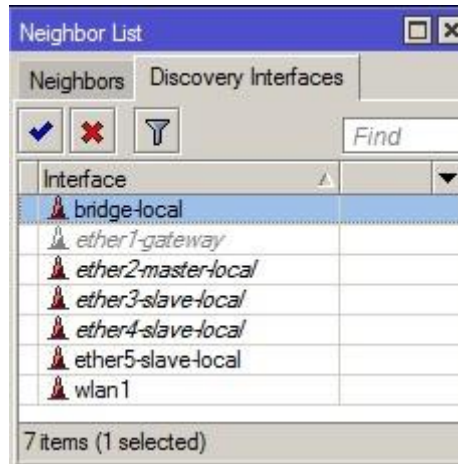
Name	Version	Build Time
routeros-tile	6.32.2	Sep/17/2015 15:20:53
advanced-tools	6.32.2	Sep/17/2015 15:20:53
dhcp	6.32.2	Sep/17/2015 15:20:53
hotspot	6.32.2	Sep/17/2015 15:20:53
ipv6	6.32.2	Sep/17/2015 15:20:53
mpls	6.32.2	Sep/17/2015 15:20:53
ppp	6.32.2	Sep/17/2015 15:20:53
routing	6.32.2	Sep/17/2015 15:20:53
security	6.32.2	Sep/17/2015 15:20:53
system	6.32.2	Sep/17/2015 15:20:53
wireless-cm2	6.32.2	Sep/17/2015 15:20:53
wireless-fp	6.32.2	Sep/17/2015 15:20:53

12 items

Configuration

Neighbour discovery

- Disable interface



- ✓ Disable MNDP on interface to IXP/Transit, some of them will handle this as a threat
- ✓ Some IX/Transit require you to turn off Proxy ARP, ICMP redirects, Directed broadcast, IEEE802 Spanning Tree, Interior routing protocol broadcast, Mac layer broadcast
- ✓ Read peering agreement


Configuration

- Disable unused physical interface
- Device name
- User / Password
 - Proper credentials
- NTP Client
 - Make sure your router time is synchronized
- Latest stable OS
- Disable LCD / Minimal information

Configuration

OSPF between devices for IGP

- for infrastructure
- loopback interface, for adjacency, not only router id



Bridge

Bridge Ports Filters NAT Hosts

+ - ✓ ✗ [Icon] [Icon] Settings Find

Name	Type	L2 MTU	Tx	Rx
R loopback0	Bridge	65535	1888 bps	0 bps

OSPF

Interfaces Instances Networks Areas Area Ranges Virtual Links Neighbors NBMA Neighbors Sham Links LSA Routes ...

1 item + - ✓ ✗ [Icon] [Icon] Find

Interface	Cost	Priority	Authentic...	Authenticatio...	Network Type	Instance	Area	Neig...	State
Ether1-TR	10	1	none	*****	broadcast	default	backbone	1	backup
Ether2-IX	10	1	none	*****	broadcast	default	backbone	1	backup
P loopback0	10	1	none	*****	broadcast	default	backbone	0	passive

OSPF Instance <default>

General Metrics MPLS Status

Name: default

Router ID: 192.168.42.240

Redistribute Default Route: never

Redistribute Connected Routes: no

Redistribute Static Routes: no

Redistribute RIP Routes: no

Redistribute BGP Routes: no

Redistribute Other OSPF Routes: no

In Filter: ospf-in

Out Filter: ospf-out

enabled default

OK Cancel Apply Disable Comment Copy Remove

OSPF Network <192.168.42.0/24>

Network: 192.168.42.0/24

Area: backbone

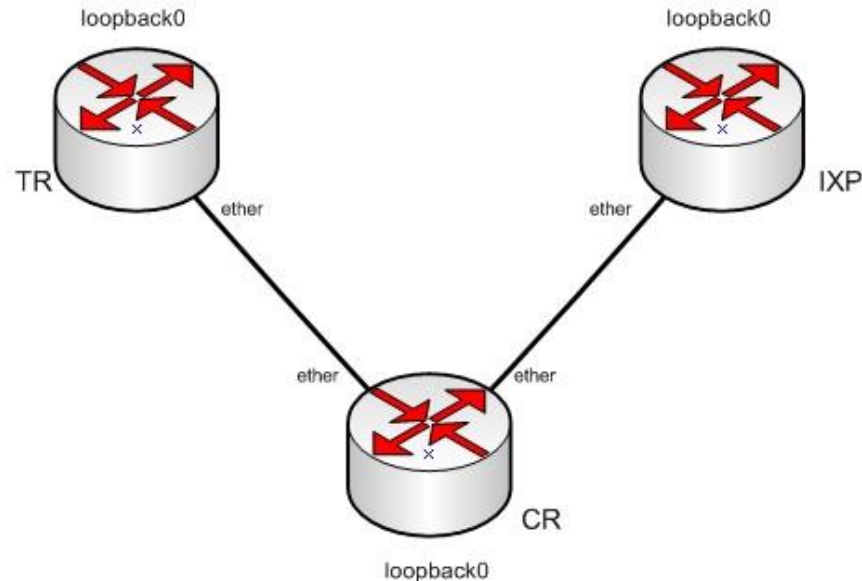
enabled

OK Cancel Apply Disable Comment Copy Remove

Configuration

iBGP between devices

- TR – CR – IXP
- Loopback interface peering
- For carry prefixes across backbone



iBGP instance

BGP Instance <default>

Name:

AS:

Router ID:

Loopback peering

BGP Peer <IX>

General Advanced Status

Name:

Instance:

Remote Address:



Remote Port:

Remote AS:

Advertise Networks

BGP

Instances VRFs Peers Networks Aggregat

Network	Synchroni...
 172.16.16.0/20	no
 192.168.0.0/16	no

Loopback interface as source

BGP Peer <TR>

General Advanced Status



Address Families: ip ipv6

Update Source:

Checking

BGP


Instances VRFs Peers Networks Aggregates VPN4 Routes Advertisements

Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID	Uptime	Prefix Co...	State
 IX	default	192.168.42.242	65530	no	no	d...	192.168.42.242	00:11:58	1	established
 TR	default	192.168.42.241	65530	no	no	d...	192.168.42.241	10:40:38		established

Configuration

eBGP between peer / other AS

- Peering
- Advertise your prefixes
- Filtering
 - In Filter -> how we send the traffic
 - Out filter -> how they will send the traffic
 - Standard regexp
 - Use template for filter
 - Organize filter using jump
- Traffic engineering, routing policy, follow BGP BCP



Peering


- Use your AS, peering IP, peer AS
- Prepare your in/out filter

The image shows a configuration window for a BGP Peer. The window has a blue header with the text "BGP Peer <IX>". Below the header are three tabs: "General", "Advanced", and "Status". The "General" tab is selected. The configuration fields are as follows:

- Name: IX
- Instance: default
- Remote Address: 192.168.42.242
- Remote Port: (empty)
- Remote AS: 65530
- TCP MD5 Key: (empty)
- Nexthop Choice: default
- Multihop
- Route Reflect
- Hold Time: 180 s
- Keepalive Time: (empty)
- TTL: default
- Max Prefix Limit: (empty)
- Max Prefix Restart Time: (empty)
- In Filter: IIX1-in
- Out Filter: IIX1-out

Advertise your prefix

- Announce your aggregate from registry



- Use blackhole type route for pull-up route
- Put on core router, not border



Announce your aggregate, for internet stability

Routing Filters

- In Filter -> how we send the traffic -> our routing table
- Out filter -> how they will send the traffic -> their route to our AS
- Template
 - In Filter
 - Discard prefix from other peering AS
 - Accept prefix from peering AS
 - Discard our own prefixes
 - Discard RFC5735 prefixes
 - Discard prefix longer than 24
 - Out Filter
 - Allow only our prefixes to be announce
 - Use jump for organize your rule
- Regexp
 - . - any single character
 - ^ - start of the as-path
 - \$ - end of the as-path
 - _ - matches comma, space, start and end of as-path

Route Filters						
#	Chain	Prefix	Prefix Length	BGP AS Path	Action	Jump Target
::: Discard prefix from other peering AS						
0	IIX1-in			_17451	discard	
1	IIX1-in			_38060	discard	
2	IIX1-in			_56258	discard	
::: Accept prefix from peering AS						
3	IIX1-in			^7597_	accept	
::: Discard our own prefixes						
4	IIX1-in				jump	Our-Discard
::: Discard RFC5735 prefixes						
5	IIX1-in				jump	RFC5735
::: Discard prefix longer than 24						
6	IIX1-in	0.0.0.0/25	25-32		discard	
7	IIX1-in				discard	
::: Allow only our prefixes to be announce						
8	IIX1-out				jump	Our-Allow
9	IIX1-out				discard	
10	Our-Allow	192.168.40.0/21	21-24		accept	
11	Our-Allow				return	
12	Our-Discard	192.168.40.0/21	21-24		discard	
13	Our-Discard				return	
14	RFC5735	0.0.0.0/8	8-32		discard	
15	RFC5735	10.0.0.0/8	8-32		discard	
16	RFC5735	127.0.0.0/8	8-32		discard	
17	RFC5735	169.254.0.0/16	16-32		discard	
18	RFC5735	192.0.0.0/24	24-32		discard	
19	RFC5735	192.0.2.0/24	24-32		discard	
20	RFC5735	192.88.99.0/24	24-32		discard	
21	RFC5735	192.18.0.0/15	15-32		discard	
22	RFC5735	198.51.100.0/24	24-32		discard	
23	RFC5735	203.0.113.0/24	24-32		discard	
24	RFC5735	224.0.0.0/4	4-32		discard	
25	RFC5735	255.255.255.255			discard	
26	RFC5735				return	
27	def-discard	0.0.0.0/0			discard	
28	def-discard				return	

Traffic Engineering, Policy Routing

- BGP Attribute
 - http://wiki.mikrotik.com/wiki/Manual:BGP_Best_Path_Selection_Algorithm
- Routing scenario, multihomed
 - Redudancy
 - Load sharing
 - Local traffic goes to and from local peer

References

- NANOG / APNIC BGP Tutorial
- BGP Filtering with Router OS – 2013 MUM Croatia by W. Maia
- Routing Security – 2011 MUM Hungary by W. Maia
- BGP and OSPF Implementations – 2011 MUM Hungary by D. Burgess

Configuration

Access Router

- Plain Static Routing for customer
- Bandwidht Manager
- Controlled by Management Server

Remote Access Server (RB750)


- Secure VPN PPTP/L2TP
- OOB - Connection from other ISP

Configuration

Bandwidht Manager

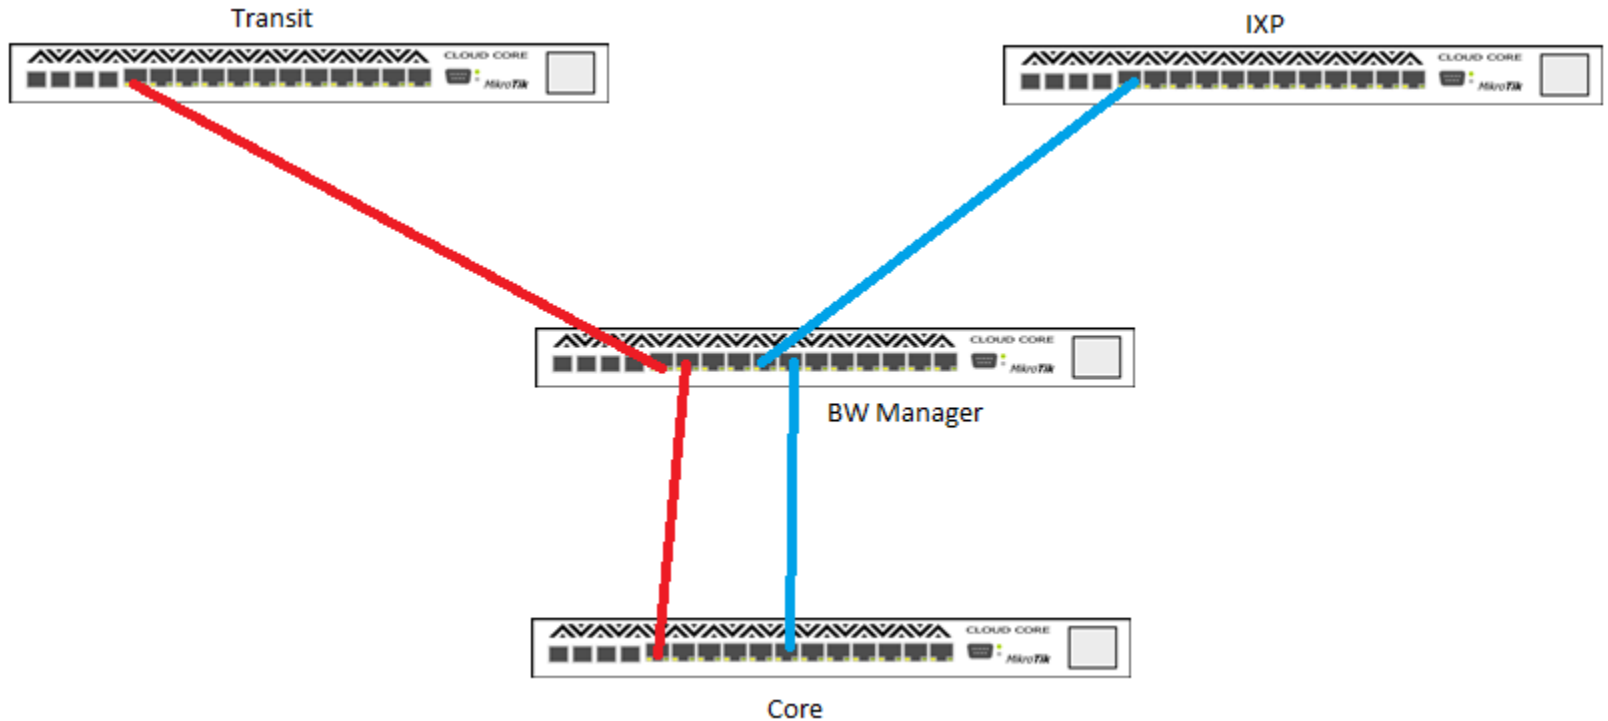
- Strategy
 - Mark packet came from AC router for Upload
 - Mark packet came from TR/IX router for Download
 - Done at Core Router
- International / Local Simple Queue / Queue Tree
- You can use transparent traffic limiter
 - <http://wiki.mikrotik.com/wiki/TransparentTrafficShaper>

Configuration



- Bridge / Routing Configuration
- International / Local Management
 - Routing List, ref :
http://mikrotik.co.id/artikel_lihat.php?id=23
 - Custom Scripting -> export routing from bgp router

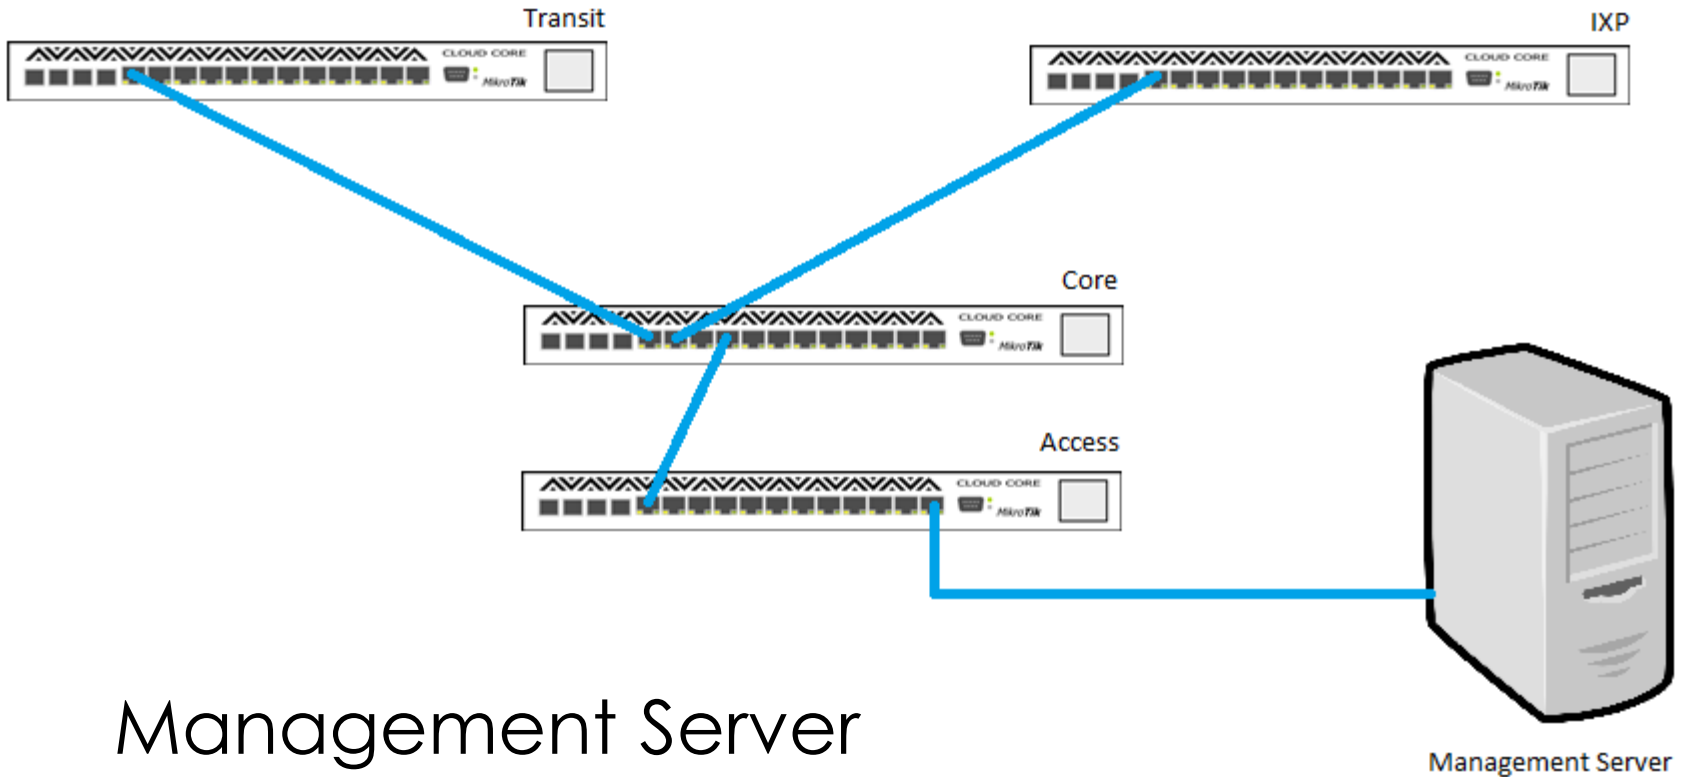
Configuration



Transparent

- Create Bridge Interface
- Marking, check packet flow diagram

Configuration



Management Server

- Don't touch my router
- Simple Mikrotik ROS API Call
- Automatic IP / VLAN / BW Allocation
- Automatic client activation / cut-off

Screenshot

Transit Router

BGP

Instances VRFs Peers Networks Aggregates VPN4 Routes Advertisements

+ - ✓ ✗ 📄 🔍 Refresh Refresh All Resend Resend All

Name	Instance	Remote Address	Remote AS	Multihop	Route Reflect	TTL	Remote ID	Uptime	Prefix C...	State
CORE	default	193.10.12.240	20012	no	no	default	193.10.12.240	152d 11:14:03	1	established
SDI	default	203.10.20.101	50002	no	no	default	193.10.12.240	27d 11:23:58	579226	established
BIZNET	default	193.10.20.100	17001	no	no	default	203.10.20.105	86d 02:28:20	597666	established

IX Router

BGP

Instances VRFs Peers Networks Aggregates VPN4 Routes Advertisements

+ - ✓ ✗ 📄 🔍 Refresh Refresh All Resend Resend All Fin

Name	Instance	Remote Address	Remote AS	Multi...	R...	TTL	In Filter	Out Filter	Remote ID	Uptime	Prefix ...	State
CORE	default	193.10.12.240	20012	no	no	default	CORE-in	CORE-out	193.10.12.240	145d 01:48:02	1	established
BIX	default	203.10.20.101	50002	no	no	default	BIX-in	BIX-out	203.10.20.101	4d 02:39:30	435	established
IIX1	default	193.10.20.100	50002	no	no	default	IIX1-in	IIX1-out	193.10.20.100	85d 17:02:05	9881	established
IIX2	default	193.10.20.105	50002	no	no	default	IIX2-in	IIX2-out	193.10.20.105	03:49:59	9984	established
IIX3	default	193.10.20.100	50002	no	no	default	IIX3-in	IIX3-out	193.10.20.100	03:49:43	10049	established
OIXP	default	193.10.20.101	50002	no	no	default	OIXP-in	OIXP-out	193.10.20.101	4d 11:08:20	11622	established
SDIIX	default	203.10.20.100	50002	no	no	default	SDIIX-in	SDIIX-out	203.10.20.100	66d 13:51:53	52057	established

Maintenance

ROS upgrade strategy

- Use stable/current only
 - RouterOS current release 6.XX
 - RouterOS bugfix release 6.XX.Y
- Read Changelog
 - Upgrade wisely
- Improving system stability

Config backup

- Simple script

Documentation

- Everything
- Log / Syslog ex: syslog-ng

additional presentation

How do the youngest country in the world ISPs run their bussiness more efficient, and more reliable with Mikrotik

History of iNet Timor

- 1999 – Referendum for Freedom
- 2000 – Telstra start cellular telephone
- 2003 – Timor Telecom : Voice (GSM/PSTN)
Telstra iNet : Data Internet
(ADSL/Dialup/Wireless)

Before Mikrotik

Network scale

- 30Mbps Upstream
- One main hub
- Dialup / ADSL Services
- 3 Wireless BTS around Dili
- VSAT Backbone
- 20 Client

Using well known product

- Cisco Router
- Cisco Switch
- Nortel/Paradyne DSLAM
- Avaya / Cisco / Breezecom
- Cisco 800 / 2500 CP Router
- Airlive CPE

Past



Problem

- Power line quality are bad, devices easy to damage
- Time to deliver replacement devices
 - From HQ (2 weeks)
 - From order to deliver, 15 day minimum
- High down time
- Expensive, almost impossible to have spare
- High cost CPE

Mikrotik

- **2006 – 2007**
 - RB230, RB132, RB133, RB532 (RouterOS v2) as Wireless Infrastructure
- **2008**
 - RB1000 (RouterOS v3), as experimental access router
- **2009**
 - RB750 as CPE router, replace Cisco 700,800,2500
- **2010**
 - RB1100 as Edge Router – Cisco replacement
 - ✓ BGP/OSPF (One default route only, One Full Routing Table)
 - RB1100 as Bandwidth Manager
 - ✓ HTB, good but complicated
 - ✓ Simple Filter Rule
 - RB1100 – As Distribution Router
 - ✓ Plain static routing
- **2013**
 - Next step with CCR1036
 - ✓ Using Simple Queue as RouterOS 6, lot easier than HTB and faster

After Mikrotik

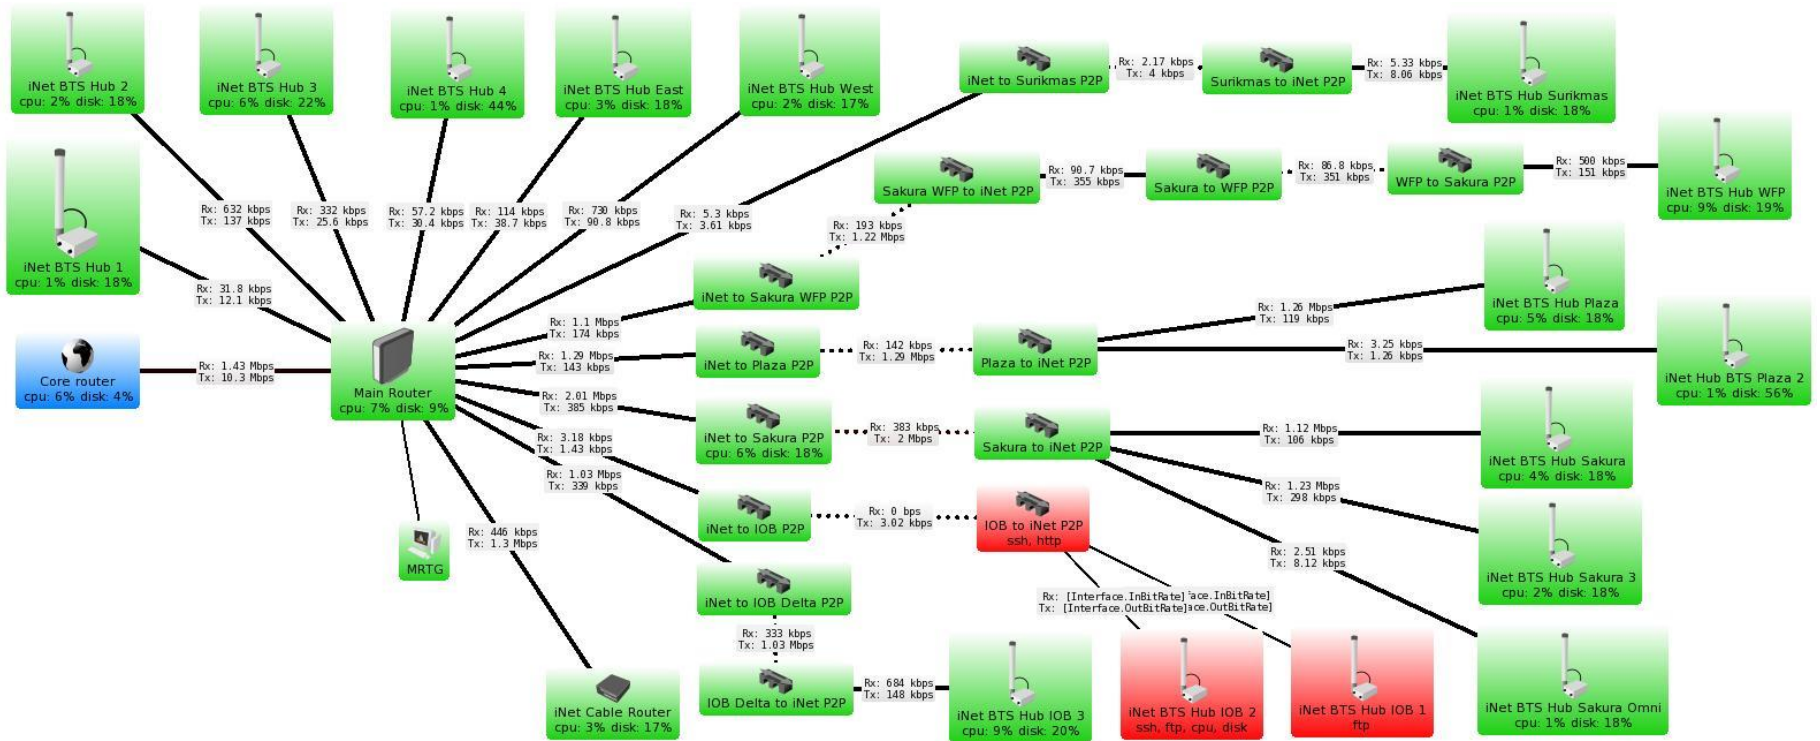
Network scale

- 150Mbps Upstream
 - VSAT & Fiber Multihoming
- Wireless & GPON Services
- 16 Wireless BTS around Dili, 5 remote BTS
- > 400 Client

Expansion

- 3G/4G LTE Access Point (Canceled 2013)
- GPON FTTH, Mikrotik ONT/ONU only
- Solar Powered remote BTS
 - 100W Solar panel + Battery
 - GSM remote switch
 - RB750UP / hEX POE Lite for controller
 - RB433, RB911, Metal






























The Dude



CCR1036

400 client simple queues ?

No Problem, we did that

39	 California Hotel	1500k	2M
29	 City Auto	1M	1M
57	 Cooyadygs	256k	256k
41	 Covec Compound	2M	2M
24	 Covec Engineering	1M	1M
109	 Covec Group	3M	3M
1	 Covec House1	768k	768k
44	 Covec House2	768k	768k
111	 Covec Office1	2M	2M
22	 Covec Office2	2M	2M
60	 DN Movers	384k	384k
55	 Dai Nippon Lda	768k	768k
16	 Dili Beach	2M	2M
95	 Dive Center	512k	512k
2	 EDTL	768k	768k
18	 ETT Groups	5M	5M
15	 BK ANZ	1M	1M
93	 BK Airport	1M	1M
43	 BK Dili Port	1M	1M
101	 Discovery Hotel	2M	2M
98	 ETT	4M	4M
92	 Gloria Jeans	2M	2M
9	 Sakib House	1M	1M
82	 Eastern Burger	1M	1M
83	 Eugenia Compound	384k	384k
70	 Fernanda House	768k	768k
58	 Fundasaun Mahein	1M	1M
4	 GFA-Lahane	512k	512k
79	 GMP-TL Group	3M	3M

421 items 1 selected) 0 B queued

Result

- Reducing Expenses
 - Reducing capex
 - Reducing customer cost
- Fast to deliver, Fast to replace, Low down time
 - We have cold spare devices
 - If we dont, we can get it less than 12 hours
- Easy to operate and maintain
 - Winbox easy to use
- Open lot of possibilities,
 - Exploit all of technology available on ROS

Last Update

iNet start using O3b

- Medium orbit Satellite, not geo stationer
- 360 minutes contact per satellite
- 300ms latency, not 500ms anymore
- Required two autotracking dish
- Using Mikrotik to do VRRP

Last..

- Don't be affraid to use RouterOS on your ISP
- Don't be embarassed if you already use ROS
- Router OS have complete features for ISP

Thank You