GrangeNet

Advanced Communication Services

Dr Greg Wickham & Chris Myers

grangenet

IPv6 Workshop
4 JULY 2004
Version 2.0

Original slides thanks to:

Abilene NOC at Indiana University
Grover Browning (gcbrowni@iu.edu)
Brent Sweeny (sweeny@iu.edu)
John Barlow (ex GrangeNet)

grangenet

IPv6 Workshop
4 JULY 2004
Version 2.0

Equipment supplied by
CISCO Systems
Juniper
ANU
GrangeNet
AARNet

Thank you for your support

grangenet

IPv6 Workshop
4 JULY 2004
Version 2.0

→ Goals

- Design and set up a functioning, globally connected, IPv6 network.
- After the workshop you should be able to engineer IPv6 networks within your campus, division or faculty.
- Topics covered:
  - Router Configuration: How do you turn on IPv6 routing on a router?
  - BGP Configuration: Configure BGP sessions between the campus and core routers.
  - Addressing: How addressing works; what types of addressing schemes are possible; how allocations are made; and current best practice.
  - Transition Issues: How do you make your IPv6 network IPv4-aware and vice versa? How extensively can you do dual stack implementations and where do you need to do translation?
  - Services: What IPv6 services should a campus or RNO currently offer? What are the best known practices for distributing IPv6 throughout a campus or state network?

grangenet

→ Services Location

- Access to building & room
- Parking & transport
- Toilets
- Food and drink location
- Instructors
  - Chris Myers 0423 784 265
  - Dr Greg Wickham 0407 854 566

grangenet

→ Schedule

- 9:00am to 12:30pm
  - IPv6 workshop intro
  - IPv6 Fundamentals and Lab
- Lunch

grangenet

IPv6 and GrangeNet

- Native IPv6 Connectivity offered to all connected members
- Deployment Status
  - Queensland GigaPoP enabled
  - New South Wales GigaPoP enabled
  - Australian Capital Territory GigaPoP enabled
  - Victoria GigaPoP enabled
- Tunnel IPv6 from overseas Via AARNet.
- IPv6 Working Group setup with AARNet
  - http://www.grangenet.net/status/wp/index.html
  - http://www.aarnet.edu.au/engineering/wgs/ipv6/

grangenet

Who's in the room

- Introductions – around the room
  - Who we are
  - Why are we here

grangenet

IPv6 Fundamentals

- Exercise #1
  - Address allocation
- Exercise #2
  - Enable IPv6 and test
- Exercise #3
  - BGP for IPv6

grangenet

Exercise #1

- Goal
  - Address allocation
- Procedure
  - IPv6 Address Overview
  - Create lab allocation of IPv6 addressing

grangenet

General approach

128 bits

grangenet

General approach

Network address (48 bits)

grangenet

General approach

Network address (48 bits)		EUI host address (64 bits)

grangenet

General approach

Network address (48 bits)	16 bits	EUI host address (64 bits)

16 bits are left for subnetting—what to do with them?

grangenet

Approaches to v6 subnetting

1. Sequentially:
   0000
   0001
   …
   FFFF

   16 bits = 65535 subnets

grangenet

Approaches to v6 subnetting

1. Sequentially
2. Follow IPv4 addressing:

   Subnets or combinations of nets & subnets, or VLANs, etc., e.g.

   - 128.8.**60**.n/24 → 003c
   - 129.9.**91**.n/24 → 005b
   - 130.10.**156**.n/24 → 009c
   - **156.56.60**.n/24 vs **129.79.60**.n/24?
     - 013c or 383c or 9c3c vs 023c or 4f3c or 813c

grangenet

Approaches to v6 subnetting

1. Sequentially
2. Follow IPv4 addressing
3. Topological/aggregating:

   perhaps reflecting wiring plants, supernets, large broadcast domains, etc. e.g.

- Main library = 0010/60
  - Floor in library = 001a
- Computing center = 0020/55
  - Student servers = 002c
- Medical school = 00c0/50
  - Etc.…

grangenet

New/different things to think about

- You can use “all 0s” and “all 1s”! (0000, ffff)
- You’re not limited to 254 hosts per subnet!
  Switch-rich LANs allow for larger broadcast domains (with tiny collision domains), perhaps thousands of hosts/LAN…
- No “secondary subnets” (though >1 addr/interface)
- No tiny subnets either (no /30, /31, /32)—plan for what you need for backbone blocks, loopbacks, etc.
- *Remember to use /64 if you want autoconfigure to work!*

grangenet

IPv6 Prefixes

Subnet Masking, VLSM, & Aggregation

- 2001:388::/32
- 2001:0388:0000:0000:0000:0000:0000:0000
- 2001:388:305c::/48
- 2001:0468:305c:0000:0000:0000:0000:0000
- 2001:388:305c:1::1/64
- 2001:0388:305c:0001:0000:0000:0000:0001

grangenet

Now try it yourselves

You have the allocation 2001:388:305C::/48 at RNO

- Allocate a block of addresses for each of:
  - –your POD uplinks
  - –your point-to-point inter pod interfaces (we will do it together)

grangenet

Exercise #2

- Goal
  - –Enable IPv6 and test
- Procedure
  - –Enable IPv6 on the routers
  - –Apply static address to each pod
  - –Enable IPv6 on your laptops
  - –Ping check the local router

grangenet

Enable IPv6 on the routers

- Enable IPv4
- Turn On Routing
- Enable IPv4
- Configure Interfaces

- Enable IPv6
- Turn On Routing
- Enable IPv6
- Configure Interfaces

grangenet

Enable IPv6 on the routers

- Step 1
- Logon to router
  - –Username APAN
  - –Password APAN
- Enter configuration mode
  - –Enable password APAN

grangenet

Enable IPv6 on the routers

- configure terminal
- IPv6 unicast-routing (Global)
- Ipv6 cef (Global)
- Interface fastethernet or ethernet 0/0
- IPv6 enable (Interface)
- IPv6 address <address> (Interface)
- Exit
- Exit
- write memory

grangenet

Apply static address to each pod

- Get back into configuration global mode
- And then add
  - IPv6 route <address> <hop address>
- For your pod and others
- Check you work with
  - show IPv6 route
  - Show IPv6 interface

grangenet

Enable IPv6 on your laptops

- On Windows XP
  - IPv6 Install
- Common commands
  - IPv6 ? – help on IPv6 commands
  - IPv6 rc – View the route cache
  - IPv6 nc – View the neighbor cache
  - IPv6 if – View interface information
  - Ping6 – ping IPv6 address
  - Netsh – command shell for network configuration

grangenet

Ping check the local router

- Using ping check the IPv6 interfaces
- Router
  - Ping IPv6 <address>
- Windows
  - Ping6 <address>

grangenet

Exercise #3

- Goal
  - BGP for IPv6
- Procedure
  - Enable BGP between pod and upstream
  - Confirm BGP is working
  - Add BGP Prefix filtering
  - Final tests (Dancing turtle!)

grangenet

Enable BGP between pod and upstream

- Router bgp <ASN>
- Address-family IPv6 unicast
- Neighbor <v6 addr> remote-as <ASNX>
- Network 2001:388:305c:XXXX::/64

grangenet

Confirm BGP is working

- Cisco show commands
- show bgp sum ! show summary of neighbors' BGP state
- show bgp ! show all v6 BGP-learned routes
- show bgp neigh [addr] routes ! what he's *sending*
- show bgp neigh [addr] advertised ! what *you're* sending
- show ipv6 route ! show all v6 routes

  – What routes are you sending? Receiving?
  – Do you see all local IPv6 nets?
  – Do you see workshop IPv6 nets?

grangenet

Added BGP Prefix-filtering

BGP Prefix-filtering (at your upstream facing you):

- –IPv6 prefix-list From-PODX permit 2001:388:305c:XXXX:/64

- And to apply it:
  - –neighbor 2001:388:305c:xxxx::1 prefix-list From-PODX in

grangenet

Final tests (Dancing turtle!)

- http://www.kame.net -- The “kame” or turtle at the top of the main page “dances” if you are connected via IPv6
- http://ipv6.research.microsoft.com -- Accessible only via IPv6

grangenet

Overview

- Overview
  - –Equipment needs
  - –Where are the Applications?
  - –DNS Issues
  - –Traffic types
  - –ICMP Types
  - –Stateless AutoConfiguration
  - –Predefined Scopes
  - –Router software versions
  - –Campus network overview

grangenet

Equipment Needs

- For a simple campus IPv6 implementation:
  - –One Router (Cisco 2600) ~$4,000
  - –One Linux Box (DNS etc) ~$2,000
- Or it could all be combined in a single PC with:
  - –Linux or BSD
  - –Zebra (BGP routing)
  - –DNS etc.

grangenet

Where are the applications?

- ‘bread-and-butter’ apps:
  - –telnet
  - –ftp
  - –http
  - –Sendmail, various pop/imap clients (one that works for win32 is “postie”, from http://www.infradig.com/infradig/download.shtml)
  - –dns/bind
  - –nntp (news)
  - –ssh
- Also:
  - –Pchar
  - –Netperf
  - –A growing list…

grangenet

DNS Issues

- BIND Versions
  - –v8 for ipv6 address support, v9 for dual-stack
- Dual-Stack Support
- Production System Support
- Separate (experimental?) v6 DNS servers?

grangenet

Traffic Types

- Unicast
- Multicast
- Anycast

- Aiiiieeeee! Where did Broadcast go?

grangenet

ICMP Types

1 – Destination Unreachable
2 – Packet Too Big
3 – Time Exceeded
4 – Parameter Problem
128 – Echo Request
129 – Echo Reply
130 - Group Membership Query
131 – Group membership Report
132 – Group Membership Reduction
133 – Router Solicitation
134 – Router Advertisement
135 – Neighbor Solicitation
136 – Neighbor Advertisement
137 - Redirect

grangenet

Stateless AutoConfiguration

- *Only* works with /64!
- Build a Link Local Address (Split MAC address, put FFFE in Middle.)
- Join All Nodes Multicast Group (FF02::1)
- Send Solicitation ICMP to All Routers (FF02::2) HOP set to 255.
- Router responds with a Router Advertisement ICMP.
- Prefixes have a Valid lifetime and a Preferred Lifetime.

grangenet

Predefined Scopes

FE80::/64	Link Scope
FEC::/10	Site Scope
FF01::	Node Scope
FF02::	Link Scope
FF05::	Site Scope
FF08::	Organization Scope
FF0E::	Internet Scope

grangenet

Router software versions

- Cisco Router Support
  - Cisco Main stream 12.3(3)

grangenet

Campus network overview

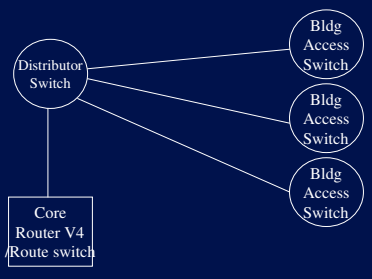

grangenet

Campus network overview

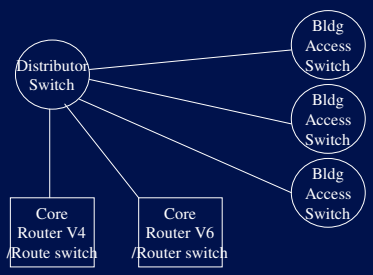

grangenet

Campus network overview

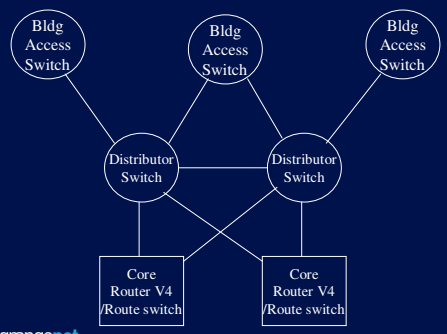

grangenet

Campus network overview

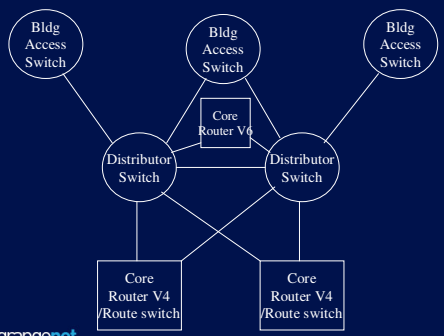

grangenet

Layer-3 Campus

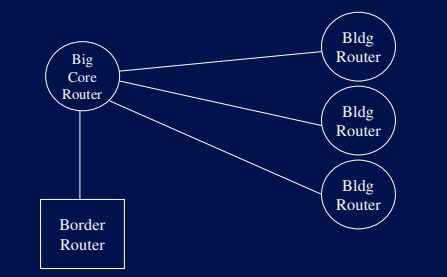

grangenet

Questions?

Thankyou

ACS@grangenet.net

http://www.grangenet.net

grangenet