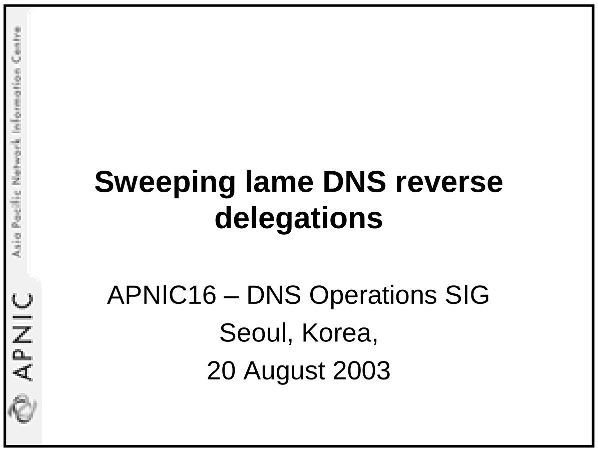

Sweeping lame DNS reverse delegations

APNIC16 – DNS Operations SIG
Seoul, Korea,
20 August 2003

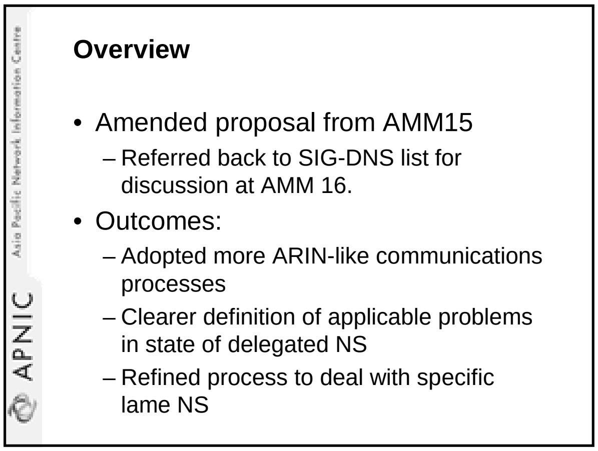

Overview

- Amended proposal from AMM15
  - Referred back to SIG-DNS list for discussion at AMM 16.
- Outcomes:
  - Adopted more ARIN-like communications processes
  - Clearer definition of applicable problems in state of delegated NS
  - Refined process to deal with specific lame NS

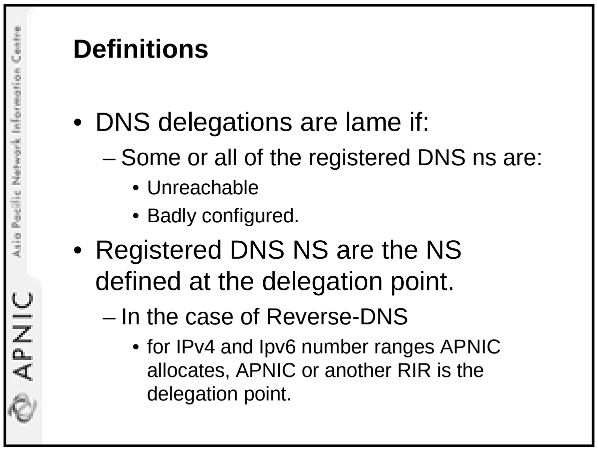

Definitions

- DNS delegations are lame if:
  - Some or all of the registered DNS ns are:
    - Unreachable
    - Badly configured.
- Registered DNS NS are the NS defined at the delegation point.
  - In the case of Reverse-DNS
    - for IPv4 and Ipv6 number ranges APNIC allocates, APNIC or another RIR is the delegation point.

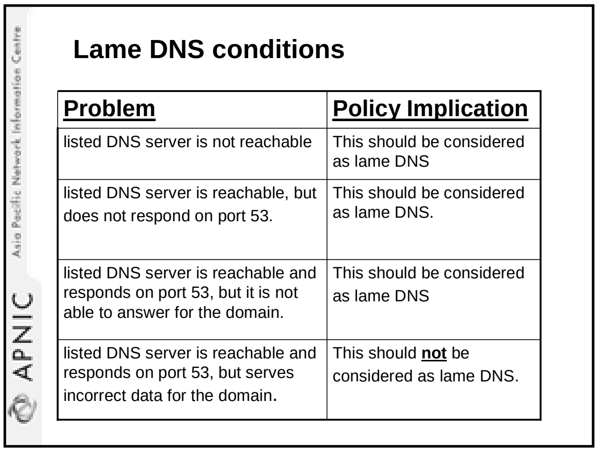

Lame DNS conditions

Problem	**Policy Implication**
listed DNS server is not reachable	This should be considered as lame DNS
listed DNS server is reachable, but does not respond on port 53.	This should be considered as lame DNS.
listed DNS server is reachable and responds on port 53, but it is not able to answer for the domain.	This should be considered as lame DNS
listed DNS server is reachable and responds on port 53, but serves incorrect data for the domain.	This should **not** be considered as lame DNS.

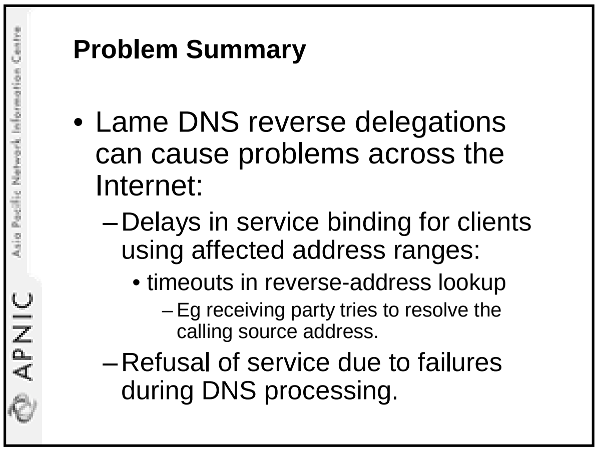

Problem Summary

- Lame DNS reverse delegations can cause problems across the Internet:
  - –Delays in service binding for clients using affected address ranges:
    - timeouts in reverse-address lookup
      - –Eg receiving party tries to resolve the calling source address.
  - –Refusal of service due to failures during DNS processing.

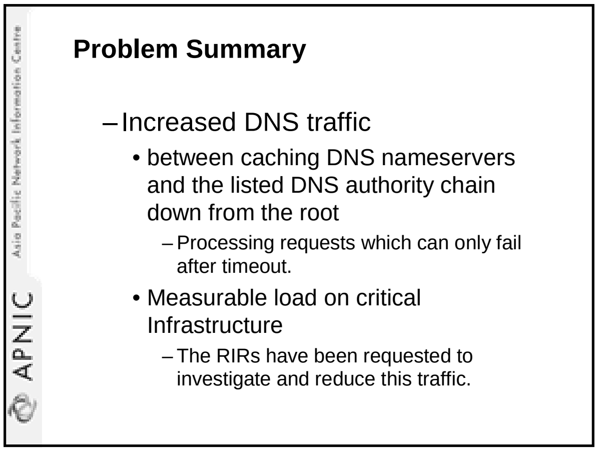

Problem Summary

- –Increased DNS traffic
  - between caching DNS nameservers and the listed DNS authority chain down from the root
    - –Processing requests which can only fail after timeout.
  - Measurable load on critical Infrastructure
    - –The RIRs have been requested to investigate and reduce this traffic.

Problem Summary

- Lame DNS reverse delegations affect
  - The users of the network in question.
  - Unrelated third parties.
- End users cannot resolve problem directly
  - Due to hierarchical nature of authoritative delegation.
- If the network administrators do not correct errors in their DNS configurations
  - RIR has to resume control of the delegated domain (pending delegate resuming control)
  - disable the listing of the misconfigured servers.

Proposal

- Identify potential lameness.
  - (two points of test, AU & JP)
- Test the DNS reverse delegation
  - (15 day test period).
- Attempt to notify the domain holder
  - (45 day notice period).
- Disable lame DNS reverse delegation.
  - (If not corrected at end of notice period)

Identify potential lameness

- Modified process
  - based on scripts used for current statistical measurement exercise
  - Run independently in Japan and Australia
  - Mark each delegated NS separately for status
  - Collate state at HQ nightly to compute aggregated lameness state
    - (pass/fail value, not lame at either location == pass)
      - Prevents single-point failure in test

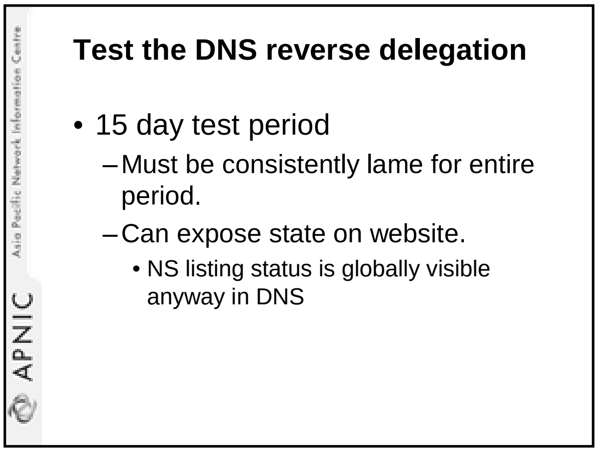

Test the DNS reverse delegation

- 15 day test period
  - Must be consistently lame for entire period.
  - Can expose state on website.
    - NS listing status is globally visible anyway in DNS

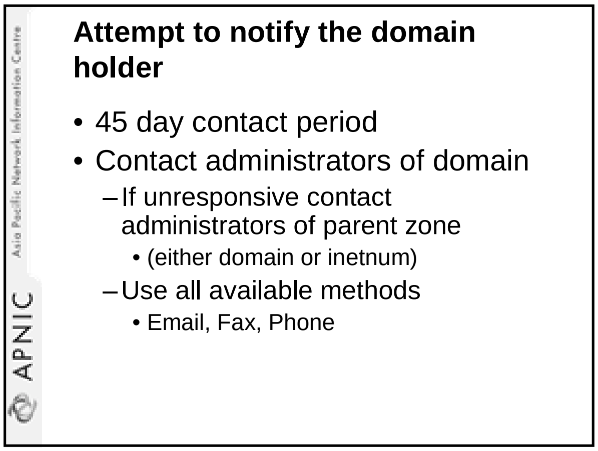

Attempt to notify the domain holder

- 45 day contact period
- Contact administrators of domain
  - If unresponsive contact administrators of parent zone
    - (either domain or inetnum)
  - Use all available methods
    - Email, Fax, Phone

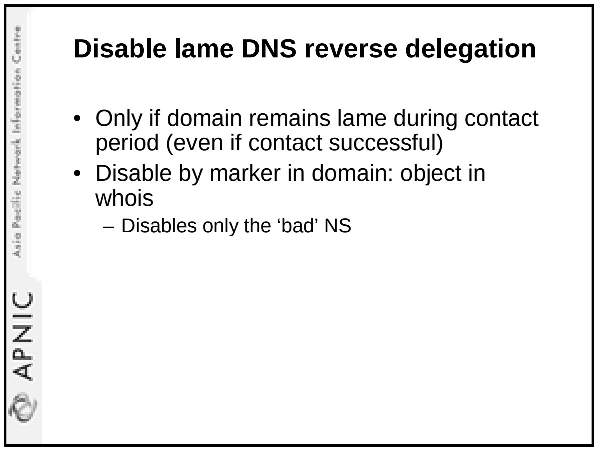

Disable lame DNS reverse delegation

- Only if domain remains lame during contact period (even if contact successful)
- Disable by marker in domain: object in whois
  - Disables only the ‘bad’ NS

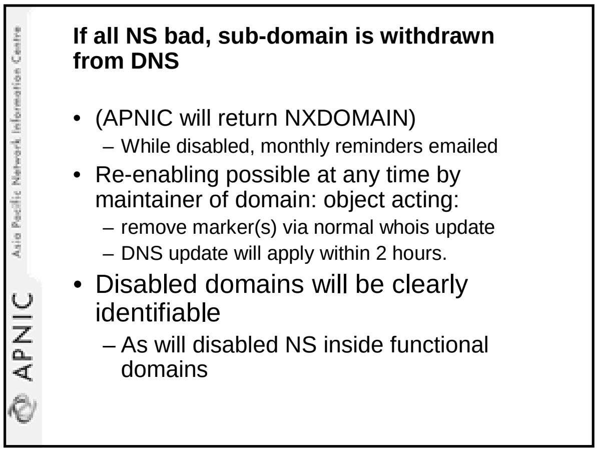

If all NS bad, sub-domain is withdrawn from DNS

- (APNIC will return NXDOMAIN)
  - While disabled, monthly reminders emailed
- Re-enabling possible at any time by maintainer of domain: object acting:
  - remove marker(s) via normal whois update
  - DNS update will apply within 2 hours.
- Disabled domains will be clearly identifiable
  - As will disabled NS inside functional domains

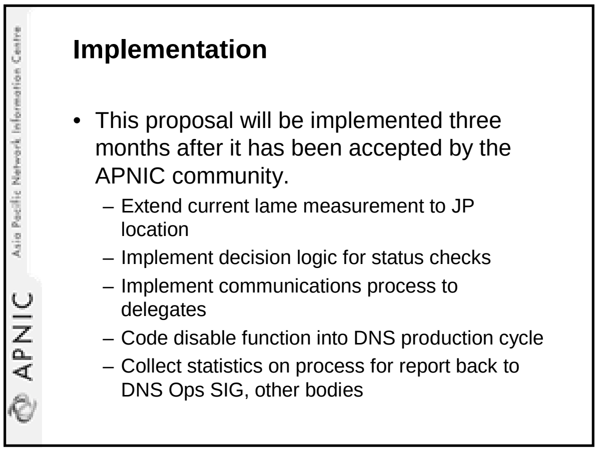

Implementation

- This proposal will be implemented three months after it has been accepted by the APNIC community.
  - Extend current lame measurement to JP location
  - Implement decision logic for status checks
  - Implement communications process to delegates
  - Code disable function into DNS production cycle
  - Collect statistics on process for report back to DNS Ops SIG, other bodies

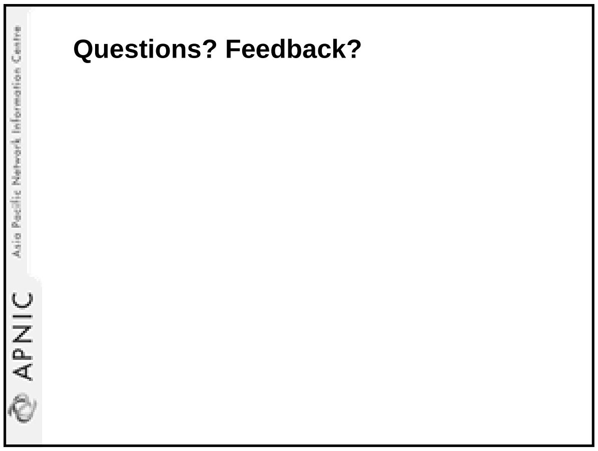

Questions? Feedback?