Nominum

DNS Performance – Areas of improvement?

Request for Discussion

APNIC 20, DNS SiG, September 8th, 2005

Mathias Körber

Nominum, Inc

Introduction

Nominum.

- Mathias Körber

  Senior Consulting Engineer
  Nominum, Inc
  Mathias.Koerber@nominum.com

  - based in Singapore
  - Support and Consulting, mostly Asia

- Background
  - 1995-2000 SingNet
  - 1996-2002 affiliated with SGNIC
  - 2000-2002 Nominum, Inc
  - 2003-2004 Lightspeed Technologies, Singapore
  - 2004-present Nominum,Inc

Nominum Overview

Nominum.

History

- Founded in 1999
- Located in Silicon Valley
  - Growing presence in EMEA and APAC
- Venture funded by top-tier investors
  - ATV, Bessemer, Globespan, Morgenthaler, VeriSign, Juniper Networks

Focus

- Commercial–grade IP name and asset management
  - DNS, DHCP, and IP management products
- For medium to large deployments
  - Telco, service providers, finance, retail, government, etc.
- Technology leadership
  - Performance
  - Security
  - High availability

Unsurpassed Domain Expertise

Nominum.

People

- **Dr. Paul Mockapetris**, *Chief Scientist and Chairman*
  - Inventor of DNS
- **David Conrad**, *Chief Technical Officer and Founder*
- **Ted Lemon**, *Senior Architect*

Products

- Developed BIND 9 and ISC-DHCP v3

Operations

- Operational support for Root Servers E and F

RFC Authored by Nominum Employees

RFC 882
RFC 883
RFC 973
RFC 1034
RFC 1035
RFC 1101
RFC 1183
RFC 1413
RFC 1414
RFC 1693
RFC 2050
RFC 2154
RFC 2669
RFC 2670
RFC 2786
RFC 2845
RFC 3007
RFC 3008
RFC 3074
RFC 3127
RFC 3225
RFC 3315
RFC 3396
RFC 3442
RFC 3597
RFC 3639
RFC 3655

The Problems we Solve

Nominum.

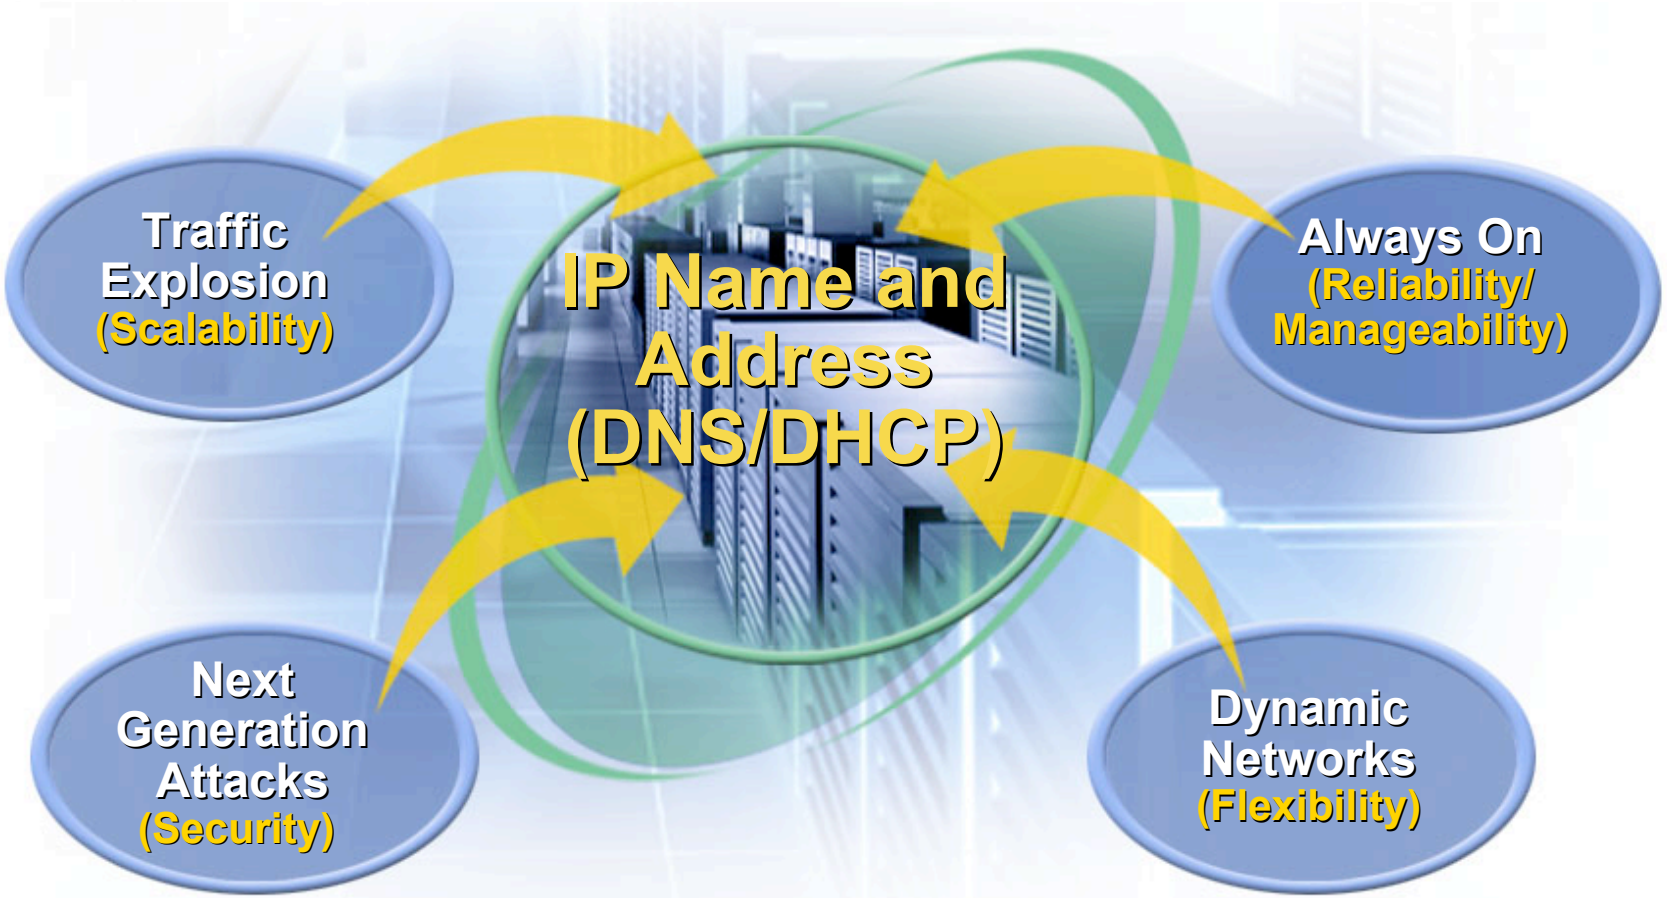

Sample Customers


Product Family

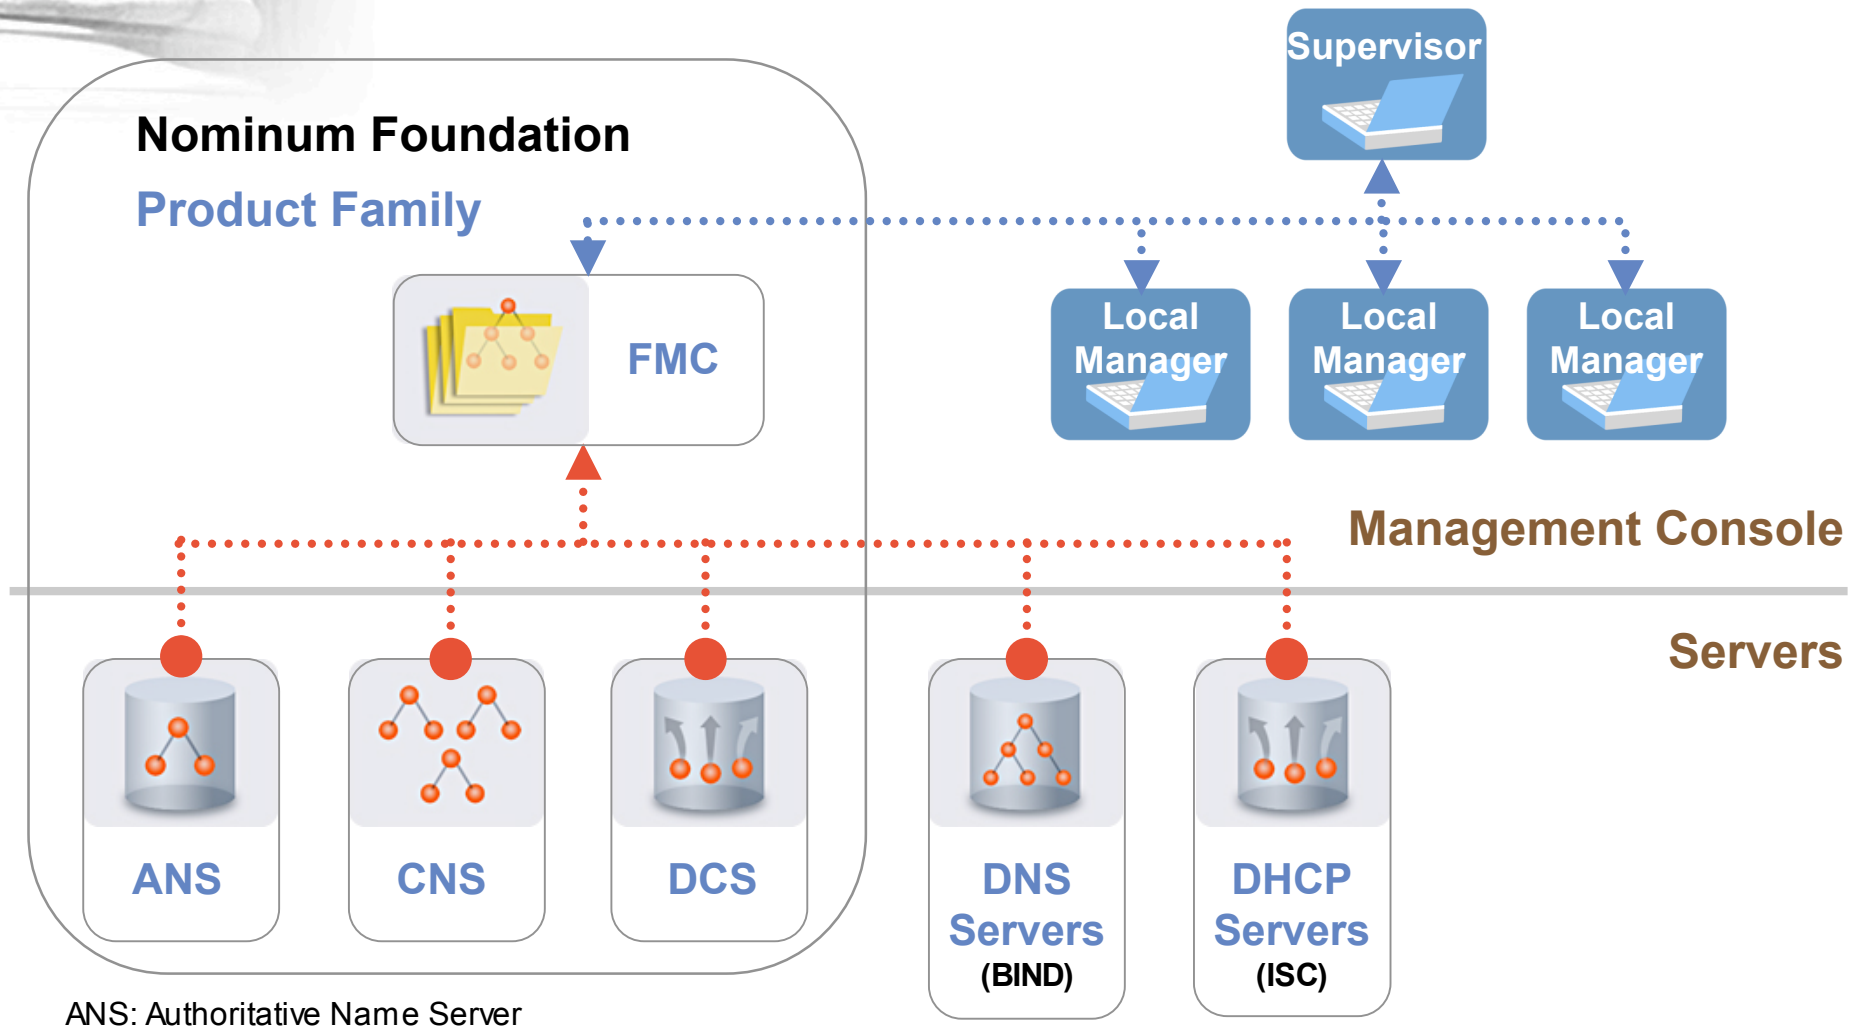

ANS: Authoritative Name Server
CNS: Caching Name Server
DCS: Dynamic Configuration Server
FMC: Foundation Management Center

Increased Demands on DNS

Nominum.

Recent (and not so recent) developments are posing increased demands on the DNS

- eg:
  - Higher query loads
  - Larger data sizes
  - More frequent updates/changes
  - Need for faster processing

Higher Query loads

Nominum.

- The internet is experiencing ever increasing DNS traffic

Part of it is due to legitimate reasons

Part is due to mistakes/accidents

and part due to malicious (ab)use

Higher Query loads - Some causes:

Nominum.

- malicious
  - Spreading viruses/trojans
  - Phishing/pharming
  - (D)DOS attacks
    - » Against nameservers themselves
    - » other parts of networks
- accidental/unintended
  - » misconfigurations
  - » windows ‘SRV’ updates etc escaping
  - » unnecessarily low TTLs
- legitimate
  - Web:
    - » more/diverse links per page
    - » access verification/logging
  - Automation
    - » RSS feeds
  - Email
    - » RBL checks
    - » Spam-checks

Higher Query loads (cont)

Nominum.

- Eg Spam checks
  - Spam is a “killer” DNS application
    - 75% of ISP DNS requests are MX Lookups (Mail Routing); Non-existent domains increase load
  - Spam itself is ‘malicious’
  - Spam checks are legitimate
  - Many spam checks require frequent DNS lookups (RBL, valid sender domain etc)

Effects or higher query loads

Nominum.

- on Caching nameservers
  - increased memory usage
  - increased CPU usage
  - low cache efficiencies
- Higher rate of dropped queries
  - clients will retry -> ever higher loads
- Higher latencies
  - interactive experience suffers
  - non-interactive processes slow down
- Less available overhead to deal with
  - additionally high traffic situations

More data in the DNS

Nominum.

- Increased domain registrations
  - Multilingual domains in addition to ‘English’ ones
- New Technologies
  - IPv6
    - Larger RR sizes
    - Larger zone depth (more zone cuts)?
  - DNSSEC
    - Highly increased zone sizes
  - ENUM
    - At least 2 NS per telephone number
      - Likely many more depending on customer usage and ENUM scheme

More data in the DNS - Effects

Nominum.

- Larger zones
- Larger RRsets
- Deeper hierarchies

More data in the DNS - Effects :

Nominum.

- On caching nameservers
  - Reduced cache efficiencies
  - Higher latencies (cache lookup times)
  - Deeper hierarchies mean
    - more recursions
    - Verification of
      - GLUE segregation
      - DNSSEC signatures

More data in the DNS - Effects

Nominum.

- On Authoritative NS:
  - Higher memory requirements
    - Large zones may not even fit
  - Longer startup/restart times
  - Large zonetransfers
    - IXFR mitigates this somewhat

  => Reduced performance

More frequent updates

Nominum.

- DNS becomes more ‘dynamic’
  - Increased use of DDNS
    - Mobile clients
    - DHCP servers updating dynamic client info
  - Faster domain registrations
  - Self-help DNS interfaces allowing individual changes
  - Number portability (ENUM)

More dynamic DNS - Effects

Nominum.

- Increased demands on master servers
  - DDNS
    - prerequisite checking
    - Integrity checking
    - atomic updating of changes into zones
  - Frequent reloading of changed zones
    - possible service interruption during each reload
    - increased memory requirements and processing
- Increased master-slave traffic
    - NOTIFY
    - IXFR/AXFR
- Increased query traffic
    - due to lower TTLs employed to make

Demand for faster results

Nominum.

- Telephony technologies
  - ENUM
  - SIP

  demand fast connection establishment
- DNS queries must complete in a few ms at most
- DNSSEC also requires speed

Solutions

Nominum.

- Add more caching nameservers
  - Split customer base across servers
  - anycast
  - L4 switch infrastructure
  - Others
- Disadvantages
  - added hardware costs
  - added administrative overheads
  - limited by resolver limits (# of nameservers)
- Optimized nameservers

ANS vs BIND Latency

Nominum.

ANS
Latency <5ms

BIND
Latency up to 200ms

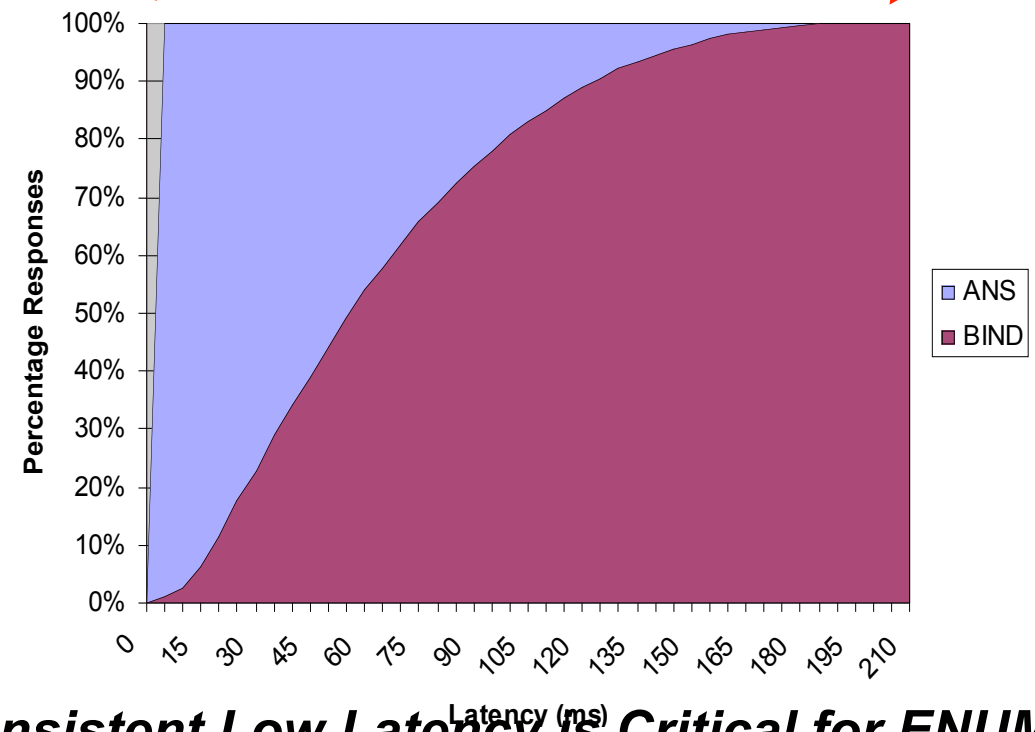

Consistent Low Latency is Critical for ENUM

Configuration:
HP DL380,
2x2.4Ghz P4 Xeons,
2GB DDR Ram,
2x36GB SCA SCSI 10K,
striped with RAID 0,
Red Hat Enterprise 3.0

Comparing BIND and ANS for ENUM

Nominum.

	BIND 9	ANS cold cache	ANS warm cache
Records in 3.5 Gigs RAM	28 MM	>200MM	>200MM
Latency	2 seconds	0.2 seconds	0.003 seconds
Queries/Sec	57	493	33,000
CPU Utilization	99% each on two CPU's	3%	12%

Configuration: 2 X AMD 1800+ CPU at 1.5GHz, 3.5 GB RAM, 100Mb Ethernet Interface
Red Hat Enterprise 3.0 Server

Solicitation

- Some problem areas identified
- Others sure to exist now or arise over time
- Some can be / have been addressed today
- How to prepare for new issues?

Solicitation

Nominum.

In what other areas do we expect ‘crunch’ ?

How can those areas be addressed

- DNS engine improvements (design, performance, usability, administrability)?
- interoperation (with what?)
- network design
- procedures

Solicitation (cont)

- Would like to learn/discuss ‘regional’
  - experiences
  - concerns
  - expectations

Solicitation (cont)

Nominum.

- Asiapac specific issues?
  - Internationalization?
  - Infrastructure issues (limited international connectivity)?
  - Growth industries?
- How should these be addressed?
- Future sessions desired (part of DNS SiG or separate)?

Contact

Mathias Körber

Available tomorrow (Friday) during APNIC member meeting

Weekend?

Email: Mathias.Koerber@nominum.com

+65 9815 7807

+1 650 381 6044