



UNDERSTANDING DHCP AND DNS

SESSION NMS-1101


Agenda

Cisco.com

- **Introduction to Names and Addresses**
- **Managing Addresses with DHCP**
 - Protocol
 - Assignment and Reliability
- **Resolving Names with DNS**
 - Protocol
 - Database
 - Reliable Operation
- **New Things**

Address Review

Cisco.com



- **IPv4 address 32 bits**
Decimal, 8-bit fields, period separation
128.9.0.33
- **IPv6 address 128 bits**
Hexadecimal, 16-bit fields, colon separation
2001:0DB8:0000:0001:02A0:C9FF:FE61:1216

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

3

Address Hierarchy and Naming

Cisco.com

- **ADDRESSES** have a topological hierarchy
- **NAMES** have a logical hierarchy
NOT NECESSARILY ALIGNED WITH EACH OTHER...



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

4

Subnet Mask

Cisco.com

Address 128.9.0.33

1000000000000100100000000000100001

Mask 255.255.255.0

11111111111111111111111100000000

- Mask separates network (1) from host (0) part of the address
- Prefix (longest match) routing—contiguous “1” bits to the left



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

5

Subnets

Cisco.com

- Each range of addresses for hosts defines a subnet e.g. 128.9.0.0/24
 - 24 is the number of ‘1’ bits in the mask for this address
 - 32–24=8 is the number of bits in host address
- Within the subnet, hosts communicate directly, using layer 2
- Special meaning for certain host addresses
 - All ones—broadcast
 - All zero—network

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

6

Special Addresses

Cisco.com

- **Multicast**
IPv4—224-239.d.d.d [RFC 2365]
IPv6—FFxx:x:x:x:x:x:x
- **Anycast [RFC 1546]**
Unicast, but with multiple advertisers
- **Site local**
IPv4—10/8, 172.16/12, 192.168/16 [RFC 1918]
IPv6—~~FE00:0:0:<subnet ID>:<interface ID>~~
- **Link local**
IPv4—169.254/16
IPv6—FE80:0:0:0:<interface ID> Removed by Decision in the ipng wg in the IETF Spring 2003
- **Loopback**
IPv4—127.0.0.1
IPv6 — 0:0:0:0:0:0:1 (::1)

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

7

Agenda

Cisco.com

- **Introduction to Names and Addresses**
 - **Managing Addresses with DHCP**
 - **Protocol**
 - Assignment and Reliability
 - **Resolving Names with DNS**
 - Protocol
 - Database
 - Reliable Operation
- **New Things**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

8

DHCP Basics

Cisco.com

- **Ideal administrator—DHCP server acts as proxy for network administrator**
- **Assignment is temporary—address is assigned with a “lease”**
- **Addresses can be reassigned when no longer in use**
- **Backup for reliability**

NMS-1101
9592_04_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

9

How DHCP Works: Obtaining an Address

Cisco.com

- **Server dynamically assigns IP address on demand**
- **Administrator creates pools of addresses available for assignment to hosts**
- **Address is assigned with lease time**
- **Client can extend lease time dynamically**
- **Server can reassign address after lease expires**
- **DHCP delivers other configuration information in options**



Here Is Your Configuration:
IP Address: 192.204.18.7
Subnet Mask: 255.255.255.0
Default Routers: 192.204.18.1, 192.204.18.3
DNS Servers: 192.204.18.8, 192.204.19.9
Lease Time: 5 days

NMS-1101
9592_04_2004_c2

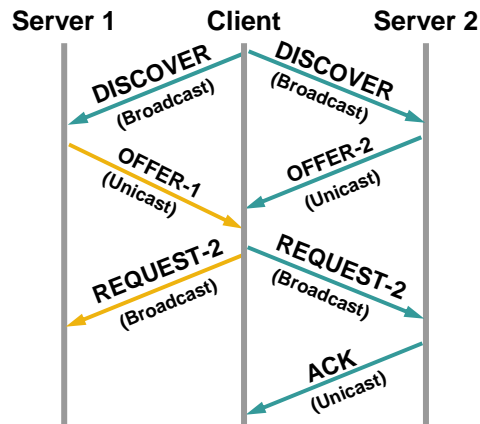
© 2004 Cisco Systems, Inc. All rights reserved.

10

How DHCP Works: Message Exchange

Cisco.com

- DHCP client broadcasts DISCOVER packet on local subnet
- DHCP servers send OFFER packet with lease information
- DHCP client selects lease and broadcasts REQUEST packet
- Selected DHCP server sends ACK packet



NMS-1101
9592_04_2004_c2

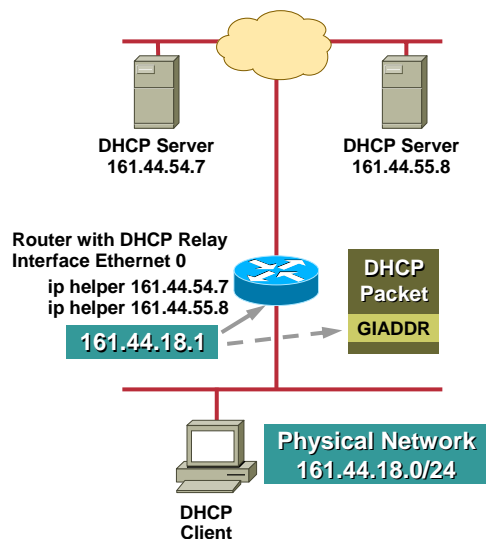
© 2004 Cisco Systems, Inc. All rights reserved.

11

DHCP Relay: Centralized DHCP Service

Cisco.com

- DHCP clients broadcasts a DISCOVER packet
- DHCP relay (IP helper address) on the router hears the DISCOVER packet and forwards (unicast) the packet to the DHCP server
- DHCP relay fills in the GIADDR field with IP address of the receiving interface of router
- DHCP relay can be configured to forward the packet to multiple DHCP servers; client will choose the "best" server
- DHCP servers use GIADDR field of DHCP packet as an index in to the list of address pools



NMS-1101
9592_04_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

12

DHCP Options for Applications

Cisco.com

- Options are registered with IANA
- Time, NIS, TCP, and IP parameters... [RFC 2131]
- Service Location Protocol (SLP) [RFC 2610]
- Novell directory services [RFC 2241]



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

13

Agenda

Cisco.com

- Introduction to Names and Addresses
- **Managing Addresses with DHCP**
 - Protocol
 - **Assignment and Reliability**
- **Resolving Names with DNS**
 - Protocol
 - Database
 - Reliable Operation
- **New Things**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

14

DHCP Reliability

Cisco.com

- **Multiple servers with split address pools**
 - Loadsharing
 - Servers answer only for configured hash (MAC)
 - RFC 3074
- **Failover**
 - Draft based on our (Cisco) design
 - Two servers can share address pools and continue to operate if one fails

NMS-1101
9592_04_2004_c2

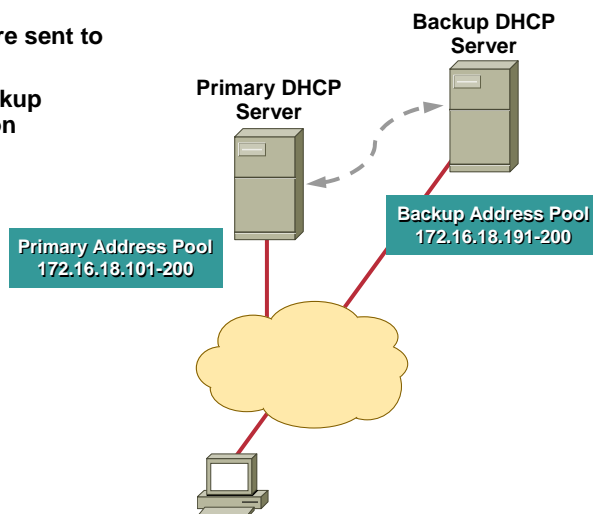
© 2004 Cisco Systems, Inc. All rights reserved.

15

DHCP Safe Failover Protocol

Cisco.com

- All DHCP requests are sent to both servers
- Primary updates backup with lease information
- Backup takes over when primary fails
- Backup server uses dedicated pool of addresses allocated by the primary to prevent duplicate IP address
- Servers synchronize when primary is up
- IETF Internet draft



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

16

How DHCP Works: DHCP Packet

Cisco.com

OP Code	Hardware Type	Hardware Length	HOPS
Transaction ID (XID)			
Seconds		Flags	
Client IP Address (CIADDR)			
Your IP Address (YIADDR)			
Server IP Address (SIADDR)			
Gateway IP Address (GIADDR)			
Client Hardware Address (CHADDR)—16 bytes			
Server Name (SNAME)—64 bytes			
Filename—128 bytes			
DHCP Options			

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

17

Summary

Cisco.com


- DHCP
- Questions?

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

18

DOMAIN NAME SERVICE



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

19

Agenda

Cisco.com

- Introduction to Names and Addresses
- Managing Addresses with DHCP
 - Protocol
 - Assignment and Reliability
- **Resolving Names with DNS**
 - **Protocol**
 - Database
 - Reliable Operation
- New Things

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

20

Domain Name Service

Cisco.com

- **DNS is a database**
And the protocol to access it
- **Distinctive features:**
 - Design for lookup queries
 - Replicated content
 - Distributed control (zones)

NMS-1101
9592_04_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

21

Name Hierarchy

Cisco.com

- **Independent of address hierarchy**
- **Names length not limited by address size (63 bytes/label, 255 bytes/FQDN)**




NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

22

DNS Servers and Resolvers

Cisco.com



- Application connects by name, the application gets the address from the resolver
- Most applications use addresses in the order provided by the resolver

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

23

TCP and UDP Ports

Cisco.com

- Port 53 for both TCP and UDP
- UDP for queries if small enough
- TCP for zone transfer
- Server can use source port of 53 when “forwarding”

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

24



Redirection and Recursion

Cisco.com

- **Redirection:**
“Take your question down the hall”
- **Recursion:**
“I’ll get back to you”
- **Resolver sets Recursion Desired (RD), server responds with Recursion Available (RA) through bits in the DNS header**

NMS-1101
9592_04_2004_c2

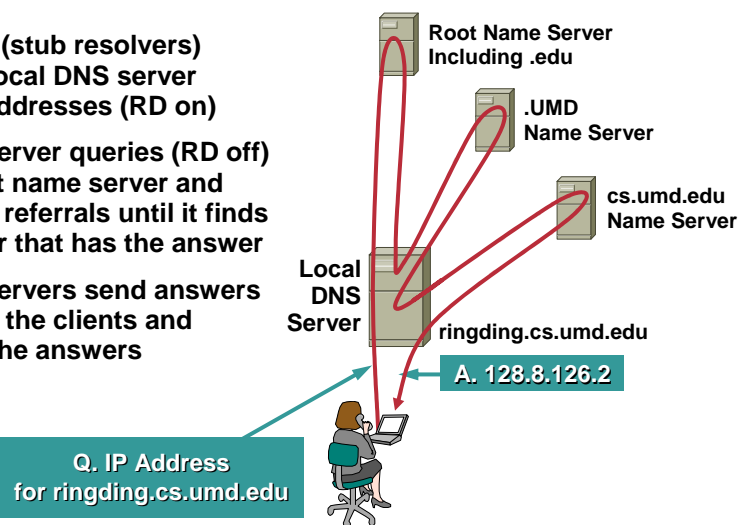
© 2004 Cisco Systems, Inc. All rights reserved.

25

DNS First Query

Cisco.com

- **Clients (stub resolvers) query local DNS server for IP addresses (RD on)**
- **Local server queries (RD off) the root name server and follows referrals until it finds a server that has the answer**
- **Local servers send answers back to the clients and cache the answers**



NMS-1101
9592_04_2004_c2

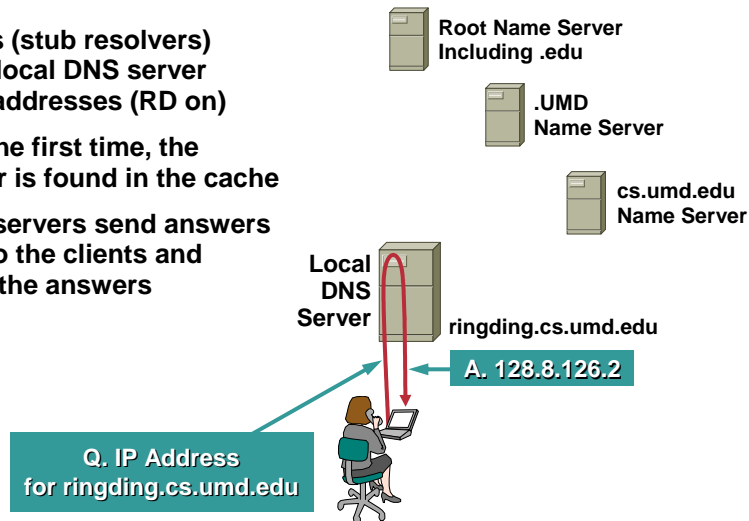
© 2004 Cisco Systems, Inc. All rights reserved.

26

DNS Subsequent Queries

Cisco.com

- Clients (stub resolvers) query local DNS server for IP addresses (RD on)
- After the first time, the answer is found in the cache
- Local servers send answers back to the clients and cache the answers



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

27

Caching and Forwarders

Cisco.com

- Caching is controlled by the Time to Live
- Negative caching (saving information that record doesn't exist) is required by RFC 2308
- The "minimum" TTL parameter in the SOA (or the TTL of the SOA RR itself if it is lower) determines the TTL for caching negative answers
- Sending a recursive query to a forwarder builds a cache for the site

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

28

Time to Live

Cisco.com

- **Changing host addresses**
 - Reduce TTL prior to change
 - Then restore to manage the load
- **CNR dynamically updates DNS TTL with 1/3 DHCP lease time**



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

29

Agenda

Cisco.com

- **Introduction to Names and Addresses**
- **Managing Addresses with DHCP**
 - Protocol
 - Assignment and Reliability
- **Resolving Names with DNS**
 - Protocol
 - **Database**
 - Reliable Operation
- **New Things**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

30

Terminology

Cisco.com

- Label (name, owner)
- Resource record (type)
- Value (encoded by type)

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.



Record Format

Cisco.com

Label		RR-Type	Value
<code><name></code>	<code>[<t1>]</code>	<code>[<class>]</code>	<code><type> <data></code>
VAXA.ISI.EDU.	IN	A	10.2.0.27
VAXA.ISI.EDU.	IN	A	128.9.0.33

Optional Fields:
We Only Care about Class = IN (Internet)
TTL ~ Time to Live in a Cache

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

32

Address Examples

Cisco.com

<name>	<type>	<data>
VAXA.ISI.EDU.	A	10.2.0.27
	A	128.9.0.33

- In the standard format for a zone description, an empty label is the same as in the previous line

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

33

IP Version 6

Cisco.com

- AAAA resource records defined in RFC 3152

v6host.example.com. AAAA 4321:0:1:2:3:4:567:89ab

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

34

Address and Canonical Name

Cisco.com

- **A (address) resource record (RR)**
The value is a 32-bit IPv4 address
- **CNAME**
The value is the name of a label
The value of a canonical name is not allowed to be the label of a CNAME record—but multiple levels of reference happen anyway

NMS-1101
9592_04_2004_c2

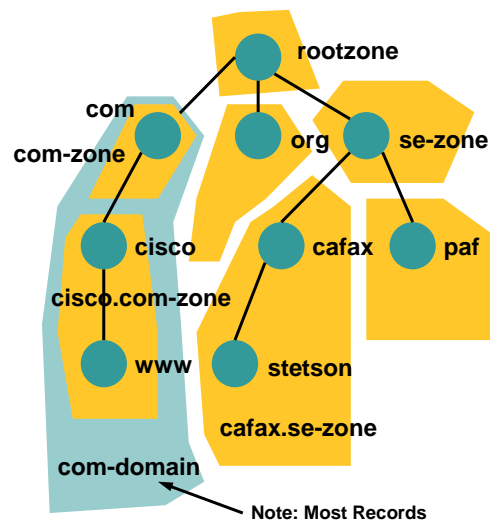
© 2004 Cisco Systems, Inc. All rights reserved.

35

Delegation Zone

Cisco.com

- Hierarchical name space
- Each node in the tree represents a domain/subdomain
- Some subdomains are defined as zones
- Each zone has a “primary” name server responsible for all lower nodes, but delegation is to all authoritative name servers
- Resource Records (RR) can, but don't have to, be defined for each node



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

36

Delegation Records

Cisco.com

- **Distributes database administration**
- **Name Server (NS) RR**
 - Refer in parent **AND** child zone to authoritative name servers for child (delegated) zone
- **Zone Start of Authority (SOA) RR**
 - Contain administrative information for delegated zone, in delegated zone only

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

37

Delegation: NS and “Glue”

Cisco.com

- **NS Resource Record (RR)**
- **“Glue” entries in parent zone when name server is in delegated zone**

<domain>	NS	<server>
SRI.COM.	NS	KL.SRI.COM.
KL.SRI.COM.	A	10.1.0.2

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

38

Delegation: SOA

Cisco.com

```
<name> [<t1>] [<class>] SOA <origin> <person> (  
  <serial>  
  <refresh>  
  <retry>  
  <expire>  
  <minimum> )
```

\$ORIGIN ARPA.

```
@ IN SOA SRI-NIC.ARPA. HOSTMASTER.SRI-NIC.ARPA. (  
  45 ;serial (sequential)  
  3600 ;refresh (1 hour regular check)  
  600 ;retry (10 minutes between check)  
  3600000 ;expire (42 days until refresh)  
  86400 ) ;minimum [negative] (a day)
```

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

39

Reverse DNS for IPv4 Addresses

Cisco.com

- Another hierarchy for in-addr.arpa.

Reverse the order in the label
because names aggregate
within suffixes rather than
(address) within prefixes



```
27.0.2.10.IN-ADDR.ARPA. PTR VAXA.ISI.EDU.  
33.0.9.128.IN-ADDR.ARPA. PTR VAXA.ISI.EDU.
```

ARPA: “Addressing and Routing Parameters Area”

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

40

Reverse DNS for IPv6 Addresses

Cisco.com

Reverse DNS for IPv6 in IP6.ARPA

```
v6host.example.com. AAAA 4321:0:1:2:3:4:567:89ab
```

```
b.a.9.8.7.6.5.0.4.0.0.0.3.0.0.0.2  
.0.0.0.1.0.0.0.0.0.0.0.1.2.3.4.IP6.ARPA.  
PTR v6host.example.com.
```

**“FOR EVERY IP ADDRESS, THERE SHOULD BE A MATCHING
PTR RECORD THE in-addr.arpa. DOMAIN.”**

RFC 1912, COMMON DNS OPERATIONAL AND CONFIGURATION ERRORS

Reverse Complication

Cisco.com

- **When the address field separations are not on dotted-decimal boundaries (classless addressing)**
- **Create CNAMEs in in-addr.arpa.**
 - Parent creates labels:
 - 0/25.1.2.10.in-addr.arpa. IN NS ns.example.com.
 - 13.1.2.10 IN CNAME 13.0/25.1.2.10.in-addr.arpa.
 - Child create PTR:
 - 13.0/25.1.2.10.in-addr.arpa. IN PTR foo.example.com.
- **Then delegate the sub-domain with the (classless) new label**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

43

Records for Applications

Cisco.com

- **MX**
- **SRV**
- **NAPTR**



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

44

MX

Cisco.com

- Mail eXchange RR
- Where the mail for the host is to be sent
- Round robin within equal preferences
- Mailers send only to lower preference numbers
po2 is only allowed to send to po1 in example below

Name	TTL	Class	MX	Preference	Target
BAZ.FOO.COM.		MX	10	PO1.FOO.COM.	
		MX	20	PO2.FOO.COM.	
		MX	20	PO3.FOO.COM.	

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

45

Wildcards

Cisco.com

- Special treatment for '*' in the label
- Any name in the query matches, and the answer is synthesized
- Most often used in mail exchange
- Create complications with DNSSEC

FOO.COM.	MX	10	RELAY.CS.NET.
*.FOO.COM.	MX	20	RELAY.CS.NET.

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

46

SRV

Cisco.com

- Generalize the MX idea
- Find hosts offering service in a domain
- Add structure to the name
- Add fields to the RR—specialize priority and weight (replace preference)
- Target must **NOT** be an alias (CNAME)
- RFC 2782

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

47

SRV

Cisco.com

- Format of RR and Example

```
_Service._Proto.Name.  SRV Priority Weight Port Target
_ldap._tcp.example.com. SRV  1  10  389  ldap1.example.com.
                        SRV  1  20  389  ldap2.example.com.
```

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

48

NAPTR

Cisco.com

- Naming Authority PointeR
- Universal resource identifier
- Regular expressions
- Replacement strings
- RFC 2915
- Used for example in ENUM

ENUM is mapping from E.164
number to URLs



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

49

Agenda

Cisco.com

- Introduction to Names and Addresses
- Managing Addresses with DHCP
 - Protocol
 - Assignment and Reliability
- Resolving Names with DNS
 - Protocol
 - Database
 - Reliable Operation
- New Things

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

50

Secondary Servers

Cisco.com

- Reliability depends on separation
- Location—physical and subnet
- Independent fate—separate power
- Separate administration if possible
- RFC 2182—best current practice

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

51

Replication

Cisco.com

- Transfer zone contents (AXFR)
- Transfer controlled by serial number and refresh parameters in SOA

Secondary query for SOA and compare serial number with what it already has; if serial is higher, client will request a zone transfer

This is repeated by client as specified in Refresh in SOA

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

52



Replication Efficiency

Cisco.com

- **Notify (new protocol operation) enables primary to inform secondary when zone has changed [RFC 1996]**

In reality, it informs client to set refresh timer to zero, so client will restart new refresh cycle immediately

- **Incremental transfer (IXFR) sends just changes to the zone [RFC 1995]**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.



Agenda

Cisco.com

- **Introduction to Names and Addresses**
 - **Managing Addresses with DHCP**
 - Protocol
 - Assignment and Reliability
 - **Resolving Names with DNS**
 - Protocol
 - Database
 - Reliable Operation
- **New Things**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

54

Dynamic Update

Cisco.com

- Atomic update of RR-set
- Base specification—RFC 2136
- Secure version—RFC 3007
- Created so that DHCP servers and clients can update DNS

<http://ops.ietf.org/dns/dynupd/secure-ddns-howto.html>

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

55

Securing Queries


Cisco.com

- TSIG
Transaction Signature—RFC 2845
- Secret-key hash of the transaction (HMAC-MD5) to the forwarder
- Pseudo RR, not cached or saved
- Only useful with local forwarders

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

56



Securing Zone Transfer

Cisco.com

- **TSIG is used where**
- **Secondary servers have an administrative relationship that can support secret keys**
- **Don't need the overhead of public keys**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.



TKEY

Cisco.com

- **Transaction KEY—not stored**
- **Use DNS to establish secret keys alternative to manual keys**
- **Modes include:**
 - Diffie-Hellman**
 - GSS-API**
 - Server or resolver assigned encrypted (encrypted using KEY RR)**
- **RFC 2930**



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

58

SIG(0)

Cisco.com

- Use DNS for client to authenticate to server
- Authenticates the transaction
- Public KEY in DNS
- Private key in client
- RFC 2931

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

59

Securing Zone Contents

Cisco.com

- DNS security
- ~~Key~~ **DNSKEY** RR—distributes public keys for records
- ~~RRSIG~~ **RRSIG** RR—authenticates (signs) one RR set
- ~~NSEC~~ **NSEC** RR—“next” record enables authentication of non existence
- DS RR—delegation signer~key in parent with which the child zone is signed
- RFC 2535—being revised

draft-ietf-dnsext-dnssec-records



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

60

Deployment of DNS Security

Cisco.com

- **Experimental only now**
- **Trust depends on the entire path from the resolver**
- **Signing all the RR sets in large zones, like .com, is an unresolved problem, which Moore's law and deployment delay may fix**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

61

Split DNS

Cisco.com

- **External holds limited contents for public**
- **Internal**
 - Isolated clients query DNS servers configured as root**
 - Internal (secondary) servers forward to external caching server for other domains**


NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

62

Internal Root

Cisco.com




NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

63

Internal Forwarding Server

Cisco.com



With Recursion OFF

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

64

Load Sharing

Cisco.com

- Resolvers use addresses in the order received, although the original concept was that they choose randomly
- DNS server can rotate the order of the (multiple) addresses of a hostname to distribute the load

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

65

Source-Dependent Answers

Cisco.com

- Return addresses in the order of “closeness” to the resolver
- Same subnet is close, but requires knowing the subnet mask
- Can look into the routing structure
- DNS support for content networking uses other metrics for which answer to give

NMS-1101
9592_04_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

66

Distributed Director

Cisco.com

- DD server is authoritative nameserver for zone with name as webserver
- When query arrives, DD server verify with DRP agents which webserver is closest to client
- DD server respond with "best" IP-address
- DD server can also do HTTP redirects in HTTP mode



NMS-1101
9592_04_2004_c2

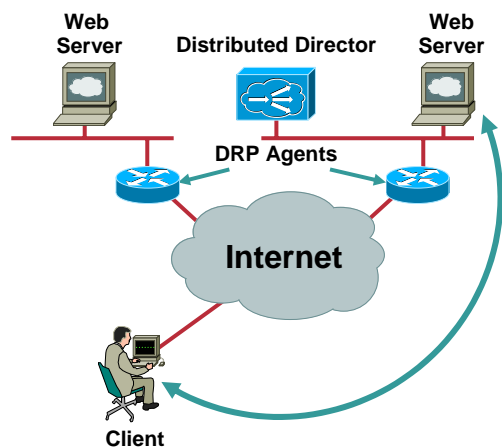
© 2004 Cisco Systems, Inc. All rights reserved.

67

DRP—Director Response Protocol

Cisco.com

- Operates with routers in the field to determine:
 - Client to server network proximity
 - Client to server link latency (RTT)
- UDP based



NMS-1101
9592_04_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

68

Anycast

Cisco.com

- **Announce the same IP address from multiple servers**
- **Needed for overloaded root servers (number of root servers limited by protocol)**
- **Only works with UDP**
- **Two versions:**
 - Anycast inside an AS (ok)
 - Anycast where the same address is announced from multiple source ASs (dangerous due to security reasons)



NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

69

Digression: Anycast on Root nameservers

Cisco.com

- **Anycast is used for root nameservers**
 - 7 of 13 root server operators use anycast today
 - Limit of 13 that fit in the reply packet (EDNS(0) expands)
- **Heavy load from bad queries...**
 - 8–10% from addresses not reachable (10/8 etc)
 - 3.6% queries about 10.in-addr.arpa etc
 - 4% queries about localhost
 - 3% about the TLD arpa
 - 2.3% recursive queries (RD bit on)
- **...and to limit impact from ddos attacks**

NMS-1101
9592_04_2004_c2

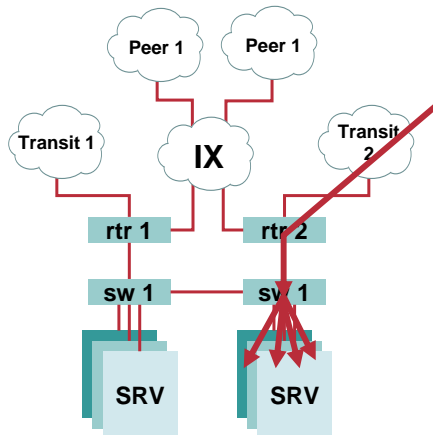
© 2004 Cisco Systems, Inc. All rights reserved.

70

Load Sharing Locally

Cisco.com

- Multiple hosts behind a content switch to make a “big” root nameserver



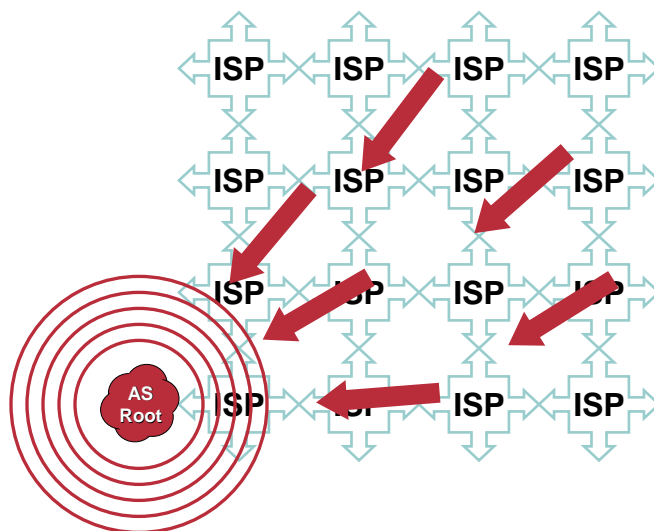
NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

71

DDOS Attack Scenario

Cisco.com




NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

72

With Anycast

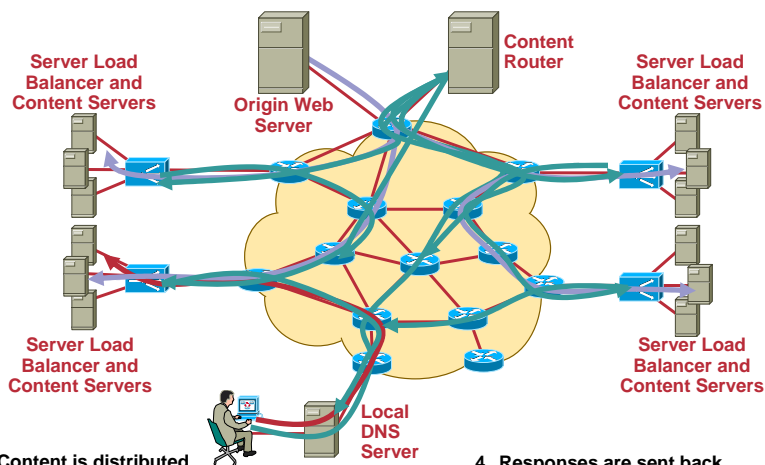


NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

73

Boomerang



1. Content is distributed
2. Client query content router via local DNS Server
3. Content servers are told to respond to DNS query
4. Responses are sent back
5. First response arriving is closest
6. Client connect to that IP address

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

74

Summary

Cisco.com

- **Addresses can be allocated automatically**
- **DNS can support more than just name to address lookup**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

75

Associated Sessions

Cisco.com

- **NMS-2101—DNS Deployment and Operation**

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.


76

Recommended Reading

Cisco.com

IP Addressing Fundamentals

ISBN: 1587050676



Available Onsite at the Cisco Company Store

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

77

Complete Your Online Session Evaluation!

Cisco.com

- WHAT:** Complete an online session evaluation and your name will be entered into a daily drawing
- WHY:** Win fabulous prizes! Give us your feedback!
- WHERE:** Go to the Internet stations located throughout the Convention Center
- HOW:** Winners will be posted on the onsite Networkers Website; four winners per day

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

78

CISCO SYSTEMS



Reverse Delegation

Cisco.com

- **RFC 2317**
2.3/12 CNAME xxx.IN-ADDR.ARPA

For delegation of

192.0.2.0/25 to organization A

192.0.2.128/26 to organization B

192.0.2.192/26 to organization C

Reverse Delegation Problem

Cisco.com

```
$ORIGIN 2.0.192.in-addr.arpa.  
;  
1 PTR host1.A.domain.  
2 PTR host2.A.domain.  
3 PTR host3.A.domain.  
;  
129 PTR host1.B.domain.  
130 PTR host2.B.domain.  
131 PTR host3.B.domain.  
;  
193 PTR host1.C.domain.  
194 PTR host2.C.domain.  
195 PTR host3.C.domain.
```

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

81

Reverse Delegation Solution

Cisco.com

```
$ORIGIN 2.0.192.in-addr.arpa.  
@ IN SOA my-ns.my.domain.  
hostmaster.my.domain. (...)  
;...  
; <<0-127>> /25  
0/25 NS ns.A.domain.  
0/25 NS some.other.name.server.  
;  
1 CNAME 1.0/25.2.0.192.in-addr.arpa.  
2 CNAME 2.0/25.2.0.192.in-addr.arpa.  
3 CNAME 3.0/25.2.0.192.in-addr.arpa.  
;  
; <<128-191>> /26  
128/26 NS ns.B.domain.  
128/26 NS some.other.name.server.too.  
;  
129 CNAME 129.128/26.2.0.192.in-addr.arpa.  
130 CNAME 130.128/26.2.0.192.in-addr.arpa.  
131 CNAME 131.128/26.2.0.192.in-addr.arpa.
```

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

82

Reverse Delegation Solution (Cont.)

Cisco.com

```
; <<192-255>> /26
192/26      NS      ns.C.domain.
192/26      NS      some.other.third.name.server.
;
193         CNAME   193.192/26.2.0.192.in-addr.arpa.
194         CNAME   194.192/26.2.0.192.in-addr.arpa.
195         CNAME   195.192/26.2.0.192.in-addr.arpa.

$ORIGIN 0/25.2.0.192.in-addr.arpa.
@          IN      SOA   ns.A.domain. hostmaster.A.domain. (...)
@          IN      NS    ns.A.domain.
@          IN      NS    some.other.name.server.
;
1          PTR     host1.A.domain.
2          PTR     host2.A.domain.
3          PTR     host3.A.domain.
```

NMS-1101
9592_04_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

83