DNSSEC at Penn

Shumon Huque
University of Pennsylvania
ESCC/Internet2 Joint Techs Conference
July 20th 2009

DNSSEC at a glance

- "DNS Security Extensions"
- A system to verify the authenticity of DNS "data" using public key signatures
  - Protocol specs: RFC 4033, 4034, 4035, 5155
- Helps detect spoofing, misdirection, cache poisoning, etc
- Some potential secondary benefits:
  - Storing cryptographic keying material in the DNS: SSHFP, IPSECKEY, CERT, DKIM etc ..

. (root) root's pubkey

www.upenn.edu
set DO bit

2

3
referral to .edu
+ DS, RRSIG

.edu
edu pubkey

recursive
resolver
(has root's pubkey)

4

5
referral to upenn.edu
+ DS, RRSIG

www.upenn.edu

1

8

6

answer 1.2.3.4
+ RRSIG

7

upenn.edu
upenn pubkey

endstation
(uses DNS stub resolver)

DNSSEC Records

DNSKEY	Contains zone public key
RRSIG	Contains DNSSEC signature
NSEC	Points to next name in zone (used for authenticated denial of existence)
DS	Delegation Signer (certifies public key for subordinate zone)
NSEC3	Newer version of NSEC (provides zone enumeration protection and opt-out)
NSEC3PARAM	NSEC3 parameters

Signed zone additions

- When signed, each zone will have:
  - 1 or more DNSKEYs at the apex
  - 1 NSEC for every DNS name
    - 1 RRSIG for every RR set (Resource Record Set)
    - 1 or more DS records for every (secured) delegation
- Exceptions:
  - Non-authoritative data like delegation NS records and glue have no signatures

Multiple DNSKEYs

- Typically: 2-level key hierarchy
- KSK: Key Signing Key
  - Signs other keys (can be stronger and kept offline; used as the trust anchor and certified by parent zone)
- ZSK: Zone Signing Key
  - Signs all data in the zone (can be lower strength and impose less compute overhead; can rollover without external impact)

A few example queries on our testbed using the **dig** tool (available on most UNIX/Linux platforms) ...

```
$ dig jabber.upenn.edu AAAA

;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 337

;; QUESTION SECTION:
;jabber.upenn.edu.              IN      AAAA

;; ANSWER SECTION:
jabber.upenn.edu.       86400   IN      AAAA  2001:468:1802:101::805b:2ac

;; AUTHORITY SECTION:
upenn.edu.              86400   IN      NS      dns2.udel.edu.
upenn.edu.              86400   IN      NS      noc2.dccs.upenn.edu.
upenn.edu.              86400   IN      NS      noc3.dccs.upenn.edu.
upenn.edu.              86400   IN      NS      dns1.udel.edu.

;; ADDITIONAL SECTION:
noc2.dccs.upenn.edu.    86400   IN      A       128.91.254.1
noc2.dccs.upenn.edu.    86400   IN      AAAA    2001:468:1802:102::805b:fe01
noc3.dccs.upenn.edu.    86400   IN      A       128.91.251.158
dns1.udel.edu.          86400   IN      A       128.175.13.16
dns2.udel.edu.          86400   IN      A       128.175.13.17
```

$ dig jabber.upenn.edu AAAA +dnssec

Authenticated Data

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 690
;; flags: qr aa rd ad; QUERY: 1, ANSWER: 2, AUTHORITY: 5, ADDITIONAL: 7

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: do; udp: 4096
;; QUESTION SECTION:
;jabber.upenn.edu.      IN AAAA
```

DNSSEC Ok

Answer & Signature

```
;; ANSWER SECTION:
jabber.upenn.edu.       86400 IN AAAA 2001:468:1802:101::805b:2ac
jabber.upenn.edu.       86400 IN RRSIG AAAA 5 3 86400 20090719232941 (
                                20090619232159 23382 upenn.edu.
                                26bOACMMoojfX/zVW1AfhWZ/LSuvn5Fo8iHxVqV/NBzT
                                JJb0LitaOQVqKCxxswH0TDQgmQiayaL6xGk0yfHo7T32
                                i1pEFbJdkbNvd4M7GQktB22lBY12Uzrd+/FmAA2xqJ2P
                                ZDBNbIjkd41oRD098BAmYfGGGDdb8Dyectx8L/Q= )

;; AUTHORITY SECTION:
upenn.edu.              86400 IN NS dns1.udel.edu.
upenn.edu.              86400 IN NS noc3.dccs.upenn.edu.
upenn.edu.              86400 IN NS dns2.udel.edu.
upenn.edu.              86400 IN NS noc2.dccs.upenn.edu.
upenn.edu.              86400 IN RRSIG NS 5 2 86400 20090719232217 (
                                20090619223616 23382 upenn.edu.
                                WWpT4uD9p5zORM+2O7pRZ46+Qo3cHj9tnjxH62Xt9QBR
                                yu9V7+3ihlIM1HCd9kjsddskT8GJ+5hEzykB8fPIjSli
                                bqG6hCnCccGdTsGzmPoGdlz95H7Nf2yfrlGLAcSCix6I
                                EJb8Aj4+OW9Zq1dmeZrnJDXSzm8joQg5+IlkzR4= )
```

RRSIG

Record Type Covered

Algorithm

#Labels

Original TTL

Signature Expiration (YYYYMMDDHHmmSS)

```
;; ANSWER SECTION:

jabber.upenn.edu.       86400 IN AAAA 2001:468:1802:101::805b:2ac

jabber.upenn.edu.       86400 IN RRSIG AAAA 5 3 86400 20090719232941 (
                                20090619232159 23382 upenn.edu.
                                26bOACMMoojfX/zVW1AfhWZ/LSuvn5Fo8iHxVqV/NBzT
                                JJbOLitaOQVqKCxxswHOTDQgmQiayaL6xGkOyfHo7T32
                                i1pEFbJdkbNvd4M7GQktB22lBY12Uzrd+/FmAA2xqJ2P
                                ZDBNbIjkd41oRD098BAmYfGGGDdb8Dyectx8L/Q= )
```

Signature Inception (YYYYMMDDHHmmSS)

Key Tag

Cryptographic Signature

flags

proto

algorithm

```
$ dig upenn.edu DNSKEY

;; ANSWER SECTION:
upenn.edu.              7200 IN DNSKEY 256 3 5 (
                                AwEAAcDt107stSjvoBA/YVPr+2gvB3v33tXr7ROZ/Jqm
                                WtNLraxQPzgXM1AhwjtdEqwCAnk01V7+Fw7K94sh6jpI
                                5bFofS7MGtd0VvNyq52bgRnusgbm1ME2Lx9+o3fy9ppv
                                7C6bahGrV3aiq9wNVPj/ccJn5AnZCOsi3grVsj6izCYH
                                ) ; key id = 46752
upenn.edu.              7200 IN DNSKEY 256 3 5 (
                                AwEAAfAHsS33kJEImVk09yFJY5hXumAo+JVVJMJpJUaj
                                l/rh0fFkdikS2oatVvxHHHqKN9Kg3DoKQss/CzCZa4zn
                                KlqYGvSl7RefKR3QLyPBGQN2aOUWxshDgOWLmOtqNpmP
                                +6Drfn8LJVTOjuwmU80laQcdA/AoOGVPE3zP16G/F+qp
                                ) ; key id = 43248
upenn.edu.              7200 IN DNSKEY 257 3 5 (
                                AwEAAek95gyBF2nurdIE2Q63VVcMlazOlQEnz0N4Ce89
                                SB4Juw2eEBerLmEanuGJbrs0oGx3SKCMyhOYL9q1ZrmC
                                NCf6PnACwv88NtrYOjHAOmOlLAvKAQv8MTBbEwTWBBw5
                                K8jUwzcaGyDjo3U+Hai+ow8Tiev0By+hrcT4DegsbEB8
                                MEQIgEUO/Kw9wbJLEdpvVXtuV2l78G75FUwmrA8jzEka
                                M7bKg/HSTIMupbwfs4IHYgbG/PkqOZYL3uxm9gncVjbh
                                4YYd4OG6koVoWteWTS8JdYq4gr9b9AEjhwAzbe7bd7pX
                                +qD70CCbh0jSOVhPvhRpCHIYZAJIwEAWs711HHM=
                                ) ; key id = 29242
```

encoded public key

Secure Delegations

- Indicated by DS (Delegation Signer) record
- Appears in the delegating zone
- Contains a hash of the public key of the delegated zone (and also a corresponding RRSIG)

DS contains hash of the public key of delegated domain. 2 DS records are shown here because 2 different hashing algorithms were used

(hypothetical example since .net isn't signed yet ..)

```
magpi.net.        3587 IN   DS 15462 5 2 (
                            9EFD69115037892179A5408F04E6EA93CBA2488B221
                            96493142E47D1AD24C3A )
magpi.net.        3587 IN   DS 15462 5 1 (
                            C020FB9E09EE30568F250E2086D52E62F2B4FA17 )
magpi.net.        3587 IN   RRSIG DS 5 5 3600 20090812170009 (
                            20090713170009 64263 dlv.isc.org.
                            M+09bX9XP79yfDhWDUNuDEg9KOEHV2eV33/dEYnutVpD
                            iZYGqJ6BWLhWZYE8Y8megYozfa5UJv/AVcdIZ51JCPI4
                            k/jlRDj60kRaWRlfCBgqOR2WPL+F20vhg3wS57bIjmRW
                            To0r/HpXemnJVdXLbrzWD5WdpYGFy1UVX+15N4o= )
```

Signature of DS record set

DNSSEC at Penn

- Strategy: deploy in a simpler DNS environment first to gain experience
- Started with MAGPI (an Internet2 GigaPoP we run)

DNSSEC at MAGPI

- Deployed in production since May 2006
- 17 zones: magpi.net, magpi.org, and 15 reverse DNS zones
- KSK 2048-bit RSASHA1, ZSK 1024-bit RSASHA1
- Rollover: ZSK pre-publish; KSK double signature
- Use RIPE key management tools
  - http://www.ripe.net/disi/dnssec_maint_tool/
- Keys present in ISC's DLV Registry

https://rosetta.upenn.edu/magpi/dnssec.html

DNSSEC at Penn

- Requirements significantly different from MAGPI
- Much larger DNS infrastructure & more data
- Dynamically updated, 24x7
  - by IT staff and by automated programs
- Can't freeze updates to sign/resign/rollover etc
- Don't want large zone reloads and transfers; want efficient incremental transfers (IXFR)

DNSSEC at Penn

- Centralized DNS operation & management
- No DNS delegations to subdivisions
- But distributed authority to edit/create data
- Home grown DNS management system

Home Grown DNS Management System

- Baggage: many hooks into non-DNS systems
- Custom Code and Protocol
  - XML-RPC and Kerberos + AuthZ system
- Interface to Name server:
  - Dynamic Update with static TSIG Key
  - This is where DNSSEC functionality is inserted

DNSSEC at Penn

- A DNSSEC testbed is up and running
- Production deployment anticipated this summer
- What we're using:
  - ISC BIND nameserver 9.6.1
  - Set of home grown tools for zone maintenance
- Co-operation with operators of offsite secondaries (Univ of Delaware in our case)

DNSSEC at Penn

- All forward and reverse zones to be signed
- 2048-bit RSASHA1 KSK
- 1024-bit RSASHA1 ZSK
- KSK rollover: double signature policy
- ZSK rollover: pre-publish policy
- (See RFC 4641 for key maintenance practices)

BIND 9.6 features we needed

- Dynamic Update with DNSSEC
  - transition zone from insecure to secure by insertion of DNSKEY records
    - key rollover via UPDATE
- Automatic resigning
- Improved dynamic update and automation features will appear in BIND 9.7

University of Pennsylvania DNSSEC Architecture

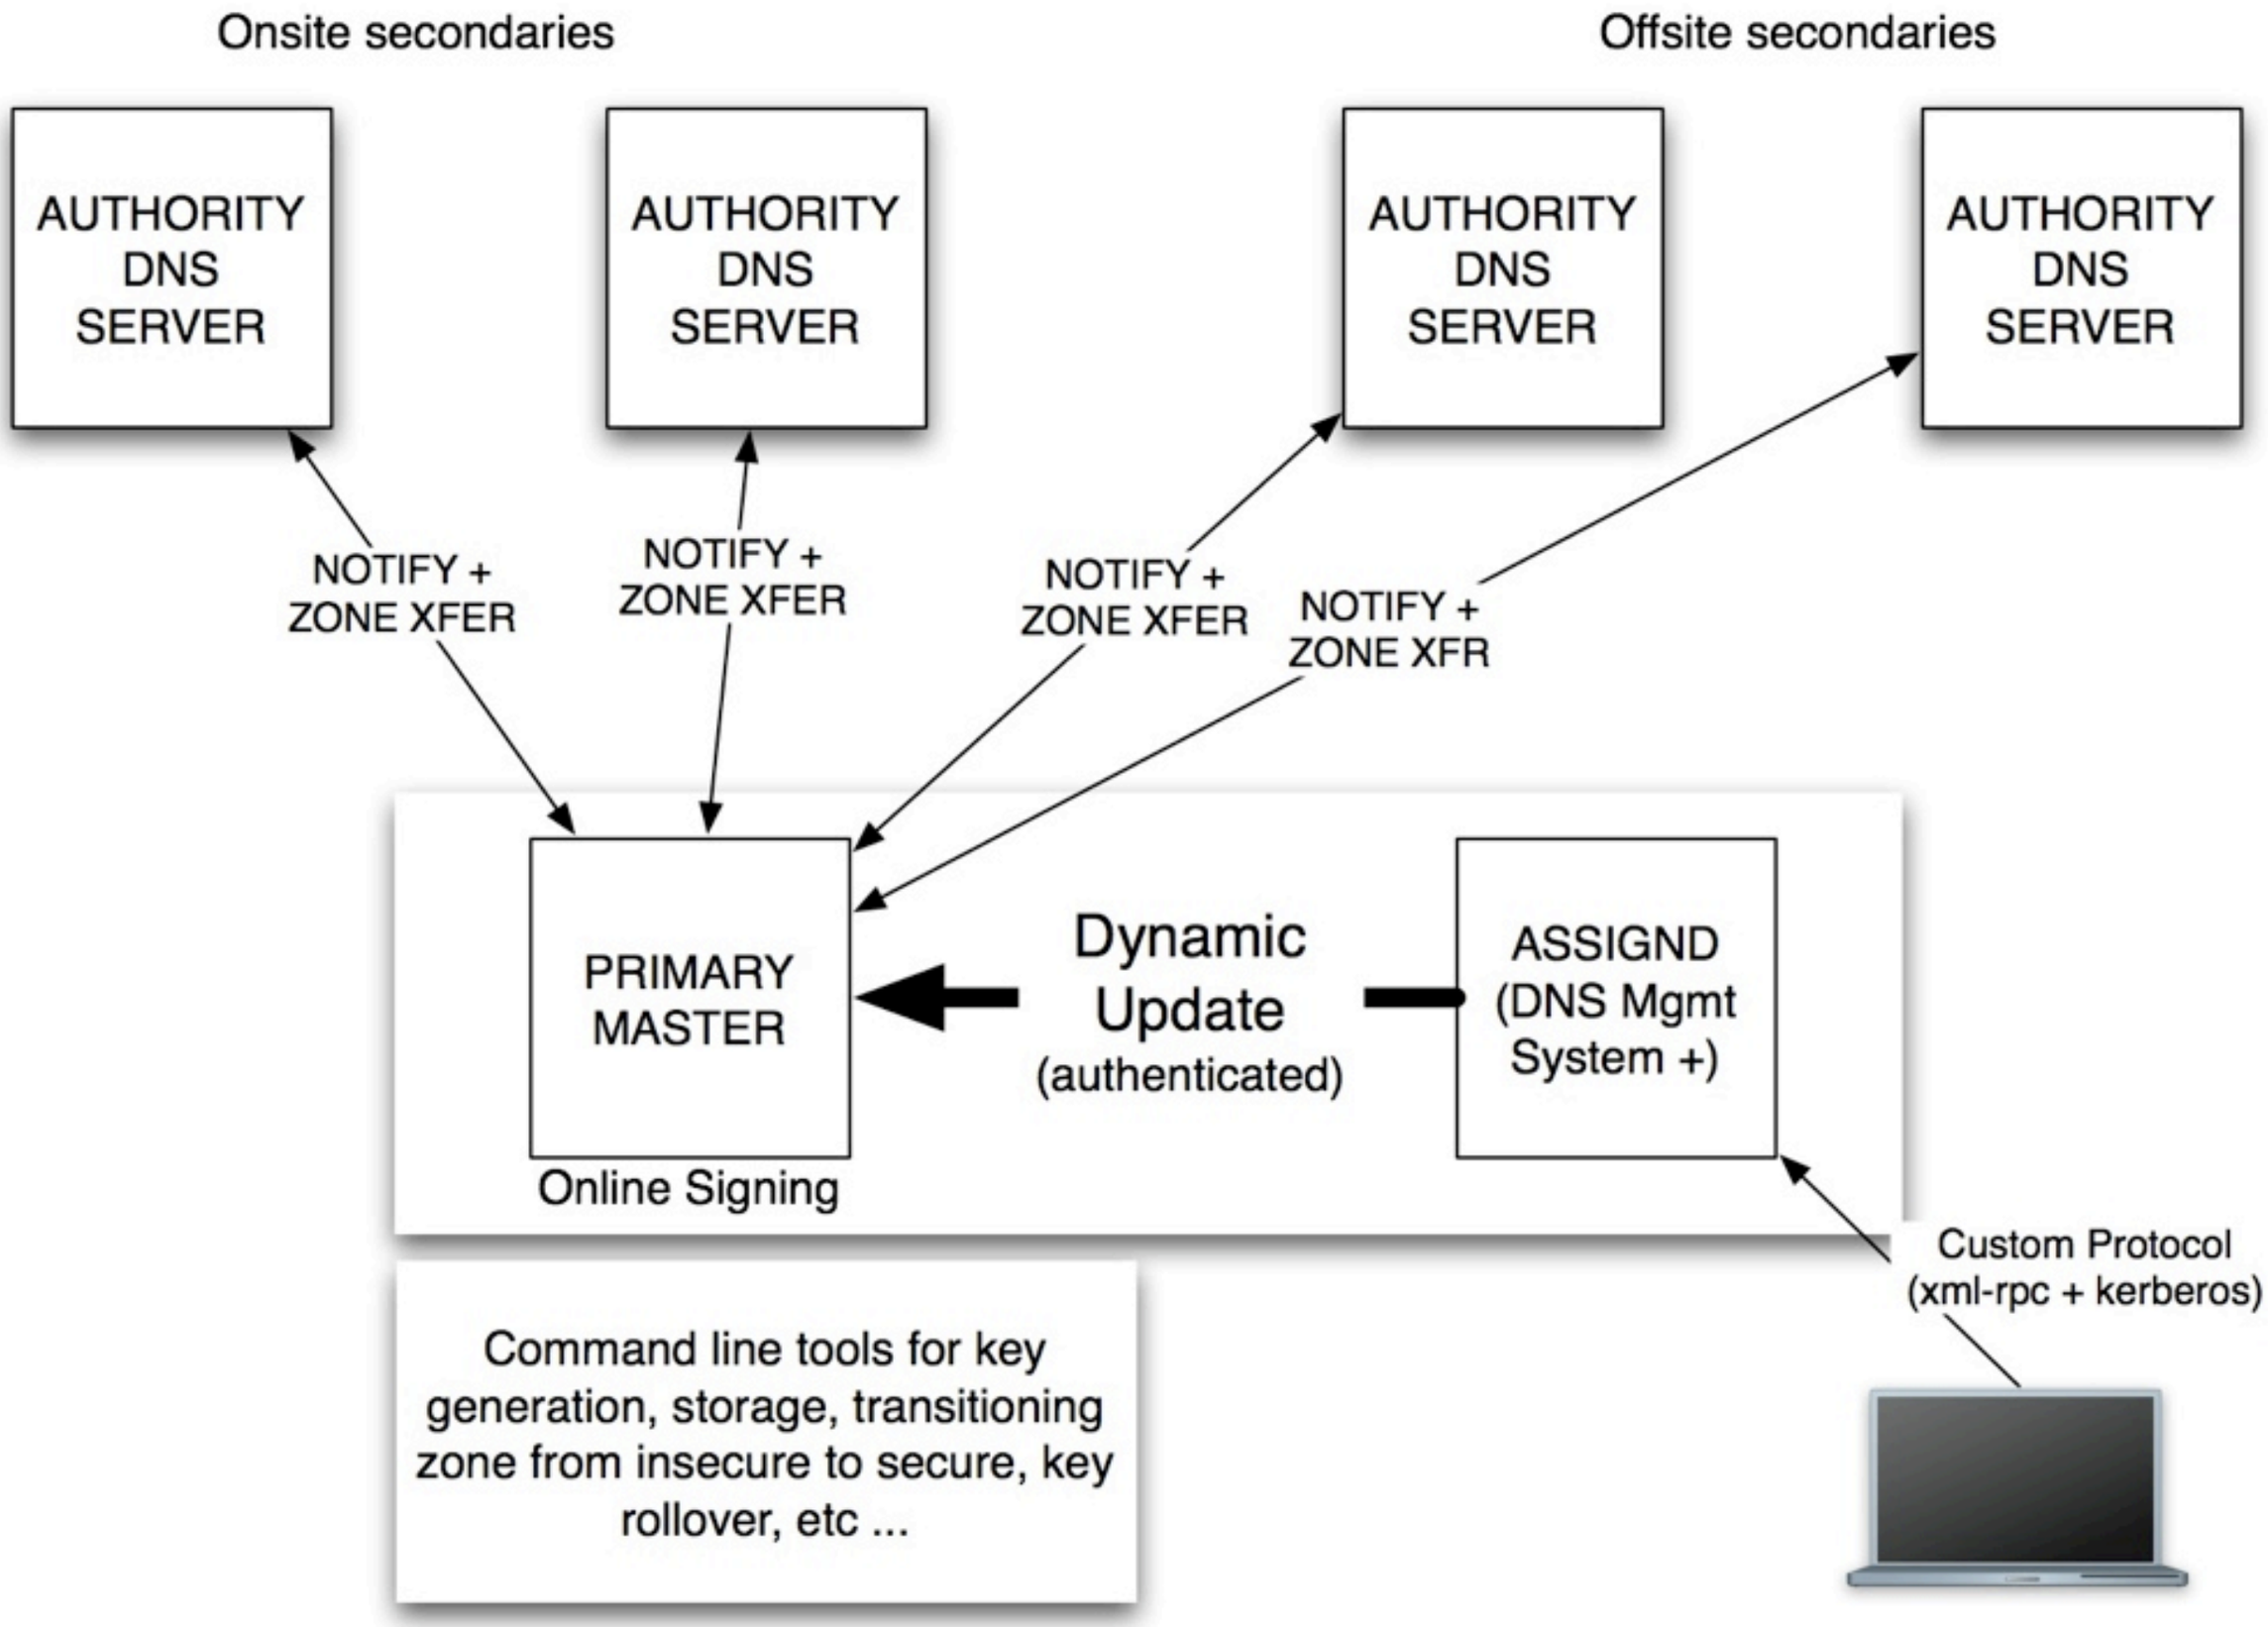

Our tools (6 python programs)

- securezone
- rollover-zsk-stage1
- rollover-zsk-stage2
- rollover-ksk-stage1
- rollover-ksk-stage2
- dnssec-keystat

Some data from our testbed deployment ...

3.9x 14.9x

Zone	Unsigned		Signed	
	#recs	#bytes	#recs	#bytes
upenn.edu	99,127	2,823,964	388,658	42,217,260
upenn.org	4	181	15	1,847
penn.edu	6	235	43	4,928
123.165.in-addr.arpa	32,348	896,828	129,332	14,418,217
130.158.in-addr.arpa	7,573	224,715	30,261	3,376,625
91.128.in-addr.arpa	26,344	811,317	105,264	11,732,214
91.130.in-addr.arpa	19,560	591,443	78,091	8,671,506
2.84.192.in-addr.arpa	132	3,311	521	57,545
0.7.4.f.7.0.6.2.ip6.arpa	7	313	21	2,039
2.0.8.1.8.6.4.0.1.0.0.2.ip6.arpa	65	2,629	258	36,748

* #bytes: number of bytes transferred by a full (AXFR) zone transfer

Record type counts in upenn.edu

RR Type	Count	% of Total
A	85,288	21.9%
AAAA	62	0.0%
CNAME	9,599	2.5%
DNSKEY	3	0.0%
MX	1,282	0.3%
NS	8	0.0%
NSEC	96,387	24.8%
RRSIG	193,137	49.7%
SOA	1	0.0%
SRV	2,866	0.7%
TXT	21	0.0%
TYPE65534	3	0.0%

Disk & Memory Consumption of nameserver process

	Unsigned	Signed
Virtual Memory	43 MB	133 MB (3x)
Resident Set Size	40 MB	129 MB (3x)
Zonefiles on disk	113 MB	233 MB (2x)

(BIND 9.6.1, authoritative only, text zone files)

Key Distribution Plans

- Secure delegations from Educause and ARIN eventually ...
- Submission to ISC DLV registry
  - By end of year, after period of testing
- HTTPS web page

DLV: DNSSEC Lookaside Validation

- A mechanism to securely locate DNSSEC trust anchors "off path"
- An early deployment aid until top-down deployment of DNSSEC is completed
- ISC's DLV Registry:
  - https://www.isc.org/solutions/dlv

DNSSEC Deployment in the Internet ...

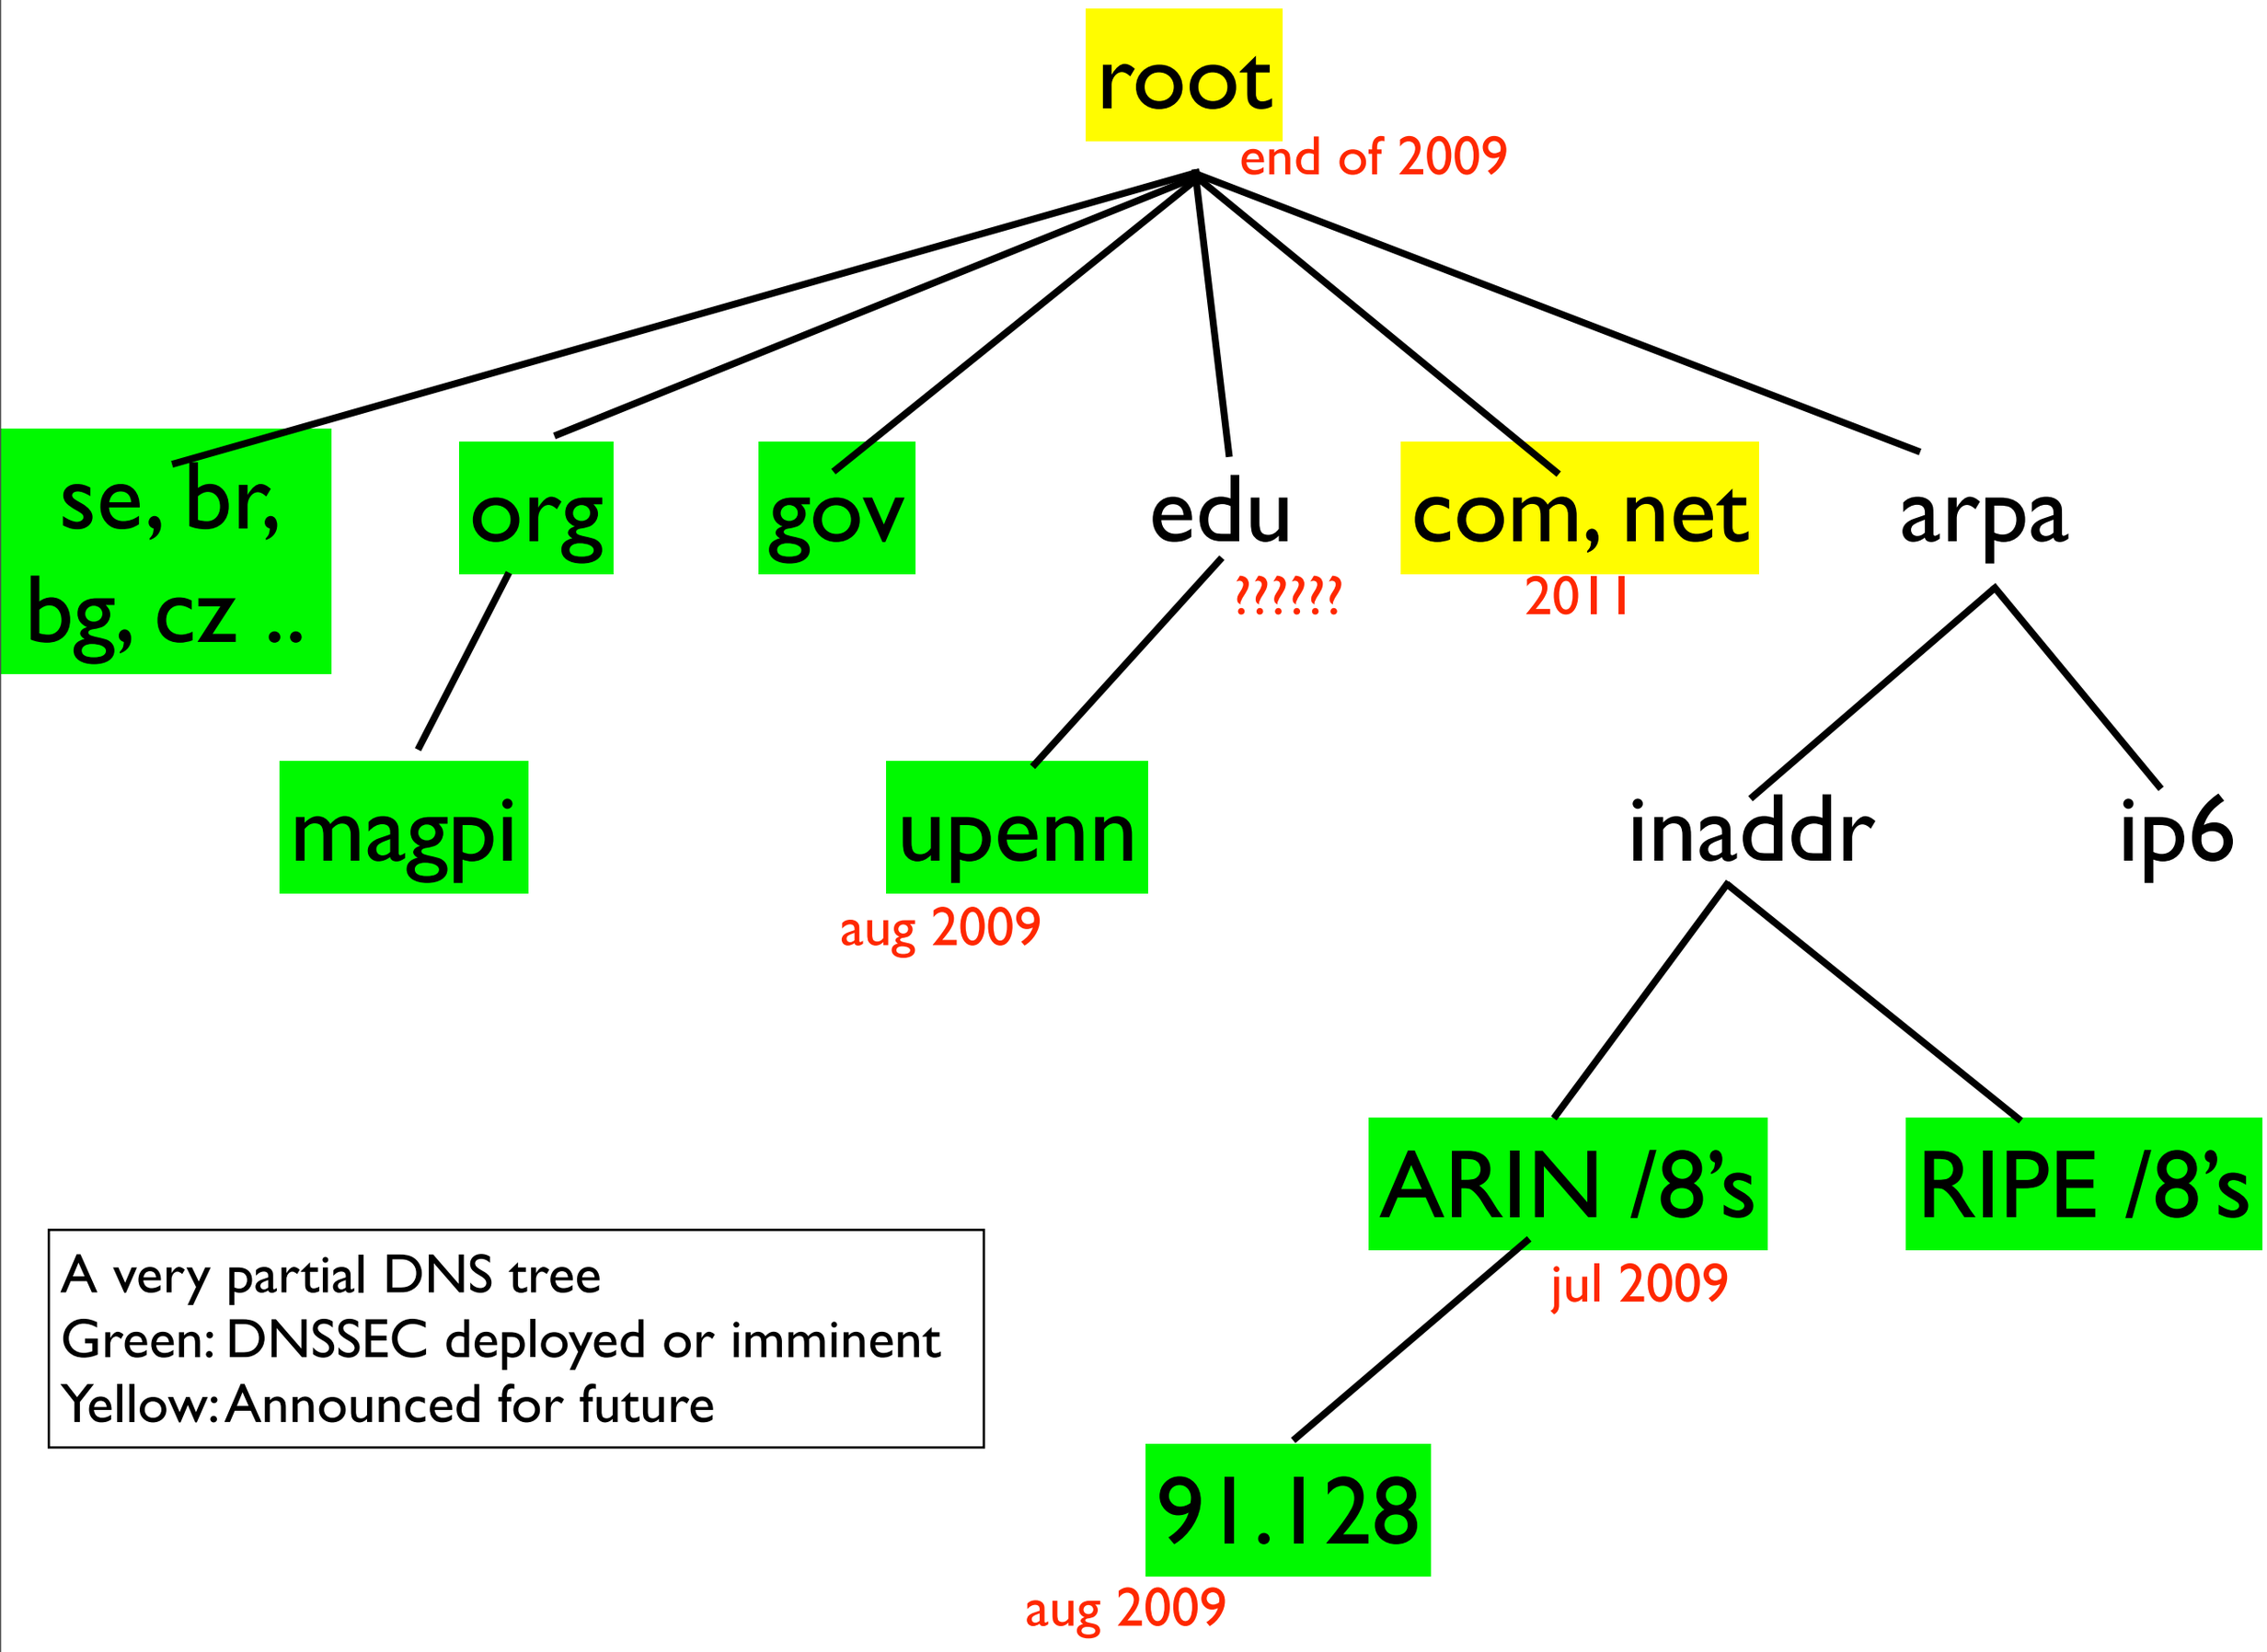

A very partial DNS tree
Green: DNSSEC deployed or imminent
Yellow: Announced for future

Notable DNSSEC Deployments to date

- Top Level Domains
  - gTLDs ORG, GOV
  - ccTLD: SE, BR, BG, CZ, PR, TH, (+ some IDNs)
- RIPE and ARIN Reverse DNS blocks
- *Note: Some don't offer secure delegation yet though! (ORG and ARIN)*

Root Signing: -- end of 2009?
http://www.icann.org/en/announcements/announcement-2-03jun09-en.htm
http://www.nist.gov/public_affairs/releases/dnssec_060309.html

.ORG -- done as of 2009-06-02; no support for secure delegation yet

.GOV -- done early 2009, http://dotgov.gov/dnssecinfo.aspx

ARIN reverse DNS: -- done as of 2009-07-01; no support for secure delegation yet
https://www.arin.net/about_us/dnssec/
https://www.arin.net/about_us/dnssec/trust_anchors.html

RIPE reverse DNS: http://www.ripe.net/reverse/dnssec/

.ARPA
http://www.iab.org/documents/correspondence/2009-06-02-Roseman-Signing-by-IANA-of-ARPA.html

.COM and .NET -- in 2011
http://www.networkworld.com/news/2009/022409-verisign-dns-security.html?hpg1=bn

SecSpider

- DNSSEC zone monitoring project
- http://secspider.cs.ucla.edu/
- Almost 12,000 signed zones as of mid July
  - (still a miniscule fraction)
- Crawling and user submissions
- Distributed polling

Who else in *our community* is doing DNSSEC?

DNSSEC Deployment in Authoritative Servers

(institution level production deployments, not subdivisions)

Org	Date	Type	Keys in DLV?
MAGPI (UPenn)	2006-06	NSEC	Yes
NANOG (Merit)	2006-08	NSEC	Yes
PSC/3ROX	2009-02/07	NSEC	Yes/No
UPenn	2009-07 (planned)	NSEC	End of Year

https://rosetta.upenn.edu/magpi/dnssec.html
http://www.merit.edu/networkresearch/dnssec.html

DNSSEC Validation in Campus Resolvers

Org	Date	Notes
Louisiana State U	2008-09	Uses ISC DLV
UC Berkeley	2008-10	Uses ISC DLV
Lawrence Berkeley Labs	????	Uses ISC DLV
U of Oregon	2009-02	Uses IANA ITAR anchors list
U of Delaware	????	Used ISC DLV (until .gov NSEC3 incident)

https://www.dnssec.uoregon.edu/

Internet2 DNSSEC Pilot Group

- List: dnssec@internet2.edu
- To join:
- https://mail.internet2.edu/wws/info/dnssec
- http://www.dnssec-deployment.org/internet2
- Monthly conference calls

Questions/Comments?

- Shumon Huque
- shuque [at] upenn.edu

Other topics (things I won't have time for, but I'm leaving the slides at the end) ..

Protection of signing Keys

- Offline not option (dynamic signing)
- Keep only KSK offline?
  - But need KSK for key rollovers
- Lock down signing server! (like KDCs?)
- Physically secured machine room, locked racks etc
- Tamper proof HSM in the future?

What about NSEC3?

- Might do it in the future ..
- Penn's DNS data is non-secret, *but I'd rather not have trivial zone enumeration. I'm slightly concerned that miscreants will be walking our zones all day just because they can*
- Looks relatively easy to transition ..
- With BIND 9.6, can transition by inserting NSEC3PARAM record into zone with Update

NSEC3 zone differences

- NSEC3 instead of NSEC records
- Owner is a cryptographic hash of the name rather than the actual name (provides zone enumeration defense)
- Not all names may have NSEC3 (opt-out feature)
- Additional apex record: NSEC3PARAM
- See RFC 5155 for details

Caveats & Concerns

- DNSSEC answers are larger
- Server side & query side impacts
- Firewalls, proxies, and other middleboxes botching EDNS0, large packets, DNSSEC records etc ...
- Many resolvers already ask for DNSSEC
  - Fallback to TCP increases?

Securing last hop (Stub resolver security)

- Validating Stub/Full Resolver on clients (goal?)
- Channel security mechanism between stub and recursive resolver:
  - TSIG
  - SIG(0)
  - IPSEC

Channel Security?

- Simple symmetric key TSIG won't work
  - Can't distribute same TSIG key to many clients, because that allows any one of them to forge answers to all others
  - Need per client keys and thus key management infrastructure
  - GSS-TSIG has chicken-egg problem (eg. DNS is often used to locate Kerberos servers)
- SIG(0) may be better (public key crypto)

Questions/Comments?

- Shumon Huque
- shuque [at] upenn.edu