ROVER
BGP Route Origin Verification via DNS

Dan Massey
Colorado State University/Maka'ala Networks

Joe Gersch
Secure64

Michael Glenn, Christopher Garner
Century Link

Introduction to Rover

- Basic Purpose: Protect against IP Hijacks
  - Origin Hijack
  - Sub-prefix Hijacks
  - Complementary technology to RPKI
    - Some similarities, some differences
- 2 Basic Components:
  - Publish
    - route origin data placed in the reverse-DNS, authenticated via DNSSEC signatures
  - Verify
    - SW tools and appliances to match unique ISP operational procedures

ROVER Design Model

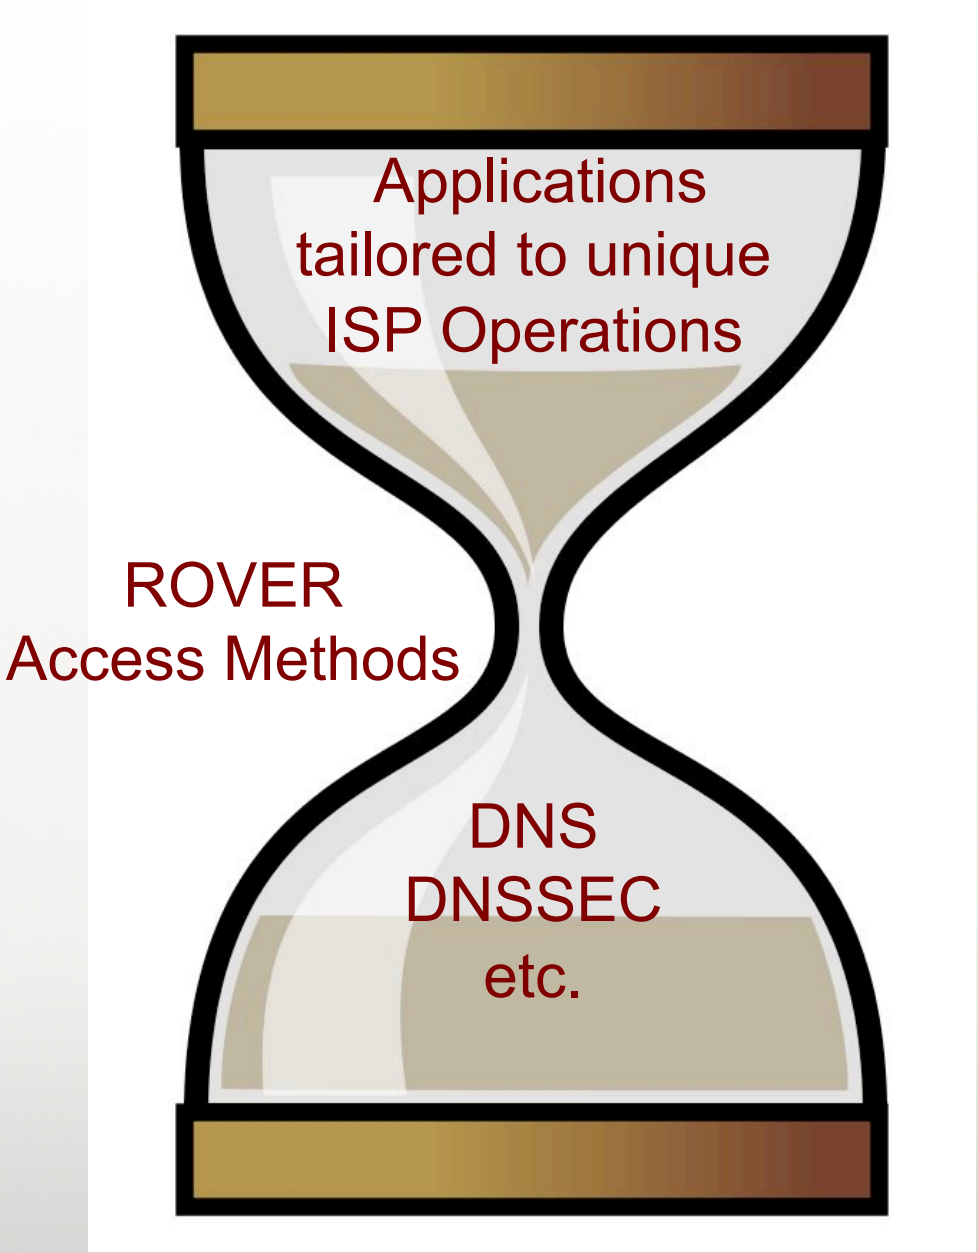

ROVER Design Model

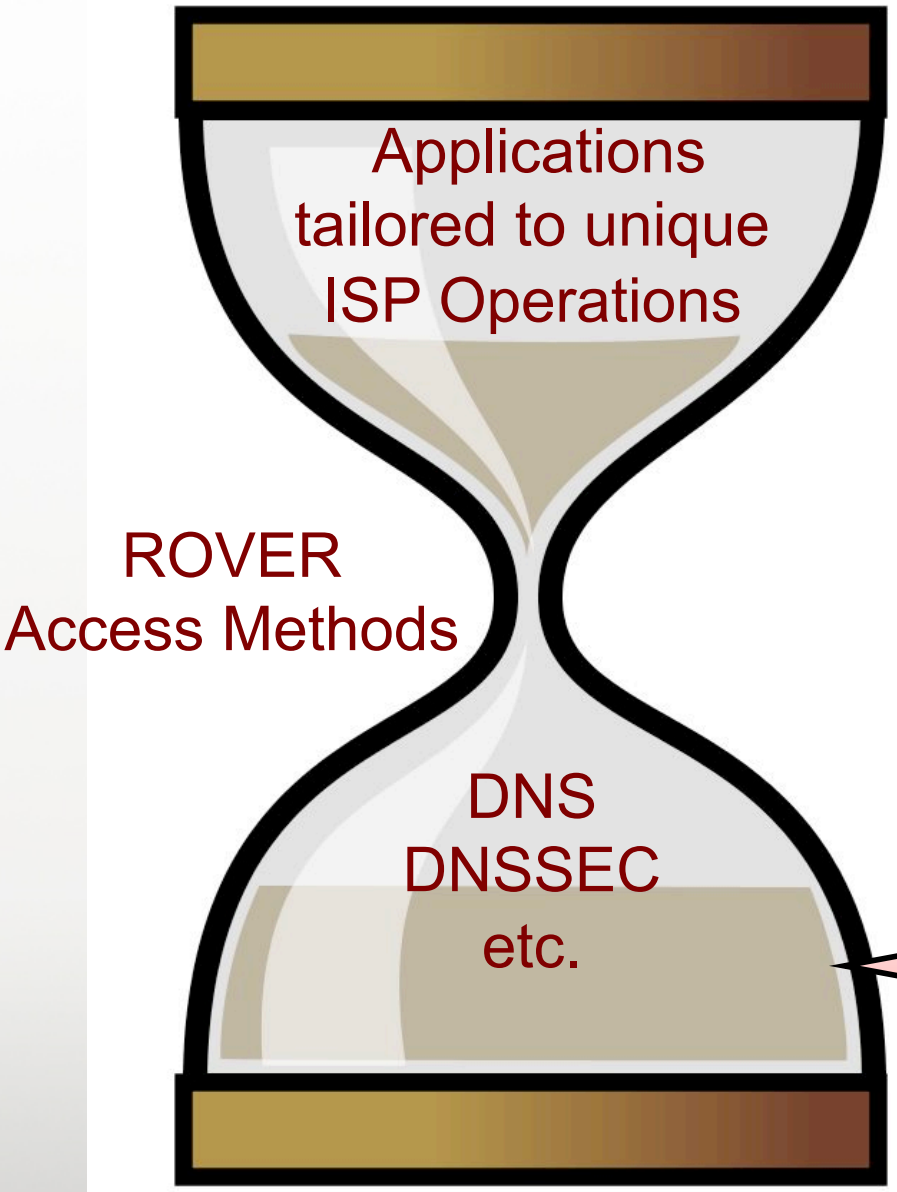

Foundations / Protocols

- pre-existing DNS infrastructure
- IN-ADDR.ARPA signed with DNSSEC
- redundancy/resiliency
- real-time updates

ROVER Design Model

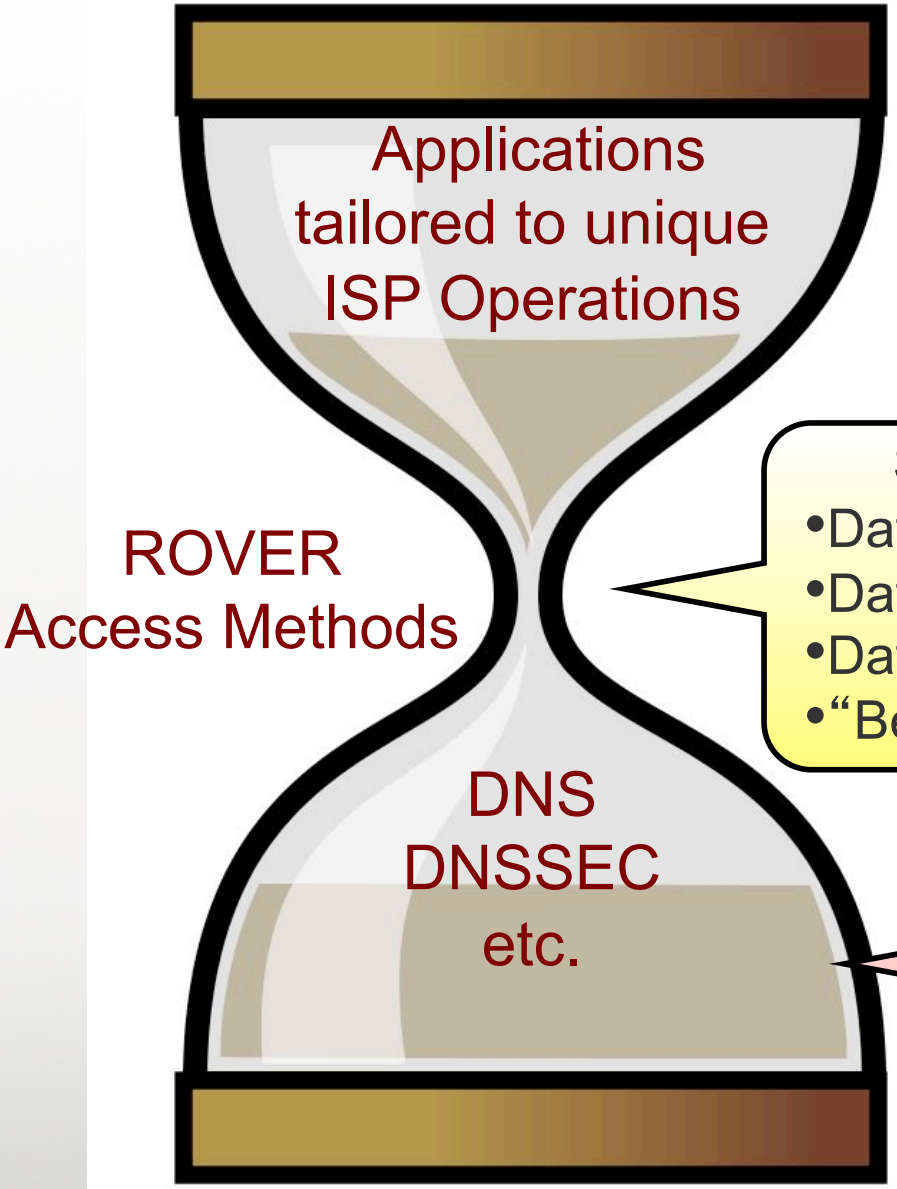

Small set of methods:
- Data Naming Convention
- Data Publishing Format
- Data Authentication
- "Best Effort" retrieval

Foundations / Protocols
- pre-existing DNS infrastructure
- IN-ADDR.ARPA signed with DNSSEC
- redundancy/resiliency
- real-time updates

ROVER Design Model

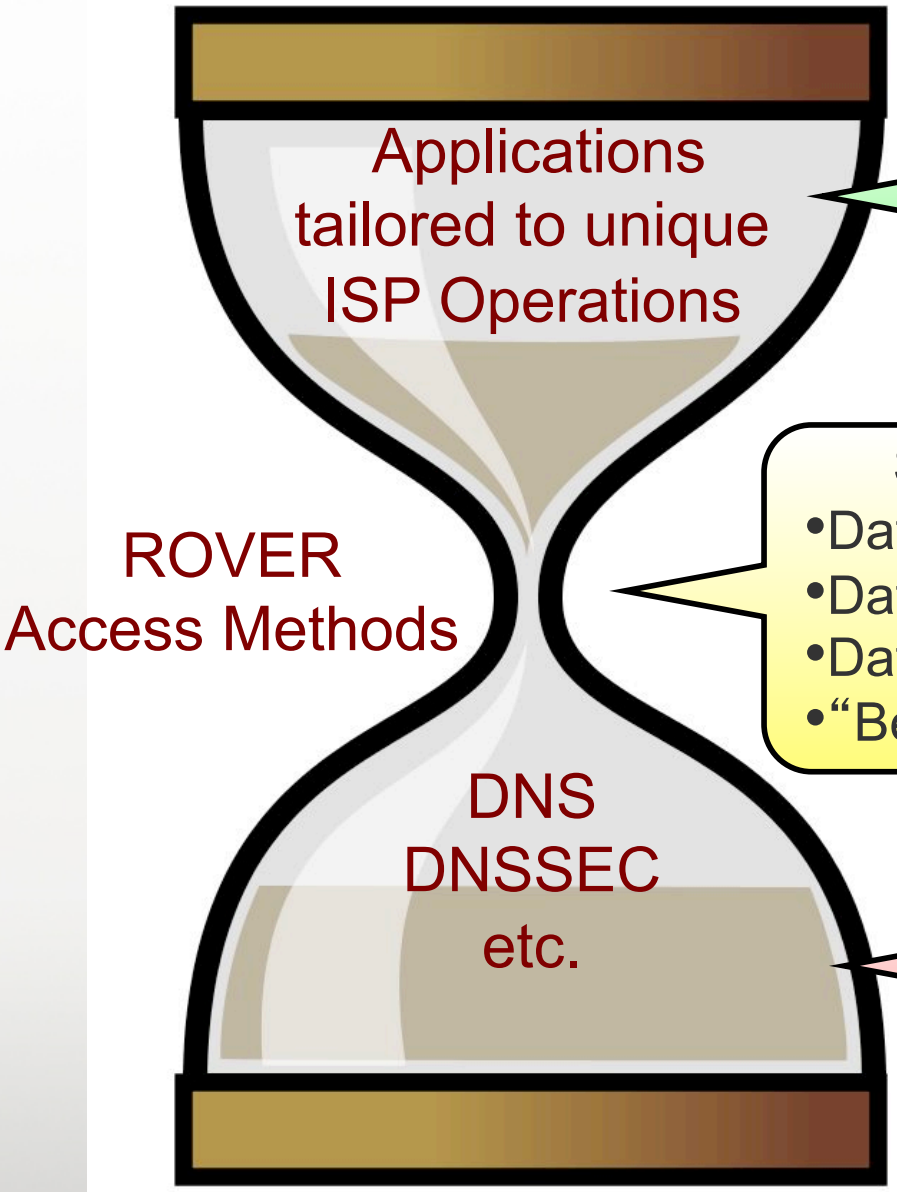

ISP Applications:
- verify routing DB
- verify prefix filters
- interface to RTR
- real-time notification/block
- (exponential retry)

Small set of methods:
- Data Naming Convention
- Data Publishing Format
- Data Authentication
- "Best Effort" retrieval

Foundations / Protocols
- pre-existing DNS infrastructure
- IN-ADDR.ARPA signed with DNSSEC
- redundancy/resiliency
- real-time updates

Publishing ROVER Data

Reverse DNS publishing method

- General-Purpose Naming convention designed to specify CIDR address blocks. Example:
  - 129.82.128.0/18 --> 0.1.m.82.129.in-addr.arpa
- 2 New DNS records
  - RLOCK: Route lock (opt in)
  - SRO: "Secure Route Origin"
- 2 Internet Drafts
  - `draft-gersch-dnsop-revdns-cidr`
  - `draft-gersch-sidr-revdns-bgp`

Three Methods of BGP Validation

- ROVER 1: Data published in the reverse DNS
  - Preferred method for publishing
  - You can do this now in your zones
- ROVER 2: Data published in a shadow reverse DNS
  - Mimics real reverse DNS
  - Allows one to experiment before moving to production
  - Tools for generating data
- COMPLETENESS: Historical BGP validation
  - Completes the picture so running at full scale
  - Clearly distinguished from real ROVER data

Examples of Published Data

- Reverse DNS (.in-addr.apra)
  - CenturyLink
  - Level3
  - Several smaller sites
- Shadow Reverse DNS
  - AT&T (tier 1 ISP)
  - HP (large enterprise)
  - CERT Australia (government)
  - Front Range Internet (regional ISP)
  - Modesto Irrigation District (critical infrastructure)
  - And many others
- Historical BGP Data
  - Global scale coverage with over 1.3 million records

Call To Action: Publish or Perish

- Two Mechanisms to Participate
  - Publish your data in the actual reverse DNS
  - Publish your data in the rover testbed
- Publishing Your Data in Reverse DNS
  - Need to enable DNSSEC in your reverse tree
  - Add RLOCK and SRO records to your existing zone(s)
  - Does not break existing zones
- Publishing in ROVER Shadow Reverse DNS
  - Auto-detects your prefixes
  - Allows you to confirm/customize entries
  - Builds zone file and stores in shadow reverse tree

Verifying ROVER Data

ROVER Verification

- The reverse DNS records can be used to:
  - Generate real-time alarms for a NOC
  - Verify route filters on a periodic basis
  - Perform real-time verifications
  - Other tools and building blocks
  - Subscribe to real-time global hijack notifications

Global Scale Rover

- Input from multiple BGP peers around the globe
  - Each prefix in an update triggers up to 3 verifications
- First verify against existing reverse DNS
  - Sending DNS queries to existing servers
  - Spans tens of thousands of existing zones
  - Typically results in NXDOMAIN (until more published)
- If nothing found, verify against shadow DNS Tree
  - Again send DNS queries, but to shadow servers
  - Larger hit rate
- If nothing found, verify against historical data
  - Over 1.3 million records providing global coverage
  - Database updates as new information learned

Searchable Web Based ROVER Database

Real-Time ROVER Verification

This page illustrates ROVER being used to verify announcements from 40+ BGP monitors located around the world. As each announcement arrives, ROVER does a DNS lookup to determine whether the route origin matches the DNS data. If no data is found, a comparision is made against a data base of historical route origins. Final results are displayed in the tables below.

This page refreshes the counters and speedometer every 10 seconds. The results table must be manually refreshed if you want to see new data.

Live BGP Feed
last update: Tue Aug 28 22:03:13 2012 (UTC)

	DNS Source: in-addr.arpa			DNS Source: Hosted in-addr.arpa			BGP History DB			
BGP announcements analyzed	VALID	ORIGIN HIJACK	SUBPREFIX HIJACK	VALID	ORIGIN HIJACK	SUBPREFIX HIJACK	VALID	ORIGIN HIJACK	NO DATA FOUND	ROUTE VALLEYS DETECTED
965400	0	0	0	52	0	0	875374	5039	84935	4099

You may sort this table in ascending/descending order by clicking on a column heading. Click on a row to see event details. Enter SEARCH FILTERs as desired.

Active BGP Hijack Events

Prefix	Event	Source	Origin	# Monito	# BGP	First Seen (UTC)	Last Seen (UTC)	Description
	ALL	ALL						
1.179.133.0/24	WARNING	HISTORY	9737	10	19	12-08-28 01:43:36	12-08-28 01:44:22	NO DATA Found; possible origin or sub-prefix hi
1.179.134.0/24	WARNING	HISTORY	9737	10	19	12-08-28 01:43:36	12-08-28 01:44:22	NO DATA Found; possible origin or sub-prefix hi
1.179.135.0/24	WARNING	HISTORY	9737	10	19	12-08-28 01:43:36	12-08-28 01:44:22	NO DATA Found; possible origin or sub-prefix hi
1.179.136.0/24	WARNING	HISTORY	9737	10	19	12-08-28 01:43:36	12-08-28 01:44:22	NO DATA Found; possible origin or sub-prefix hi
1.186.160.0/24	WARNING	HISTORY	45769	11	98	12-08-23 22:31:33	12-08-26 16:25:48	NO DATA Found; possible origin or sub-prefix hi
1.186.161.0/24	WARNING	HISTORY	45769	11	97	12-08-23 22:31:33	12-08-26 16:25:48	NO DATA Found; possible origin or sub-prefix hi
1.186.162.0/24	WARNING	HISTORY	45769	11	97	12-08-23 22:31:33	12-08-26 16:25:48	NO DATA Found; possible origin or sub-prefix hi
1.186.163.0/24	WARNING	HISTORY	45769	11	97	12-08-23 22:31:33	12-08-26 16:25:48	NO DATA Found; possible origin or sub-prefix hi
1.186.164.0/24	WARNING	HISTORY	45769	11	97	12-08-23 22:31:33	12-08-26 16:25:48	NO DATA Found; possible origin or sub-prefix hi

BGPMON data rate

The rate of the real-time BGP data stream.

Note: this page is undergoing extensive revisions. The monitor may or may not be running, and table formats will change.

New: Telnet to rover.secure64.com 40000 to see an XML data stream.

http:rover.secure64.com

BGP ROVER: Route Origin Verification

SECURE64

Learn More | Show Zones | Publish Route Origins | Verify Route Origin | Live BGP Feed

Return to Events Table

BGP Monitors Detecting Event

	Monitor IP	Monitor AS Name	# BGP	Active	First Seen (UTC)	Last Seen (UTC)	Map ID
+	192.43.217.144	14041: AS14041 - University Corporati	3	☑	12-08-24 19:12:52	12-08-28 21:43:52	?
−	65.49.129.101	3043: INTERNET-OPERATING-SERVICES -	4	☐	12-08-24 19:13:03	12-08-28 21:41:22	A
	time	path					
	12-08-28 21:41:22	WITHDRAW					
	12-08-27 19:51:26	[3043, 174, 6453, 55644, 45271]					
	12-08-24 19:13:15	[3043, 174, 6453, 55644, 45271]					
	12-08-24 19:13:03	[3043, 6086, 22773, 2828, 6453, 55644, 45271]					
+	164.128.32.11	3303: SWISSCOM Swisscom (Switzerland	5	☑	12-08-23 23:30:56	12-08-28 00:42:00	B
+	202.167.228.37	38809: NXGNET-AS-AP Nextgen Networks	5	☑	12-08-24 19:11:16	12-08-27 19:51:51	C

ROVER XML Real-Time Data Feed

jgersch — Python — 92×29

```
<ROVER version="0.1">
    <EVENT>
        <HEADER>
            <PREFIX>
                92.124.64.0/18
            </PREFIX>
            <OWNER_ASN>
                41440
            </OWNER_ASN>
            <AS_PATH>
                [28289, 53131, 16735, 12956, 3257, 12389, 41440, 41440, 41440, 5573]
            </AS_PATH>
        </HEADER>
        <EVENT_STATS event_id="7631" num_BGP_announcements="12" num_peers_involved="8">
            <FIRST_SEEN datetime="2012-08-28T21:53:03Z" timestamp="1346190783"/>
            <LAST_SEEN datetime="2012-08-28T21:53:29Z" timestamp="1346190809"/>
        </EVENT_STATS>
        <BGPMON_DATA seq_id="2128112124" seq_num="1533369785">
            <TIME datetime="2012-08-28T21:53:06Z" timestamp="1346190786"/>
            <COLLECTOR as_name="Americana Digital Ltda." asn="28289" ip="189.36.224.1" latit
ude="-10.0" longitude="-55.0"/>
        </BGPMON_DATA>
        <INCIDENT>
            <HIJACK source="HISTORY" type="ORIGIN HIJACK">
                origin does not match historical origin(s): 41440
            </HIJACK>
        </INCIDENT>
    </EVENT>
</ROVER>
```

telnet rover.secure64.com 40000

Call To Action: Verification

- Verify your own BGP updates
  - Anyone can issue DNS queries to reverse DNS
  - ROVER testbed also open to queries

    dig <prefix>.in-addr.arpa.secure64.com

- ROVER Verification
  - ROVER Website Analysis Tools
  - XML Stream of Alerts

- Interested in feedback on how to adapt to your operations

Questions

http://rover.secure64.com

Additional Information

Reverse DNS Challenges…

- You didn't build a new PKI
  - Yes, we intentionally did not introduce a new PKI
  - We work the existing reverse DNS tree that has been in operation for decades
- Can you store CIDR prefixes in the reverse DNS?
  - Yes, we use the existing structure on octet boundaries (/8, /16, /24) for IPv4 and on nibbles for IPv6
  - We introduce naming for structure for non-octet boundaries
- Does DNSSEC really work and is it deployed?
  - Yes, DNSSEC is well established at this point
  - ARIN, RIPE, and others have deployed DNSSEC

DNS Management Challenges…

- Can I have different groups manage PTR and routing records in DNS?
  - Yes, you can separate use the “m” label in the name to create different zones for routing records?
- So I have to create new zones?
  - No, you can store records in your existing zones. Create new zones only if it aides your operational practices.

Avoiding Dependencies

- Can a low-level protocol like BGP depend on a higher-level protocol?
  - no, not if there is a hard dependency
  - yes, if the dependency has a “fail-safe”
- Rover uses “best effort” data retrieval with world-wide data distribution, redundancy and local caching. Applications can use query retries with exponential back-off.
- If the data is unreachable, the default is that routing works just as it works today.

Pre-Loading Verification Data

- Do I need to download a database of all authorized routes?
  - no, there are different verification modules and it is not assumed you have a pre-generated copy
- Could I preload a database of authorized routes?
  - Yes, you can use a past routing table, RouteViews table, IRR, or any number of other mechanisms to obtain a list of prefixes and verify them ahead of time