The DNS Camel

Or

How many features can we add to this protocol before it breaks?

Bert Hubert / bert.hubert@powerdns.com

RFC	Type	Status	Title	Bgnd	Prot	Names	Ops	RR	Proxy	Stub	Auth	Res	Xfr	DDNS	DNSSEC
882		Obsolete	Domain Names – Concepts and Facilities	x		x	x				x				
883		Obsolete	Domain Names – Implementation and Specification		x		x	x			x	x			
920			Domain Requirements				x								
973		Obsolete	Domain System Changes and Observations			x		x			x	x			
1032			Domain Administrators Guide				x								
1033			Domain Administrators Operations Guide				x								
1034	Standard		Domain Names – Concepts and Facilities	x		x	x			x	x	x			
1035	Standard		Domain Names – Implementation and Specification		x	x		x			x	x	x		
1101			DNS Encoding of Network Names and Other Types			x									
1123	Standard		Requirements for Internet Hosts – Application and Support	x							x	x			
1178	Informational		Choosing a Name for Your Computer				x								
1183	Experimental		New DNS RR Definitions					x							
1348	Experimental	Obsolete	DNS NSAP RRs					x							
1401	Informational		Correspondence between the IAB and DISA on the use of DNS throughout the Internet	x											
1535	Informational		A Security Problem and Proposed Correction With Widely Deployed DNS Software									x			
1536	Informational		Common DNS Implementation Errors and Suggested Fixes							x		x			
1537	Informational	Obsolete	Common DNS Data File Configuration Errors				x								
1591	Informational		Domain Name System Structure and Delegation				x								
1611	Historic	Historic	DNS Server MIB Extensions				x								
1612	Historic	Historic	DNS Resolver MIB Extensions				x								

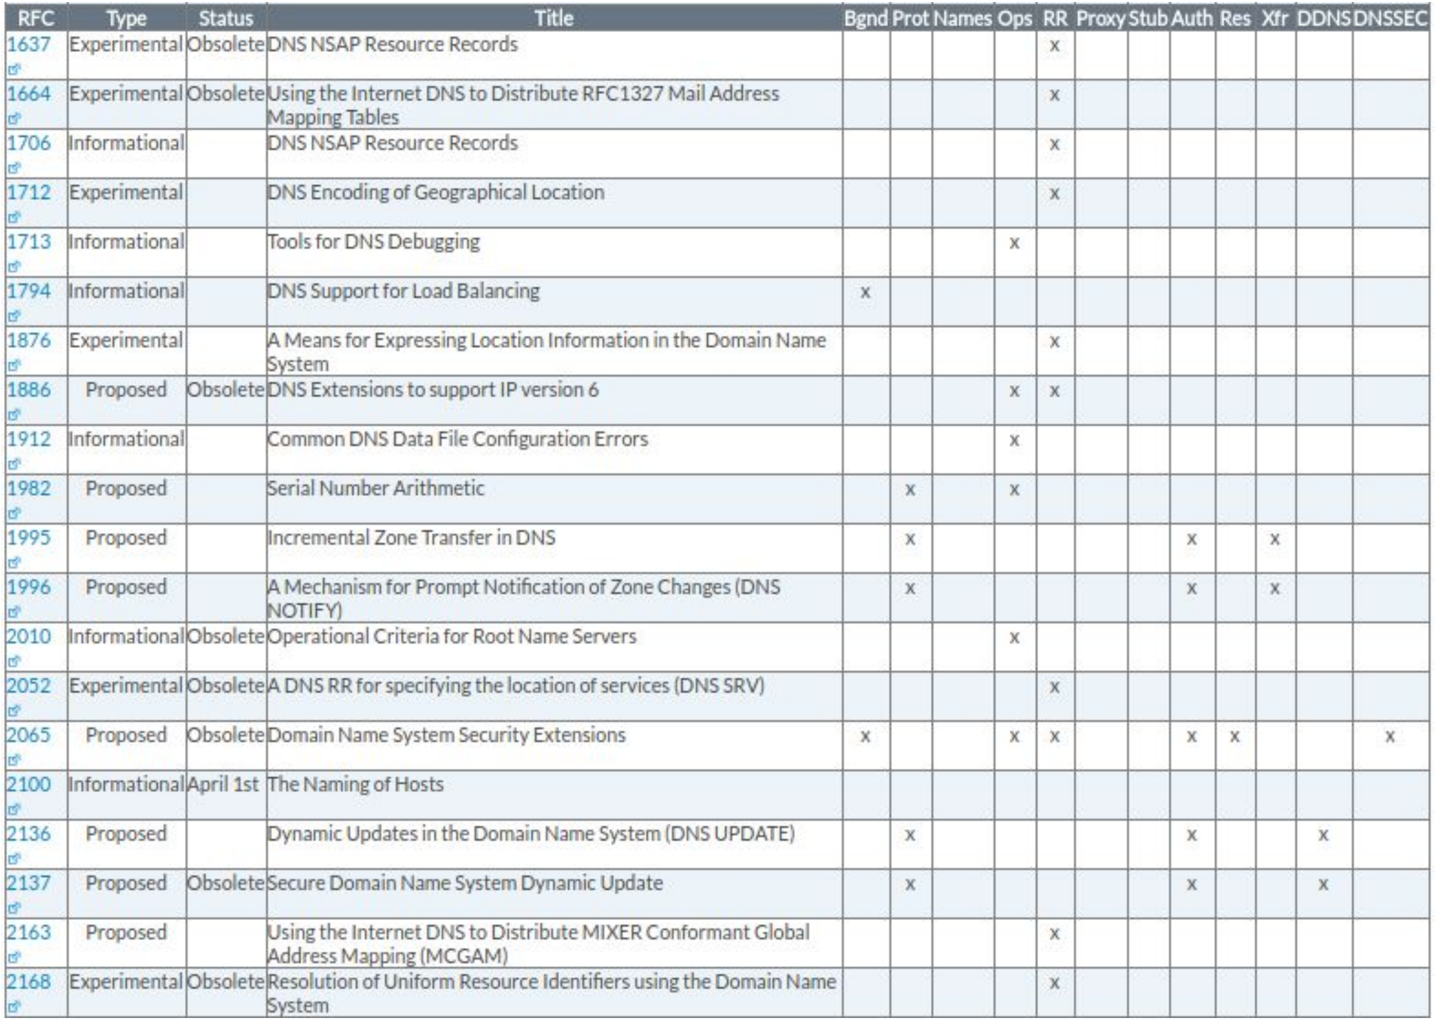

RFC	Type	Status	Title	Bgnd	Prot	Names	Ops	RR	Proxy	Stub	Auth	Res	Xfr	DDNS	DNSSEC
1637	Experimental	Obsolete	DNS NSAP Resource Records					x							
1664	Experimental	Obsolete	Using the Internet DNS to Distribute RFC1327 Mail Address Mapping Tables					x							
1706	Informational		DNS NSAP Resource Records					x							
1712	Experimental		DNS Encoding of Geographical Location					x							
1713	Informational		Tools for DNS Debugging				x								
1794	Informational		DNS Support for Load Balancing	x											
1876	Experimental		A Means for Expressing Location Information in the Domain Name System					x							
1886	Proposed	Obsolete	DNS Extensions to support IP version 6				x	x							
1912	Informational		Common DNS Data File Configuration Errors				x								
1982	Proposed		Serial Number Arithmetic		x		x								
1995	Proposed		Incremental Zone Transfer in DNS		x						x		x		
1996	Proposed		A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)		x						x		x		
2010	Informational	Obsolete	Operational Criteria for Root Name Servers				x								
2052	Experimental	Obsolete	A DNS RR for specifying the location of services (DNS SRV)					x							
2065	Proposed	Obsolete	Domain Name System Security Extensions	x			x	x			x	x			x
2100	Informational	April 1st	The Naming of Hosts												
2136	Proposed		Dynamic Updates in the Domain Name System (DNS UPDATE)		x						x			x	
2137	Proposed	Obsolete	Secure Domain Name System Dynamic Update		x						x			x	
2163	Proposed		Using the Internet DNS to Distribute MIXER Conformant Global Address Mapping (MCGAM)					x							
2168	Experimental	Obsolete	Resolution of Uniform Resource Identifiers using the Domain Name System					x							

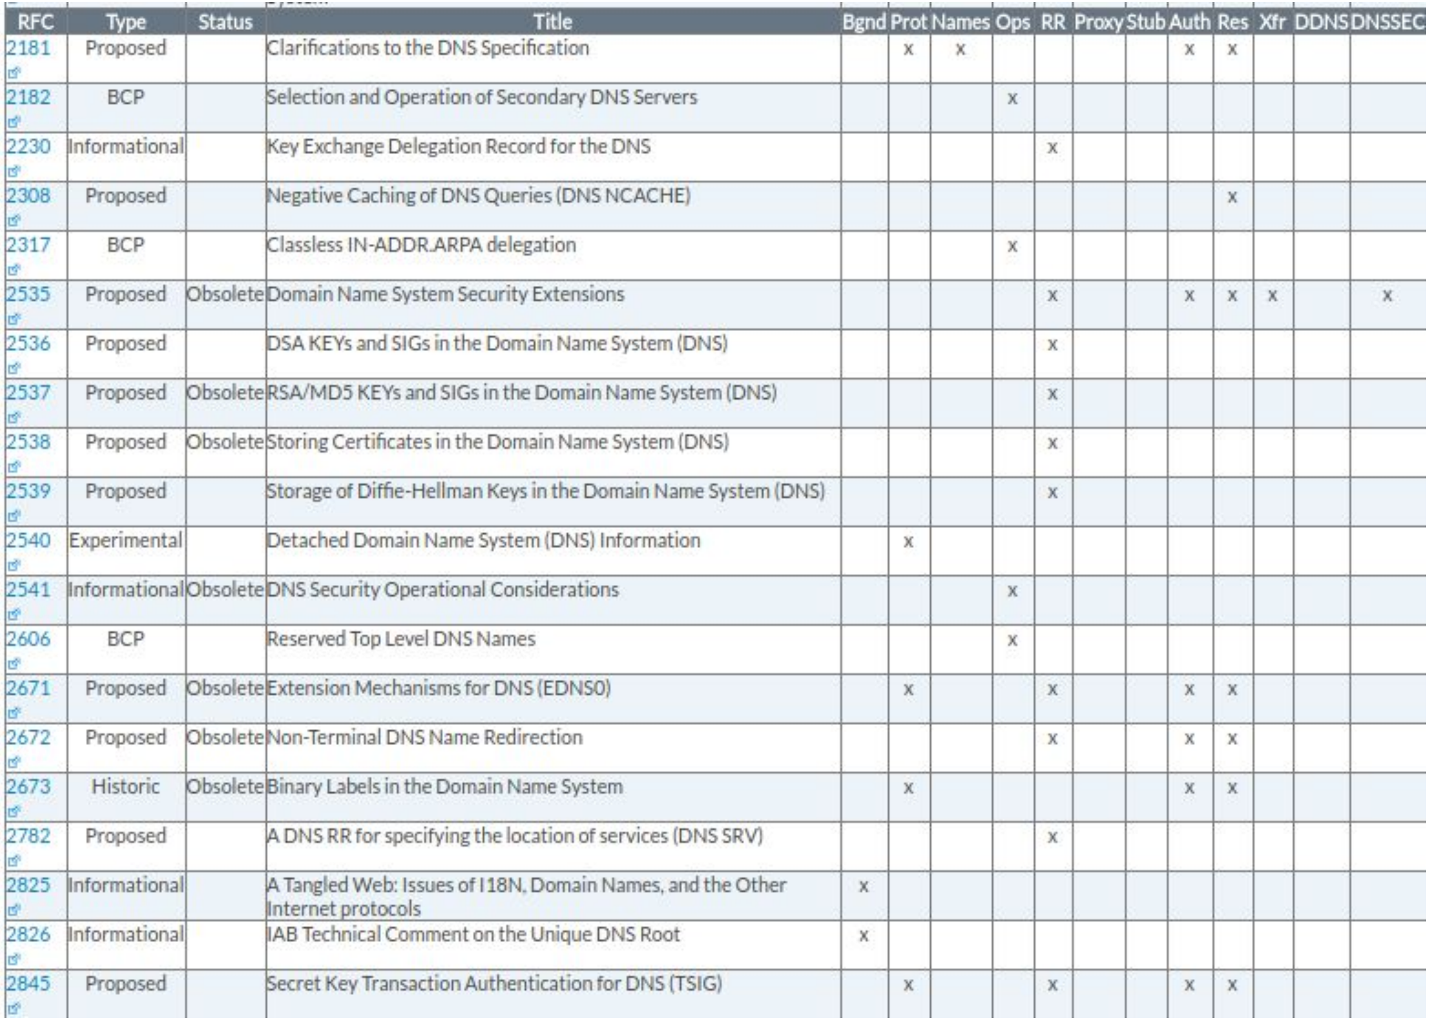

RFC	Type	Status	Title	Bgnd	Prot	Names	Ops	RR	Proxy	Stub	Auth	Res	Xfr	DDNS	DNSSEC
2181	Proposed		Clarifications to the DNS Specification		x	x					x	x			
2182	BCP		Selection and Operation of Secondary DNS Servers				x								
2230	Informational		Key Exchange Delegation Record for the DNS					x							
2308	Proposed		Negative Caching of DNS Queries (DNS NCACHE)									x			
2317	BCP		Classless IN-ADDR.ARPA delegation				x								
2535	Proposed	Obsolete	Domain Name System Security Extensions					x			x	x	x		x
2536	Proposed		DSA KEYs and SIGs in the Domain Name System (DNS)					x							
2537	Proposed	Obsolete	RSA/MD5 KEYs and SIGs in the Domain Name System (DNS)					x							
2538	Proposed	Obsolete	Storing Certificates in the Domain Name System (DNS)					x							
2539	Proposed		Storage of Diffie-Hellman Keys in the Domain Name System (DNS)					x							
2540	Experimental		Detached Domain Name System (DNS) Information		x										
2541	Informational	Obsolete	DNS Security Operational Considerations				x								
2606	BCP		Reserved Top Level DNS Names				x								
2671	Proposed	Obsolete	Extension Mechanisms for DNS (EDNS0)		x			x			x	x			
2672	Proposed	Obsolete	Non-Terminal DNS Name Redirection					x			x	x			
2673	Historic	Obsolete	Binary Labels in the Domain Name System		x						x	x			
2782	Proposed		A DNS RR for specifying the location of services (DNS SRV)					x							
2825	Informational		A Tangled Web: Issues of I18N, Domain Names, and the Other Internet protocols	x											
2826	Informational		IAB Technical Comment on the Unique DNS Root	x											
2845	Proposed		Secret Key Transaction Authentication for DNS (TSIG)		x			x			x	x			

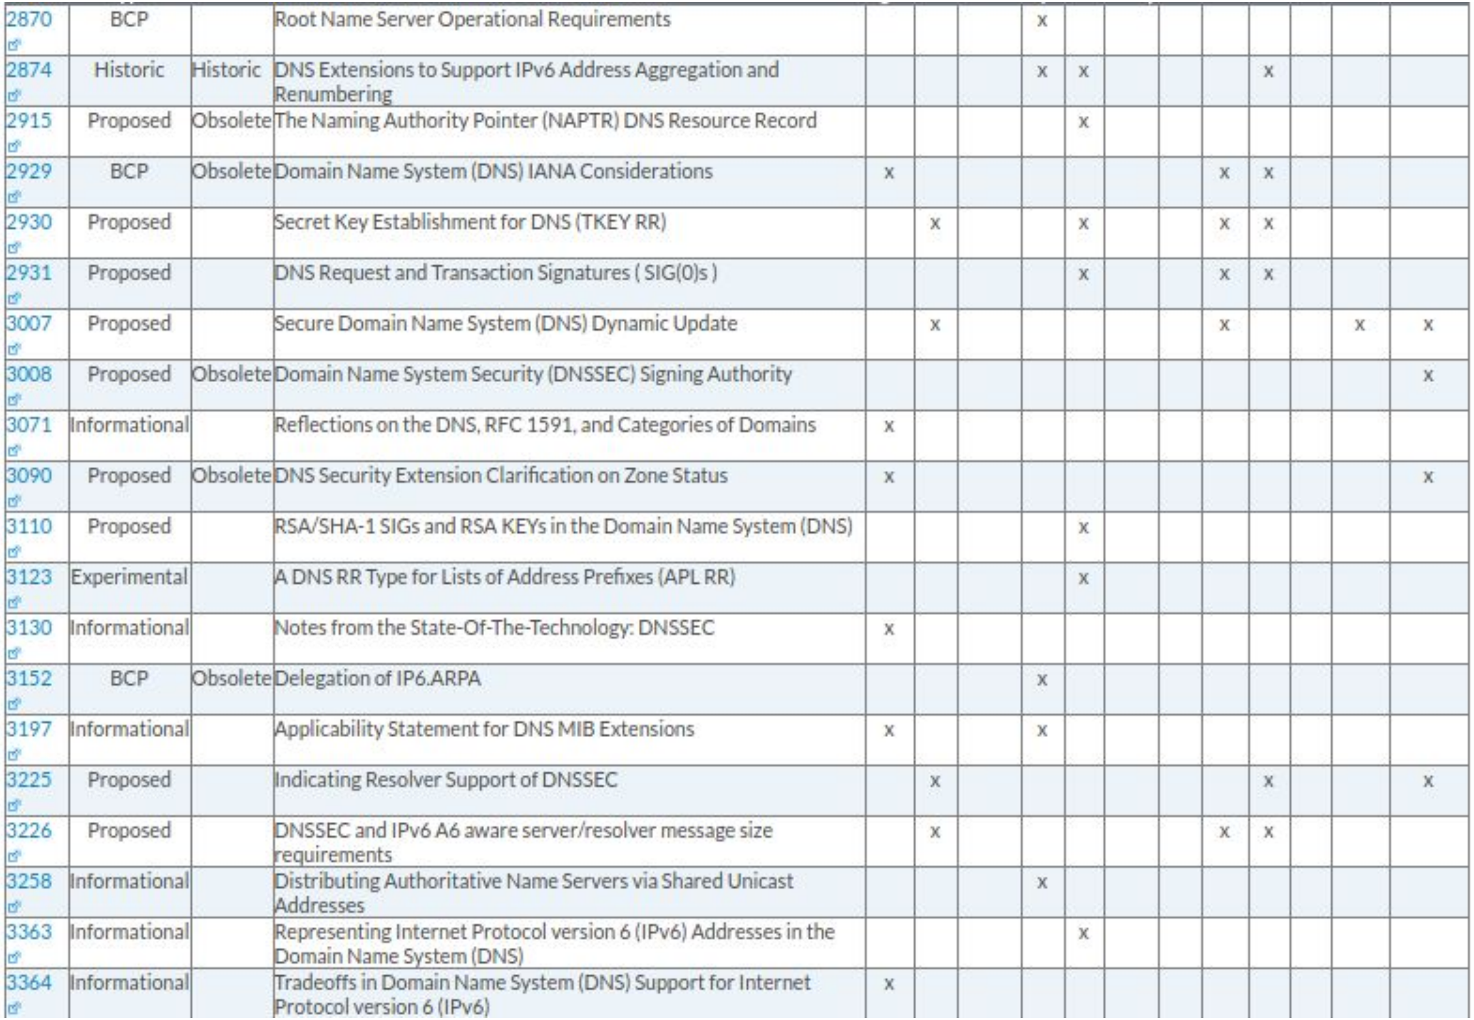

2870	BCP		Root Name Server Operational Requirements				x								
2874	Historic	Historic	DNS Extensions to Support IPv6 Address Aggregation and Renumbering				x	x				x			
2915	Proposed	Obsolete	The Naming Authority Pointer (NAPTR) DNS Resource Record					x							
2929	BCP	Obsolete	Domain Name System (DNS) IANA Considerations	x							x	x			
2930	Proposed		Secret Key Establishment for DNS (TKEY RR)		x			x			x	x			
2931	Proposed		DNS Request and Transaction Signatures (SIG(0)s)					x			x	x			
3007	Proposed		Secure Domain Name System (DNS) Dynamic Update		x						x			x	x
3008	Proposed	Obsolete	Domain Name System Security (DNSSEC) Signing Authority												x
3071	Informational		Reflections on the DNS, RFC 1591, and Categories of Domains	x											
3090	Proposed	Obsolete	DNS Security Extension Clarification on Zone Status	x											x
3110	Proposed		RSA/SHA-1 SIGs and RSA KEYs in the Domain Name System (DNS)					x							
3123	Experimental		A DNS RR Type for Lists of Address Prefixes (APL RR)					x							
3130	Informational		Notes from the State-Of-The-Technology: DNSSEC	x											
3152	BCP	Obsolete	Delegation of IP6.ARPA				x								
3197	Informational		Applicability Statement for DNS MIB Extensions	x			x								
3225	Proposed		Indicating Resolver Support of DNSSEC		x							x			x
3226	Proposed		DNSSEC and IPv6 A6 aware server/resolver message size requirements		x						x	x			
3258	Informational		Distributing Authoritative Name Servers via Shared Unicast Addresses				x								
3363	Informational		Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)					x							
3364	Informational		Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)	x											

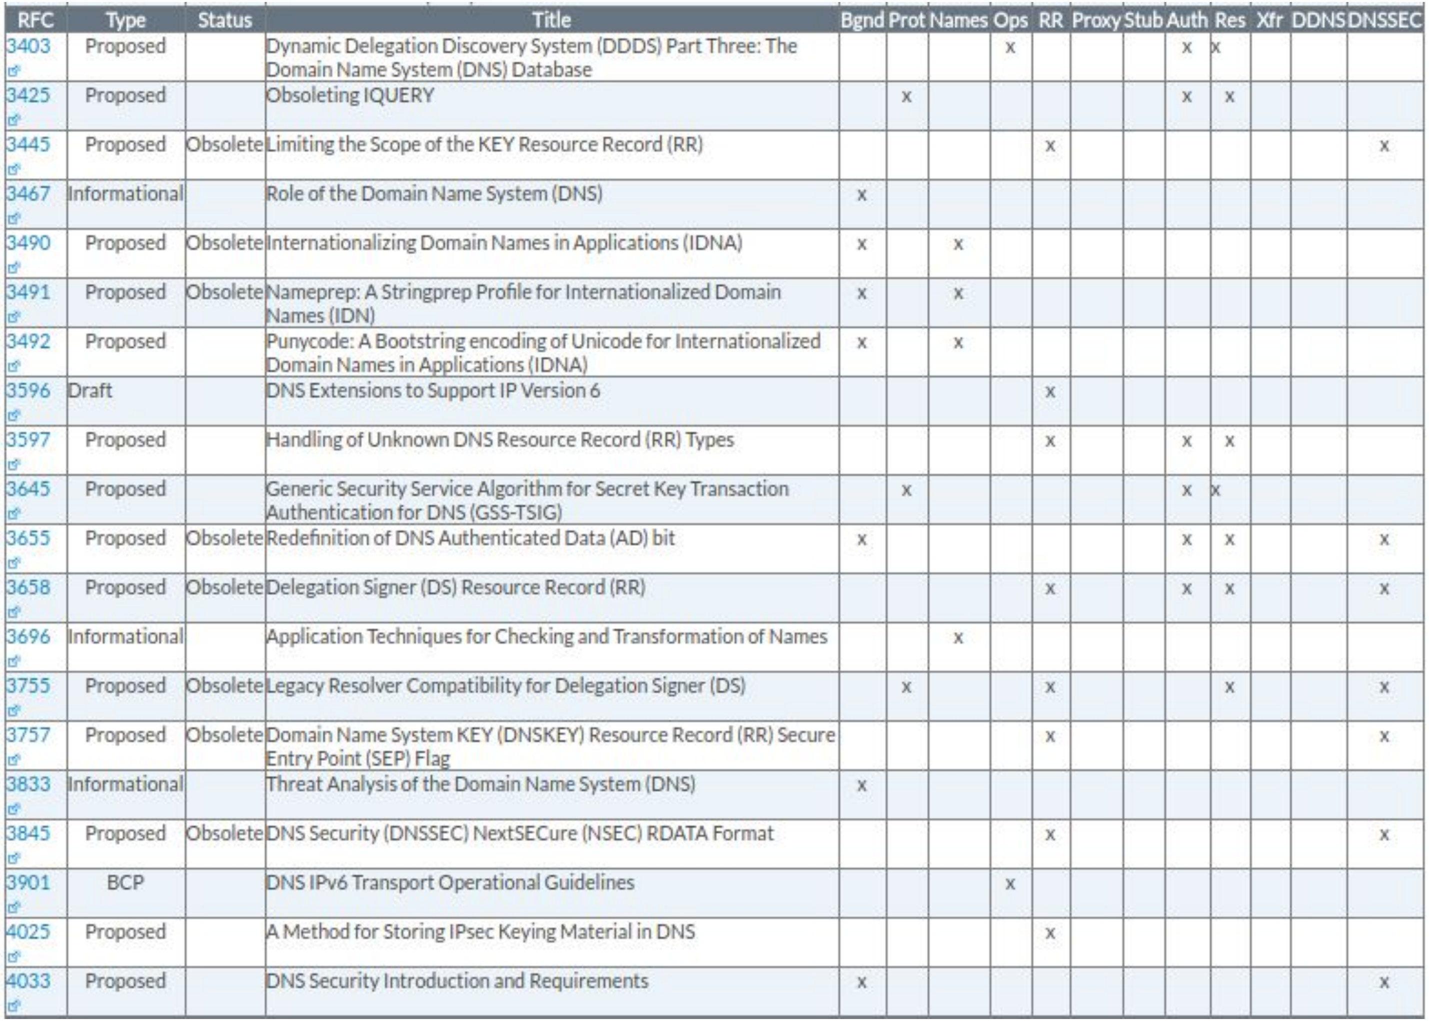

RFC	Type	Status	Title	Bgnd	Prot	Names	Ops	RR	Proxy	Stub	Auth	Res	Xfr	DDNS	DNSSEC
3403	Proposed		Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database				x				x	x			
3425	Proposed		Obsoleting IQUERY		x						x	x			
3445	Proposed	Obsolete	Limiting the Scope of the KEY Resource Record (RR)					x							x
3467	Informational		Role of the Domain Name System (DNS)	x											
3490	Proposed	Obsolete	Internationalizing Domain Names in Applications (IDNA)	x		x									
3491	Proposed	Obsolete	Nameprep: A Stringprep Profile for Internationalized Domain Names (IDN)	x		x									
3492	Proposed		Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)	x		x									
3596	Draft		DNS Extensions to Support IP Version 6					x							
3597	Proposed		Handling of Unknown DNS Resource Record (RR) Types					x			x	x			
3645	Proposed		Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)		x						x	x			
3655	Proposed	Obsolete	Redefinition of DNS Authenticated Data (AD) bit	x							x	x			x
3658	Proposed	Obsolete	Delegation Signer (DS) Resource Record (RR)					x			x	x			x
3696	Informational		Application Techniques for Checking and Transformation of Names			x									
3755	Proposed	Obsolete	Legacy Resolver Compatibility for Delegation Signer (DS)		x			x				x			x
3757	Proposed	Obsolete	Domain Name System KEY (DNSKEY) Resource Record (RR) Secure Entry Point (SEP) Flag					x							x
3833	Informational		Threat Analysis of the Domain Name System (DNS)	x											
3845	Proposed	Obsolete	DNS Security (DNSSEC) NextSECure (NSEC) RDATA Format					x							x
3901	BCP		DNS IPv6 Transport Operational Guidelines				x								
4025	Proposed		A Method for Storing IPsec Keying Material in DNS					x							
4033	Proposed		DNS Security Introduction and Requirements	x											x

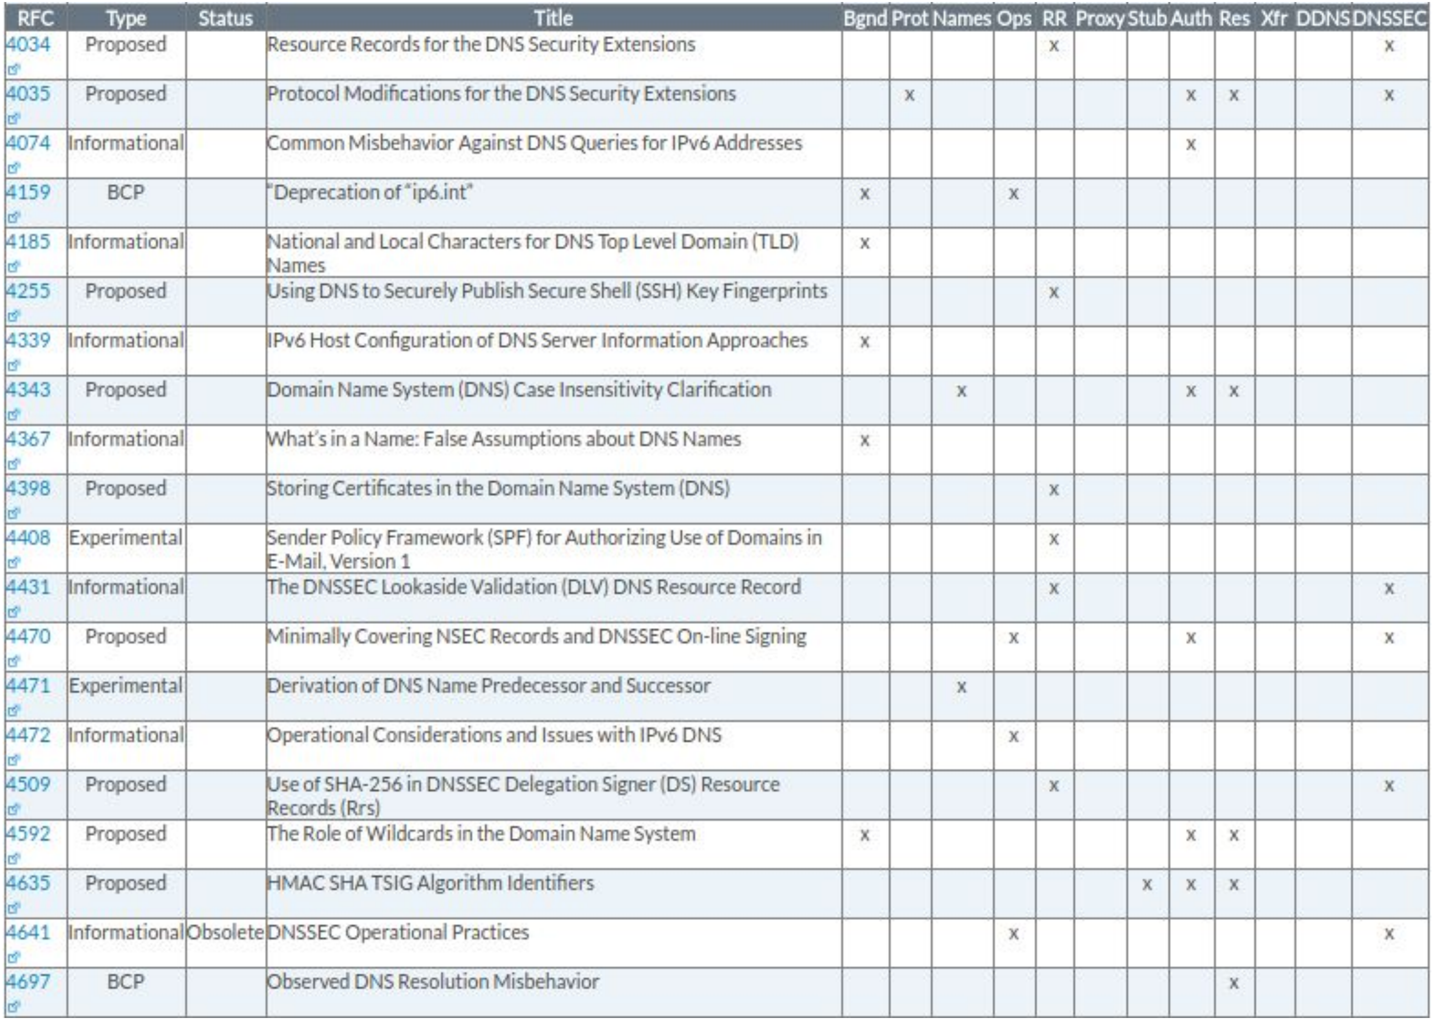

RFC	Type	Status	Title	Bgnd	Prot	Names	Ops	RR	Proxy	Stub	Auth	Res	Xfr	DDNS	DNSSEC
4034	Proposed		Resource Records for the DNS Security Extensions					x							x
4035	Proposed		Protocol Modifications for the DNS Security Extensions		x						x	x			x
4074	Informational		Common Misbehavior Against DNS Queries for IPv6 Addresses								x				
4159	BCP		“Deprecation of “ip6.int”	x			x								
4185	Informational		National and Local Characters for DNS Top Level Domain (TLD) Names	x											
4255	Proposed		Using DNS to Securely Publish Secure Shell (SSH) Key Fingerprints					x							
4339	Informational		IPv6 Host Configuration of DNS Server Information Approaches	x											
4343	Proposed		Domain Name System (DNS) Case Insensitivity Clarification			x					x	x			
4367	Informational		What's in a Name: False Assumptions about DNS Names	x											
4398	Proposed		Storing Certificates in the Domain Name System (DNS)					x							
4408	Experimental		Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1					x							
4431	Informational		The DNSSEC Lookaside Validation (DLV) DNS Resource Record					x							x
4470	Proposed		Minimally Covering NSEC Records and DNSSEC On-line Signing				x				x				x
4471	Experimental		Derivation of DNS Name Predecessor and Successor			x									
4472	Informational		Operational Considerations and Issues with IPv6 DNS				x								
4509	Proposed		Use of SHA-256 in DNSSEC Delegation Signer (DS) Resource Records (Rrs)					x							x
4592	Proposed		The Role of Wildcards in the Domain Name System	x							x	x			
4635	Proposed		HMAC SHA TSIG Algorithm Identifiers							x	x	x			
4641	Informational	Obsolete	DNSSEC Operational Practices				x								x
4697	BCP		Observed DNS Resolution Misbehavior									x			

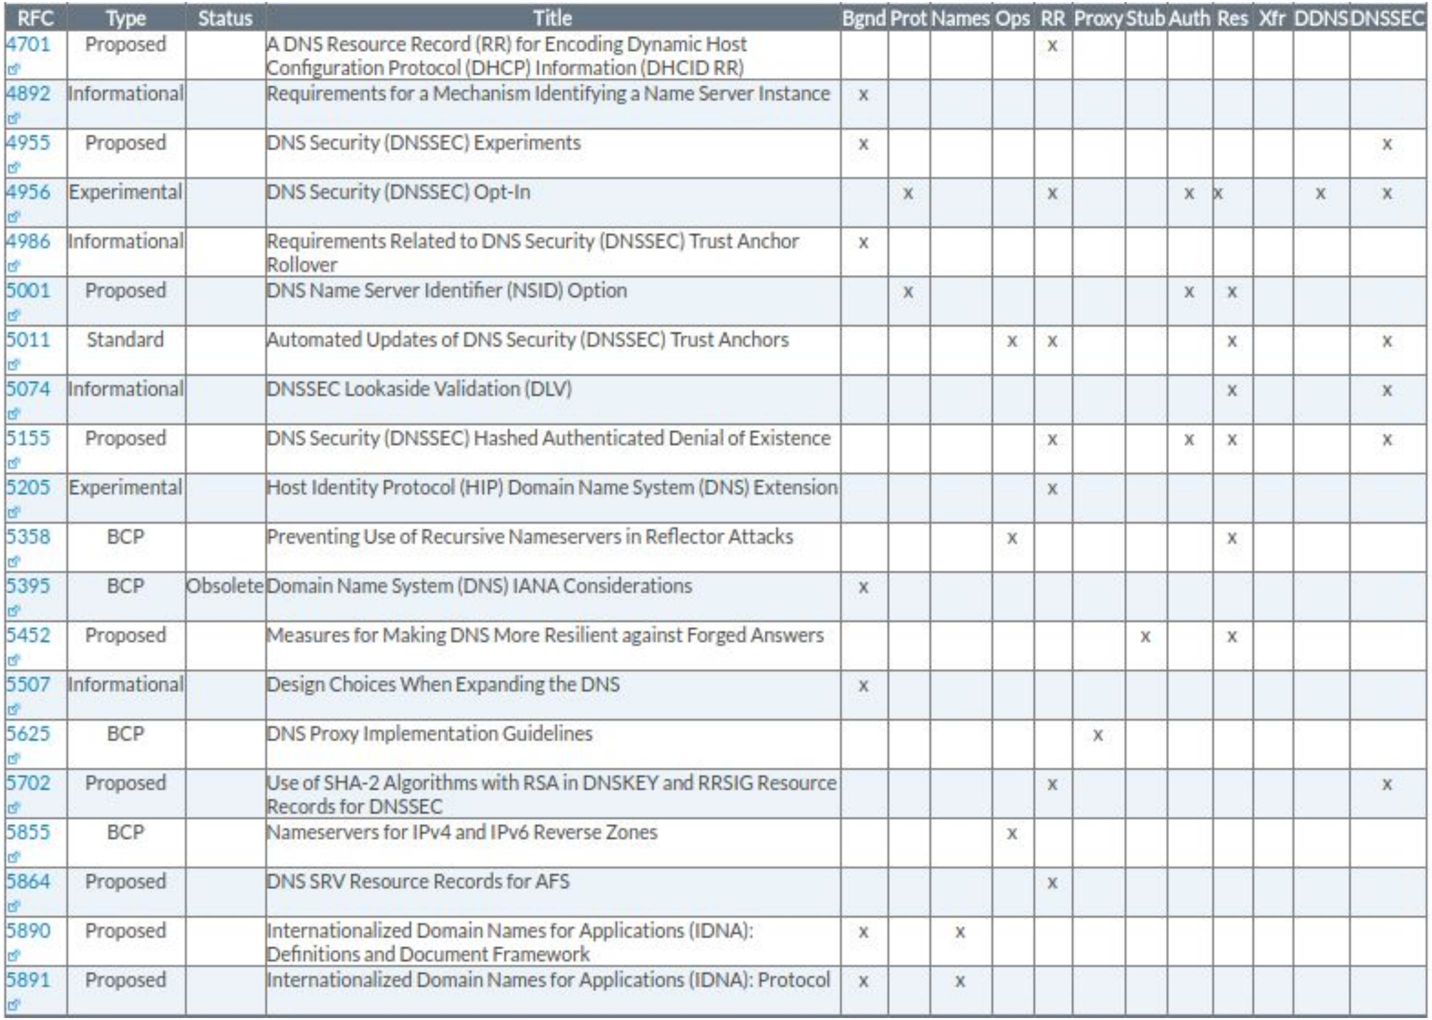

RFC	Type	Status	Title	Bgnd	Prot	Names	Ops	RR	Proxy	Stub	Auth	Res	Xfr	DDNS	DNSSEC
4701	Proposed		A DNS Resource Record (RR) for Encoding Dynamic Host Configuration Protocol (DHCP) Information (DHCID RR)					x							
4892	Informational		Requirements for a Mechanism Identifying a Name Server Instance	x											
4955	Proposed		DNS Security (DNSSEC) Experiments	x											x
4956	Experimental		DNS Security (DNSSEC) Opt-In		x			x			x	x		x	x
4986	Informational		Requirements Related to DNS Security (DNSSEC) Trust Anchor Rollover	x											
5001	Proposed		DNS Name Server Identifier (NSID) Option		x						x	x			
5011	Standard		Automated Updates of DNS Security (DNSSEC) Trust Anchors				x	x				x			x
5074	Informational		DNSSEC Lookaside Validation (DLV)									x			x
5155	Proposed		DNS Security (DNSSEC) Hashed Authenticated Denial of Existence					x			x	x			x
5205	Experimental		Host Identity Protocol (HIP) Domain Name System (DNS) Extension					x							
5358	BCP		Preventing Use of Recursive Nameservers in Reflector Attacks				x					x			
5395	BCP	Obsolete	Domain Name System (DNS) IANA Considerations	x											
5452	Proposed		Measures for Making DNS More Resilient against Forged Answers							x		x			
5507	Informational		Design Choices When Expanding the DNS	x											
5625	BCP		DNS Proxy Implementation Guidelines						x						
5702	Proposed		Use of SHA-2 Algorithms with RSA in DNSKEY and RRSIG Resource Records for DNSSEC					x							x
5855	BCP		Nameservers for IPv4 and IPv6 Reverse Zones				x								
5864	Proposed		DNS SRV Resource Records for AFS					x							
5890	Proposed		Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework	x		x									
5891	Proposed		Internationalized Domain Names for Applications (IDNA): Protocol	x		x									

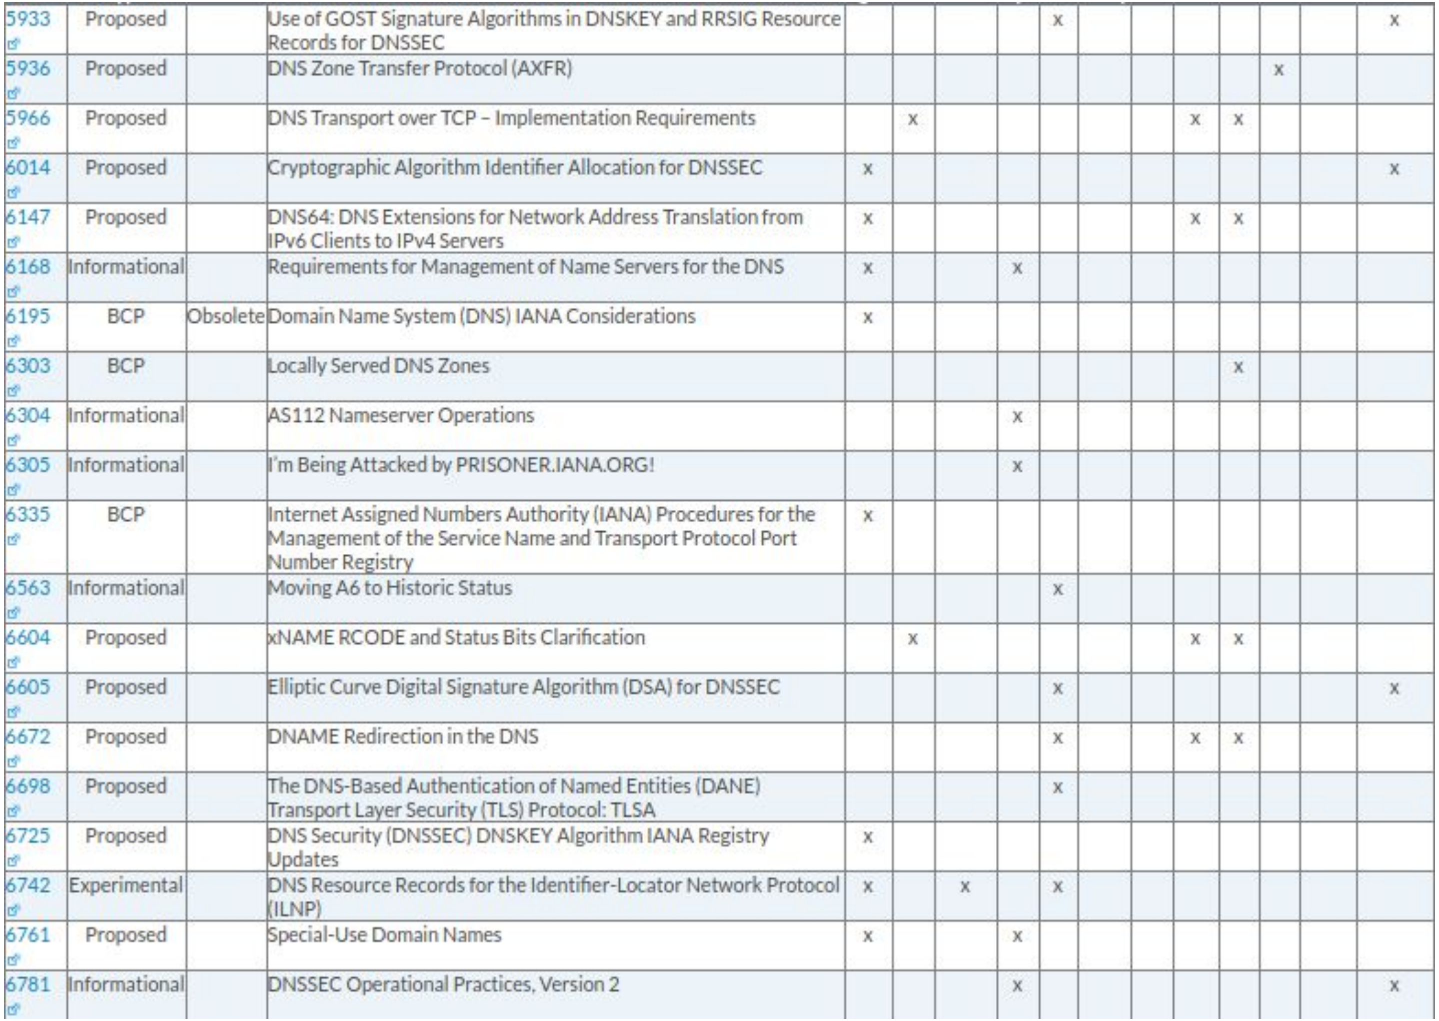

5933	Proposed		Use of GOST Signature Algorithms in DNSKEY and RRSIG Resource Records for DNSSEC					x							x
5936	Proposed		DNS Zone Transfer Protocol (AXFR)										x		
5966	Proposed		DNS Transport over TCP – Implementation Requirements		x						x	x			
6014	Proposed		Cryptographic Algorithm Identifier Allocation for DNSSEC	x											x
6147	Proposed		DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers	x							x	x			
6168	Informational		Requirements for Management of Name Servers for the DNS	x			x								
6195	BCP	Obsolete	Domain Name System (DNS) IANA Considerations	x											
6303	BCP		Locally Served DNS Zones									x			
6304	Informational		AS112 Nameserver Operations				x								
6305	Informational		I'm Being Attacked by PRISONER.IANA.ORG!				x								
6335	BCP		Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry	x											
6563	Informational		Moving A6 to Historic Status					x							
6604	Proposed		xNAME RCODE and Status Bits Clarification		x						x	x			
6605	Proposed		Elliptic Curve Digital Signature Algorithm (DSA) for DNSSEC					x							x
6672	Proposed		DNAME Redirection in the DNS					x			x	x			
6698	Proposed		The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA					x							
6725	Proposed		DNS Security (DNSSEC) DNSKEY Algorithm IANA Registry Updates	x											
6742	Experimental		DNS Resource Records for the Identifier-Locator Network Protocol (ILNP)	x		x		x							
6761	Proposed		Special-Use Domain Names	x			x								
6781	Informational		DNSSEC Operational Practices, Version 2				x								x

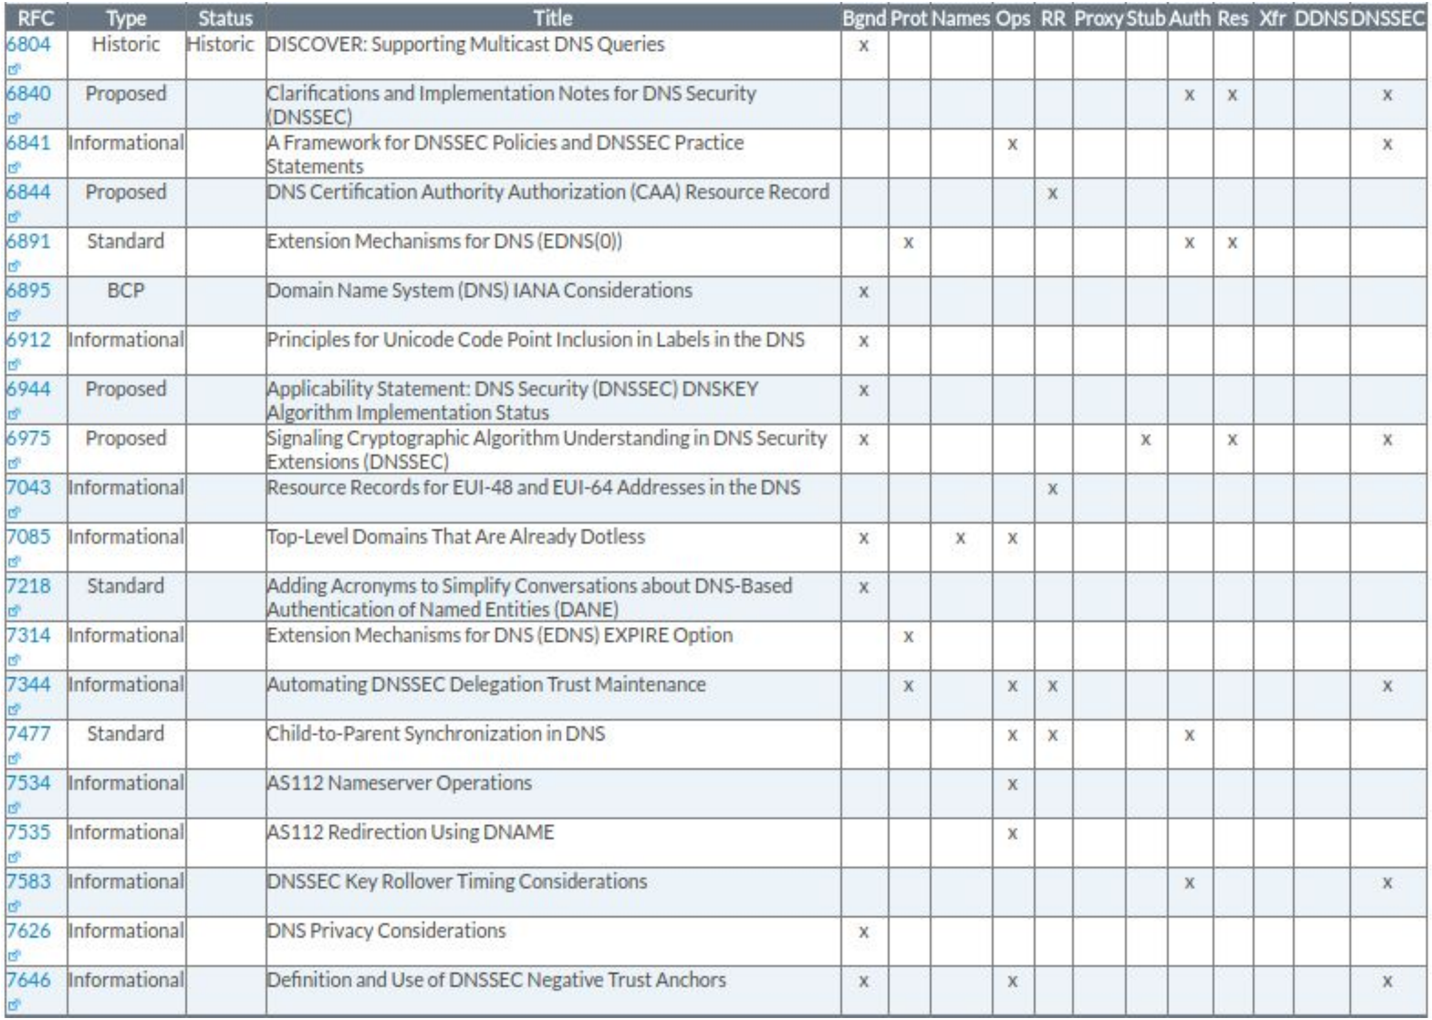

RFC	Type	Status	Title	Bgnd	Prot	Names	Ops	RR	Proxy	Stub	Auth	Res	Xfr	DDNS	DNSSEC
6804	Historic	Historic	DISCOVER: Supporting Multicast DNS Queries	x											
6840	Proposed		Clarifications and Implementation Notes for DNS Security (DNSSEC)								x	x			x
6841	Informational		A Framework for DNSSEC Policies and DNSSEC Practice Statements				x								x
6844	Proposed		DNS Certification Authority Authorization (CAA) Resource Record					x							
6891	Standard		Extension Mechanisms for DNS (EDNS(0))		x						x	x			
6895	BCP		Domain Name System (DNS) IANA Considerations	x											
6912	Informational		Principles for Unicode Code Point Inclusion in Labels in the DNS	x											
6944	Proposed		Applicability Statement: DNS Security (DNSSEC) DNSKEY Algorithm Implementation Status	x											
6975	Proposed		Signaling Cryptographic Algorithm Understanding in DNS Security Extensions (DNSSEC)	x						x		x			x
7043	Informational		Resource Records for EUI-48 and EUI-64 Addresses in the DNS					x							
7085	Informational		Top-Level Domains That Are Already Dotless	x		x	x								
7218	Standard		Adding Acronyms to Simplify Conversations about DNS-Based Authentication of Named Entities (DANE)	x											
7314	Informational		Extension Mechanisms for DNS (EDNS) EXPIRE Option		x										
7344	Informational		Automating DNSSEC Delegation Trust Maintenance		x		x	x							x
7477	Standard		Child-to-Parent Synchronization in DNS				x	x			x				
7534	Informational		AS112 Nameserver Operations				x								
7535	Informational		AS112 Redirection Using DNAME				x								
7583	Informational		DNSSEC Key Rollover Timing Considerations								x				x
7626	Informational		DNS Privacy Considerations	x											
7646	Informational		Definition and Use of DNSSEC Negative Trust Anchors	x			x								x

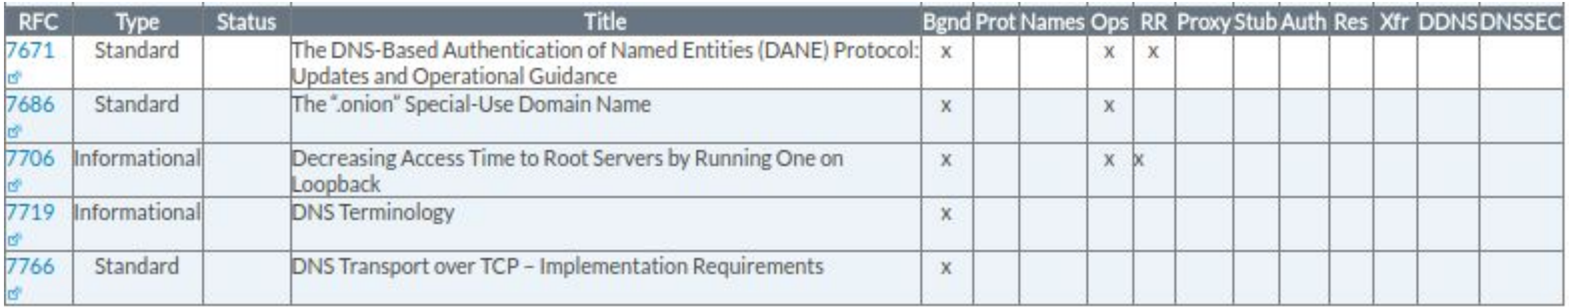

RFC	Type	Status	Title	Bgnd	Prot	Names	Ops	RR	Proxy	Stub	Auth	Res	Xfr	DDNS	DNSSEC
7671	Standard		The DNS-Based Authentication of Named Entities (DANE) Protocol: Updates and Operational Guidance	x			x	x							
7686	Standard		The ".onion" Special-Use Domain Name	x			x								
7706	Informational		Decreasing Access Time to Root Servers by Running One on Loopback	x			x	x							
7719	Informational		DNS Terminology	x											
7766	Standard		DNS Transport over TCP – Implementation Requirements	x											

185 RFCs

2781 pages / 166891 lines

888233 words

This is 2 times “The C++ Programming Language” (4th ed)

Good words on this are in RFC 8324

In the field stub resolver

```
char resppacket[512];

unsigned int ip_address;

char *ptr=resppacket+12;

/* receive */

while(!(*ptr==0xc0 && *(ptr+1)==0x0c)) ptr++;

memcpy(&ip_address, ptr+6, 4);
```

Did not read 1 of those 2781 pages

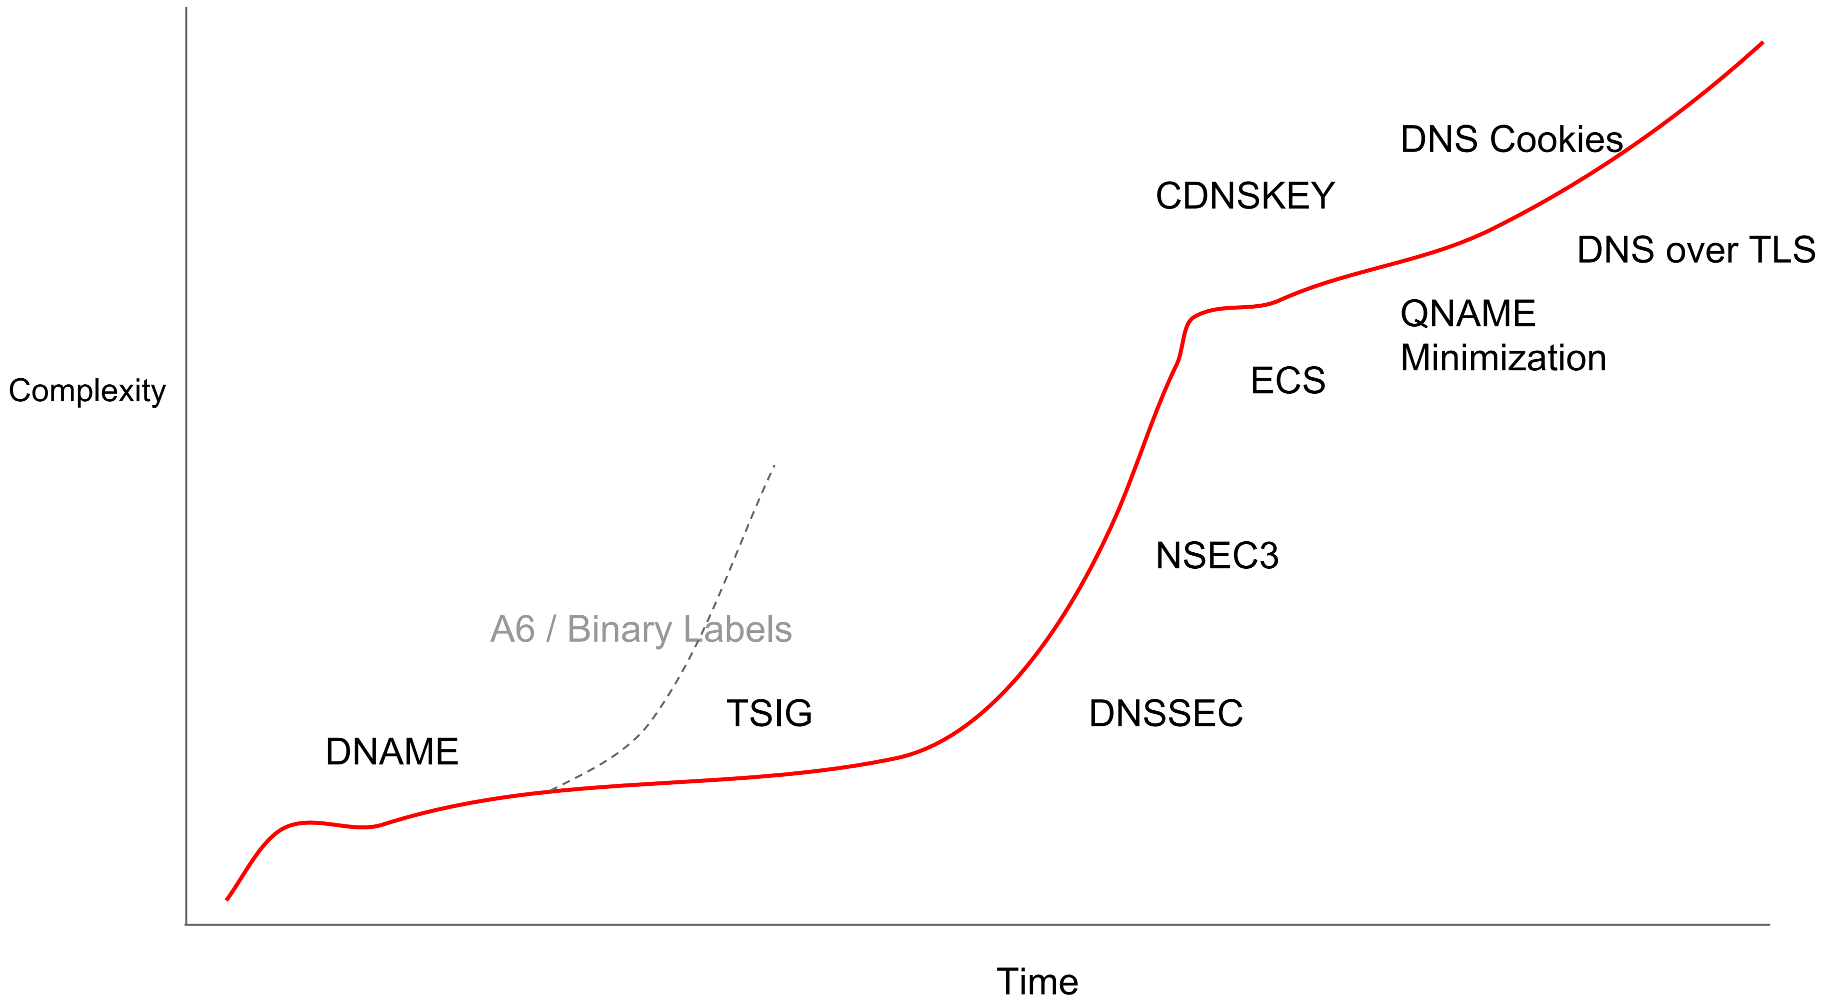
Complexity
Time
DNAME
A6 / Binary Labels
TSIG
DNSSEC
NSEC3
ECS
CDNSKEY
DNS Cookies
QNAME
Minimization
DNS over TLS

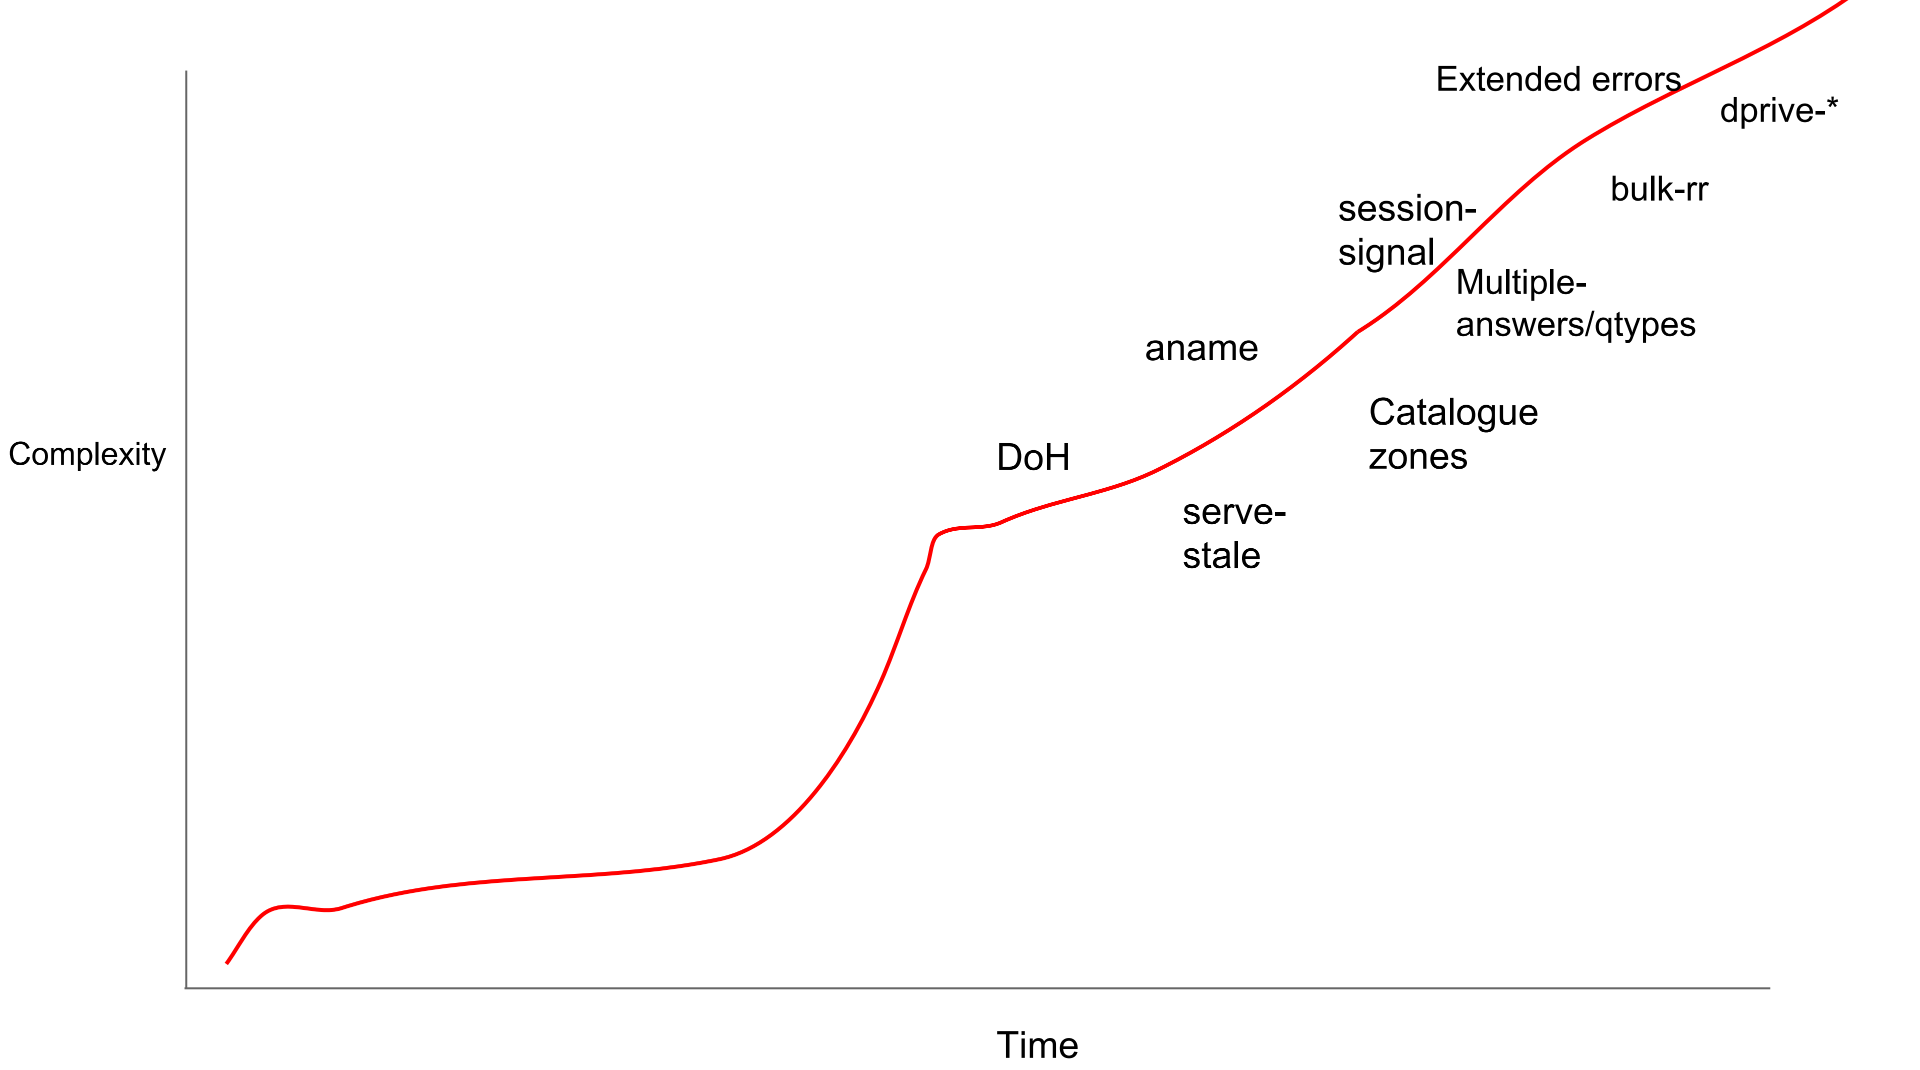
Complexity
Time
DoH
serve-stale
aname
Catalogue zones
session-signal
Multiple-answers/qtypes
Extended errors
bulk-rr
dprive-*

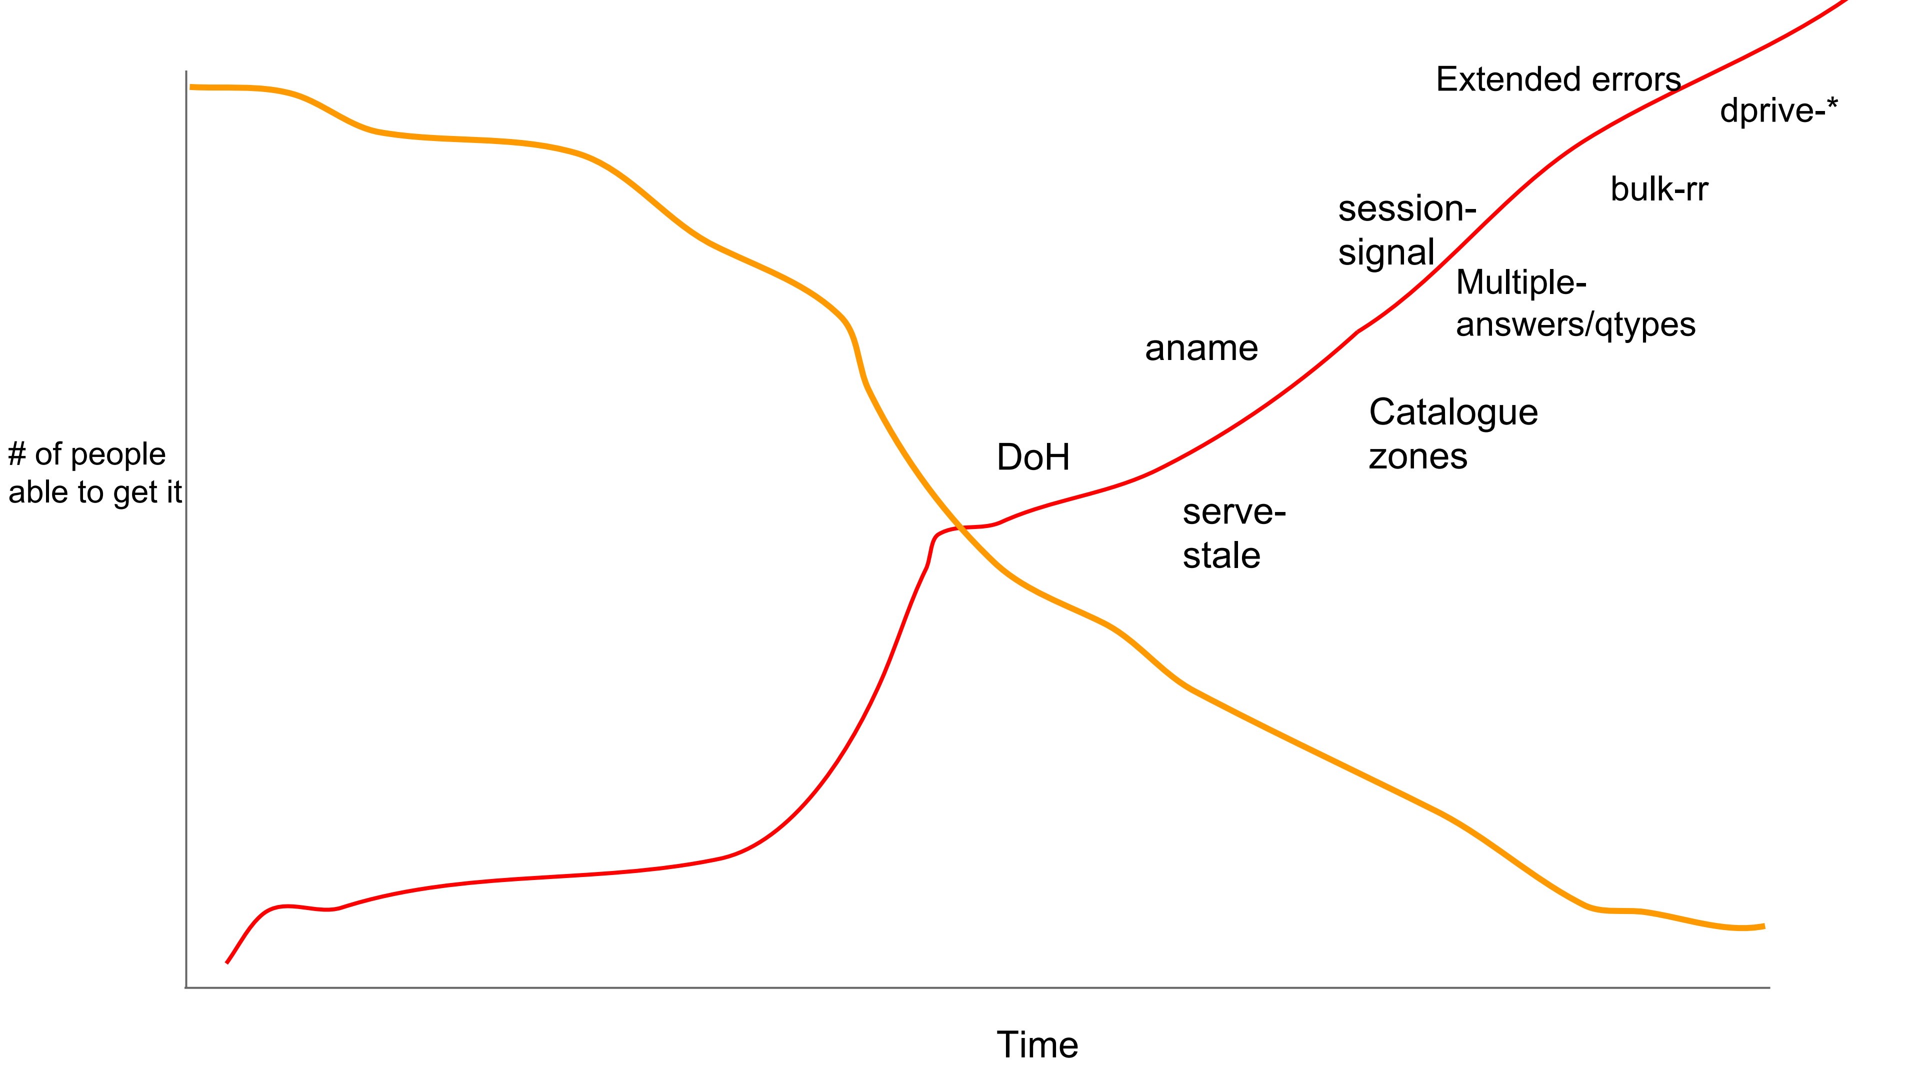
Extended errors
dprive-*
bulk-rr
session-signal
Multiple-answers/qtypes
aname
Catalogue zones
DoH
serve-stale
# of people able to get it
Time

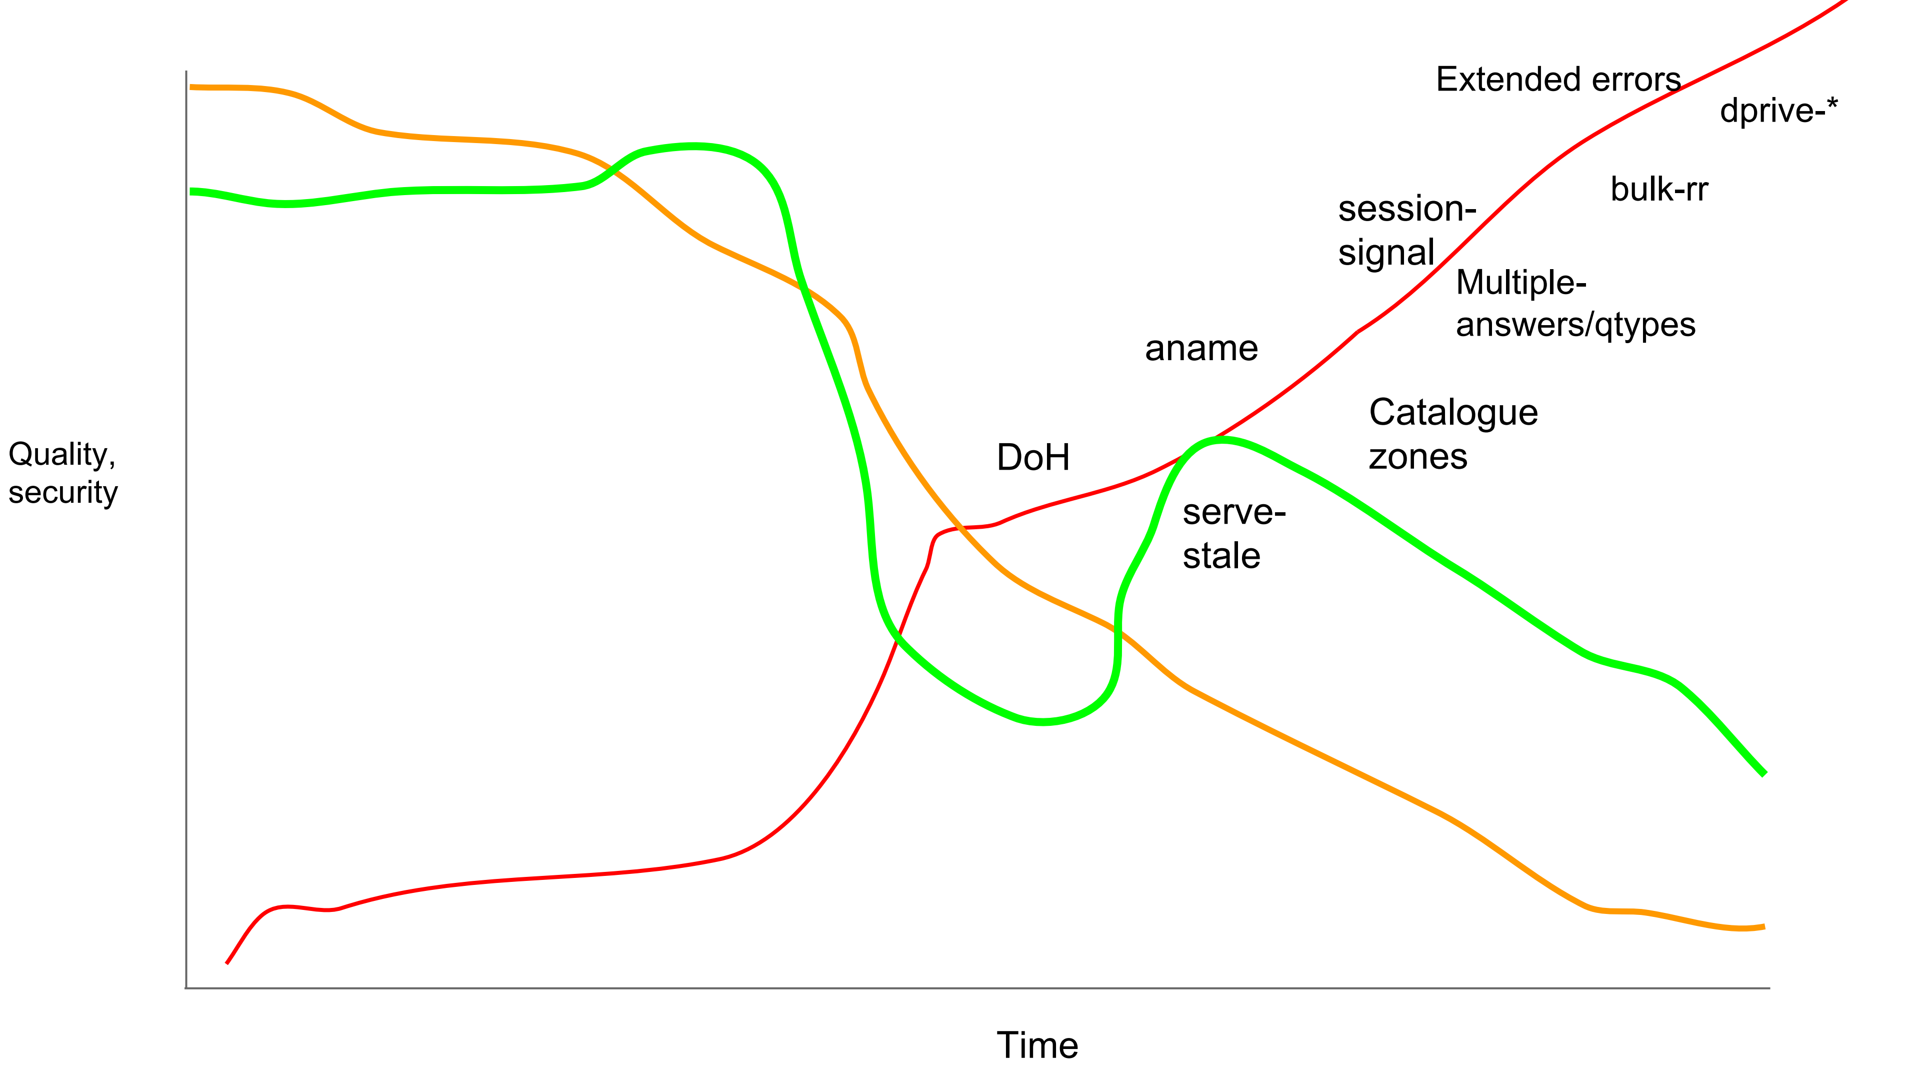
Quality, security
Time
Extended errors
dprive-*
bulk-rr
session-signal
Multiple-answers/qtypes
aname
Catalogue zones
DoH
serve-stale

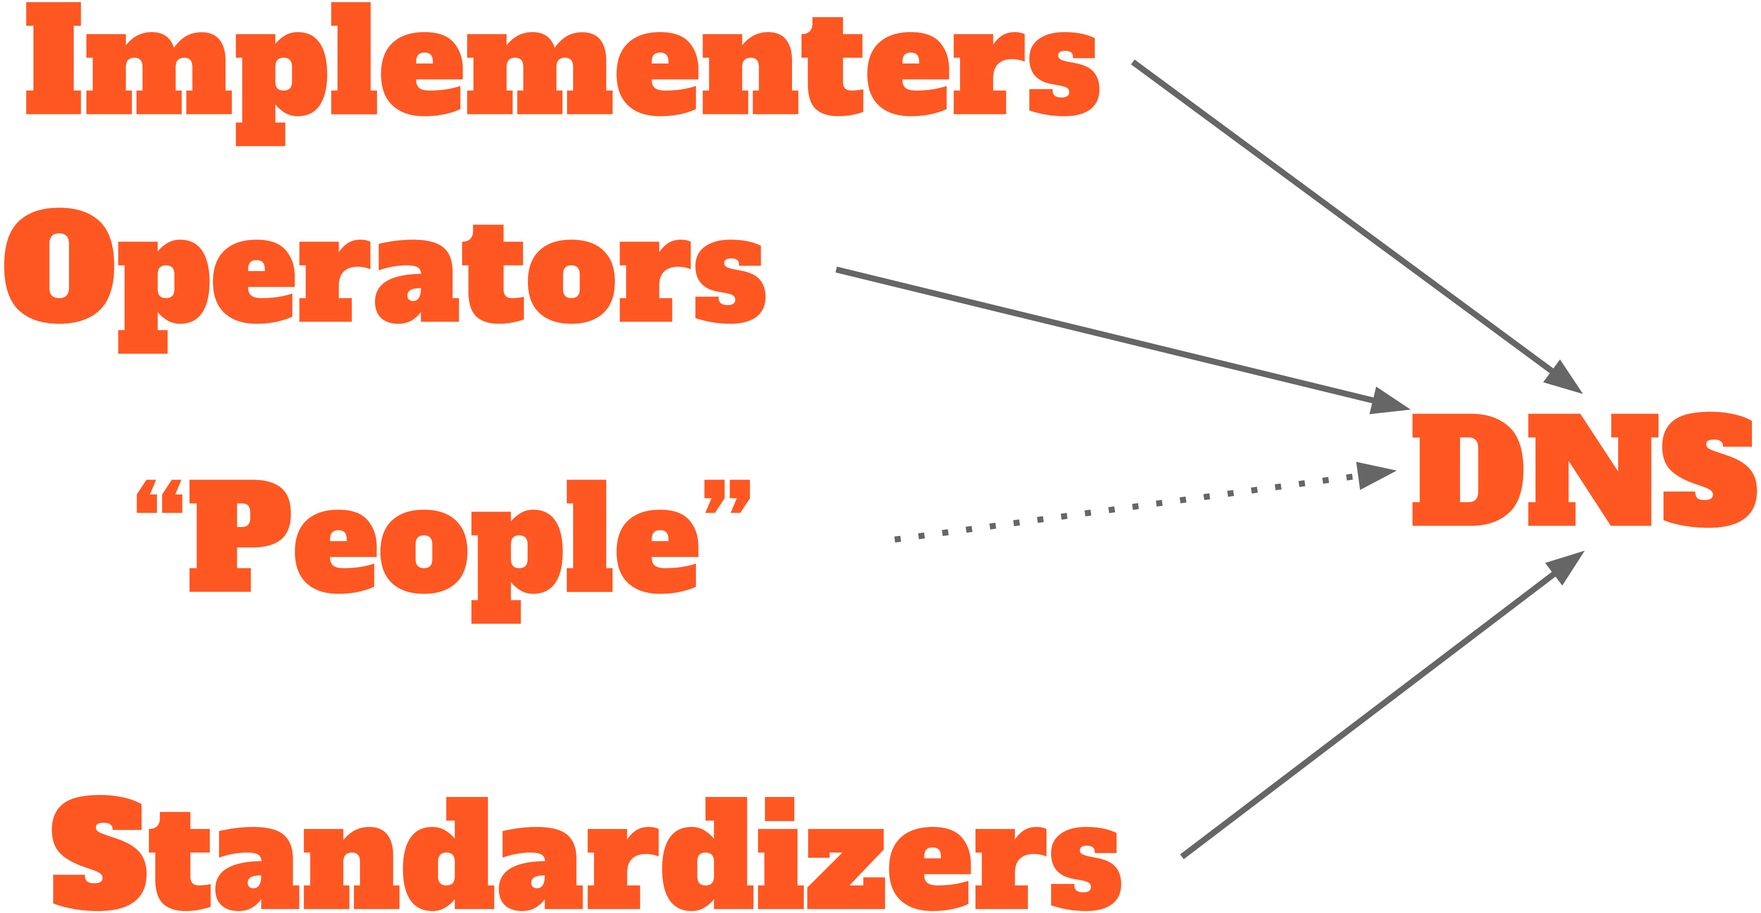

Implementers

Operators

“People”

Standardizers

DNS

Implementors

- We should be AWED by the quality of open source implementations
  a. bind, knot, kresd, unbound, NSD, there is SO much great software out there
  b. Perhaps one of the best served protocols on the internet!
- Very gifted programmers, among the smartest in the world
- So far, they (we) have been able to implement most things, eventually correctly
- **For us, saying “no, this is too complicated” is very hard**
  a. Pride
  b. “One of the other implementations will do it”
  c. Always fun to work on new challenges
- We do not have well developed “product management”
  a. Any individual committer can decide “cool feature, let’s do it”

Operators

- Commercial access provider operators are
  1. On call 24/7
  2. Being measured solely on availability, performance
  3. May actually be penalized by their governments if they do the right thing
- Typically resource constrained, understaffed
- Have no “buy in” from the rest of the access provider to work on privacy enhancing features
  1. In fact…
- Weakly represented in the standards making process
  1. With some notable exceptions
- **Typically turn off anything that could cause problems at 3AM**

ccTLD / root / authoritative operators

- ccTLD/gTLD/root operators are well represented
  a. Significant authoritative hosters (“tens of millions of domains”) are not
- Notably, authoritative implementation of features is rather simpler usually
  a. “Just serve the data”
  b. Almost stateless
- Easy to load balance - even a server that answers 20% of questions will provide good service to the internet
  a. .BE and .NL servers have been down for **hours** or **months** without anyone noticing
- Notably, the one contribution from the operational community, that is widely deployed, did not get standardized (RRL)

Standardizers

- **Like implementers, among the smartest people in the world**
    a. Share enthusiasm for hard challenges
- On a mission to turn the internet into “how things SHOULD be and what the code MUST do to achieve that”
- Try very hard to think of everything
- **Typically not on call 24/7**
- **Undervalue operational trade-offs**
- Simultaneously optimists (on what can be achieved) and pessimists (how folks will mess it up unless everything pinned down by standard)

Unexpected interaction of features

- DNAME needs DNSSEC special casing
- EDNS Client Subnet leads to zero cache hit rates
  - And associated, non-standardized, workarounds
- Qname minimization turns out to need a ton of probing
- Outbound TLS usage leads to ton of probing
- DNS cookies lead to ton of probing
- Multiple answers/qtypes lead to ton of probing
- **Most features are not orthogonal to the other features**
  - Especially on the resolver side!

Net result

- Push to enhance DNS further and further from standards community
- Little push-back from implementation community
- Commercial operational community very weakly represented “and they don’t want anything new anyhow”
- **Proposed features that SHOULD make the internet better are very likely to be accepted and implemented**
  - With little open discussion on how hard this will be
- Given relatively constant base of developers, increase in feature volume will mean **decrease in quality**
- Eventually, glut of features will cause statis

Proposal

- Think long and hard who wants a feature and who would benefit
- Conversely, who would bear the costs?
    - In terms of development, operational stability/quality impact, downstream complexity
- Involve development community more comprehensively
    - It is not enough for ‘bert’ or ‘wouter’ or ‘ondrej’ to feel that it could in theory be done
- Developer community develop some spine & “product management”
- Work ever harder to involve operational community
    - Not easy for them to come to IETF and similar venues
    - Not authorized to speak
    - No travel budget
- Thank you.