

DNS Core Census Data Dictionary

22 December 2021

DNS Core Census Version 0.1.0

[Definition of Tables and Files, Version 0.1.0](#)

[Tables and File Names](#)

[Conventions](#)

[Column Definitions](#)

[Common Columns \(Appearing in Multiple Tables\)](#)

[INDEX](#)

[CENSUS_START](#)

[CENSUS_END](#)

[CENSUS_SOURCES](#)

[Table ADDRESSES](#)

[INDEX](#)

[CENSUS_START](#)

[CENSUS_END](#)

[ADDRESS](#)

[CENSUS_SOURCES](#)

[ADDRESS_FAMILY](#)

[ROUTE_ORIGINS](#)

[REGISTERED_USE_SET](#)

[GLUE_USE_SET](#)

[AUTH_USE_SET](#)

[Table AUTONOMOUS_SYSTEM_NUMBERS](#)

[INDEX](#)

[CENSUS_START](#)

[CENSUS_END](#)

[ASN](#)

[CENSUS_SOURCES](#)

[ASN_NAME](#)

[Table DNSKEY_RECORDS](#)

[INDEX](#)

[CENSUS_START](#)

[CENSUS_END](#)

[OWNER](#)

[FLAGS](#)

[PROTOCOL](#)

[ALGORITHM](#)

[KEY_BASE64](#)

[KEY_TAG](#)

[Table DS_RECORDS](#)

[INDEX](#)

[CENSUS_START](#)

[CENSUS_END](#)

[OWNER](#)

[ALGORITHM](#)

[DIGEST_TYPE](#)

[KEY_TAG](#)

[DIGEST](#)

[Table IDN_PRACTICES](#)

[INDEX](#)

[CENSUS_START](#)

[CENSUS_END](#)

[TOPIC](#)

[KIND](#)

[VERSION](#)

[DATE](#)

[URL](#)

[Table NAMESERVERS](#)

[INDEX](#)

[CENSUS_START](#)

[CENSUS_END](#)

[NAME](#)

[CENSUS_SOURCES](#)

[REGISTERED_IPV4ADDRESS_SET](#)

[REGISTERED_IPV6ADDRESS_SET](#)

[REGISTERED_USE_SET](#)

[GLUE_IPV4ADDRESS_SET](#)

[GLUE_IPV6ADDRESS_SET](#)

[AUTH_IPV4ADDRESS_SET](#)

[AUTH_IPV4ADDRESS_SET_SIGNATURES](#)

[AUTH_IPV6ADDRESS_SET](#)

[AUTH_IPV6ADDRESS_SET_SIGNATURES](#)

[CUT_USE_SET](#)

[AUTH_USE_SET](#)

[Table ROUTE_ORIGINS](#)

[INDEX](#)

[CENSUS_START](#)

[CENSUS_END](#)

[IDENTIFIER](#)
[CENSUS_SOURCES](#)
[BGP_PREFIX](#)
[ASN](#)
[ROA_STATUS](#)

[Table RRSIG_RECORDS](#)

[INDEX](#)
[CENSUS_START](#)
[CENSUS_END](#)
[OWNER](#)
[TYPE_COVERED](#)
[ALGORITHM](#)
[LABELS](#)
[ORIGINAL_TTL](#)
[EXPIRATION_STRING](#)
[INCEPTION_STRING](#)
[KEY_TAG](#)
[SIGNER_OBJ](#)
[SIGNATURE_BASE64](#)

[Table ZONES](#)

[INDEX](#)
[CENSUS_START](#)
[CENSUS_END](#)
[ALABEL](#)
[CENSUS_SOURCES](#)
[ULABEL](#)
[CENSUS_CATEGORY](#)
[CENSUS_JURISDICTION](#)
[SOA_RR_RNAME](#)
[SOA_SET_SIGNATURES](#)
[SPECIAL_USE_LISTED](#)
[ROOT_DB_ATSIGNID](#)
[ROOT_DB_REGISTRATION_DATE](#)
[ROOT_DB_STATUS](#)
[ROOT_DB_META_CLASS](#)
[ROOT_DB_META_STATUS](#)
[ROOT_DB_META_ELIGIBILITY](#)
[ROOT_DB_META_ENTRY_INTO_SERVICE](#)
[ROOT_DB_META_ICANN_REGISTRAR_SYSTEM](#)
[ROOT_DB_META_ISO3166_CODE](#)
[ROOT_DB_META_USAGE](#)

[ROOT_DB_REGISTRANT_ORGANIZATION](#)
[ROOT_DB_REGISTRANT_ORGANIZATION_JURISDICTION](#)
[ROOT_DB_ADMIN_ORGANIZATION](#)
[ROOT_DB_ADMIN_ORGANIZATION_JURISDICTION](#)
[ROOT_DB_TECH_ORGANIZATION](#)
[ROOT_DB_TECH_ORGANIZATION_JURISDICTION](#)
[GTLD_CONTRACTS_GTLD](#)
[GTLD_CONTRACTS_ULABEL](#)
[GTLD_CONTRACTS_SPECIFICATION13](#)
[GTLD_CONTRACTS_APPLICATION_ID](#)
[GTLD_CONTRACTS_CONTRACT_SIGNATURE_DATE](#)
[GTLD_CONTRACTS_ZONE_DELEGATION_DATE](#)
[GTLD_CONTRACTS_CONTRACT_HAS_BEEN_TERMINATED](#)
[GTLD_CONTRACTS_REMOVAL_DATE](#)
[GTLD_CONTRACTS_CLASS_DOMAIN_NAME_LIST](#)
[GTLD_CONTRACTS_3PLUSLD_REGISTRATIONS](#)
[GTLD_CONTRACTS_REGISTRY_OPERATOR](#)
[GTLD_CONTRACTS_REGISTRY_OPERATOR_JURISDICTION](#)
[IDN_CCTLD_FAST_TRACK_STRING_IN_ENGLISH](#)
[IDN_CCTLD_FAST_TRACK_STRING_LANGUAGE](#)
[IDN_CCTLD_FAST_TRACK_STRING_SCRIPT](#)
[IDN_CCTLD_FAST_TRACK_VARIANT_STRINGS](#)
[IDN_CCTLD_FAST_TRACK_OPERATOR_ALSO_CCTLD](#)
[IDN_CCTLD_FAST_TRACK_STATUS](#)
[AGREEMENTS_TERMINATION_ROOT_ZONE_REMOVAL](#)
[AGREEMENTS_TERMINATION_ROOT_ZONE_REMOVAL_DATE](#)
[AGREEMENTS_TERMINATION_AGREEMENT_TERMINATED](#)
[AGREEMENTS_TERMINATION_AGREEMENT_TERMINATED_DATE](#)
[AGREEMENTS_TERMINATION_TERMINATION_IN_PROGRESS](#)
[AGREEMENTS_TERMINATION_LOG](#)
[IANA_IDN_TABLES_REGISTRY](#)
[IANA_IDN_TABLES_LISTING](#)
[UN_M49_GLOBAL_CODE](#)
[UN_M49_GLOBAL_NAME](#)
[UN_M49_REGION_CODE](#)
[UN_M49_REGION_NAME](#)
[UN_M49_SUB_REGION_CODE](#)
[UN_M49_SUB_REGION_NAME](#)
[UN_M49_INTERMEDIATE_REGION_CODE](#)
[UN_M49_INTERMEDIATE_REGION_NAME](#)
[UN_M49_M49_CODE](#)

[UN_M49_ISO_ALPHA3_CODE](#)
[ICANN GEOGRAPHIC REGION](#)
[REGISTERED_NAMESERVER_SET](#)
[REGISTERED_DS_SET](#)
[AUTH_DS_SET](#)
[AUTH_DS_SET_SIGNATURES](#)
[AUTH_NAMESERVER_SET](#)
[AUTH_NAMESERVER_SET_SIGNATURES](#)
[AUTH_DNSKEY_SET](#)
[AUTH_DNSKEY_SET_SIGNATURES](#)
[CUT_NAMESERVER_SET](#)
[NON_DELEGATED_NAMES](#)
[NON_DELEGATED_CNAMES](#)
[NON_DELEGATED_DNAMES](#)
[DELEGATED_SUB_TLDS](#)
[SIZE_DATA_NS_OWNERS](#)
[SIZE_DATA_DS_OWNERS](#)
[SIZE_DATA_NSEC_OWNERS](#)
[SIZE_DATA_NSEC3_OWNERS](#)
[SIZE_DATA_IDN_NS_OWNERS](#)
[SIZE_DATA_IDN_DS_OWNERS](#)

[Contact Information](#)

Definition of Tables and Files, Version 0.1.0

The DNS Core Census V.0.1.0 is made available in three forms. Two forms are via the HTTP in a series of tables, one in CSV and the other in JSON. The third form consists of tables in a relational database. The contents of each form are equivalent, but the naming of the tables and naming of the files differs. Accessibility of each form may vary over time.

A previous version of the DNS Core Census was made public via a web URL, it is called version 0.0.2. Subsequent developmental versions, 0.0.3 and 0.0.4 were not made public.

The most significant difference between versions 0.1.x and 0.0.x is a shift from a data-type rich, hierarchical representation in JSON to a table-based, flattened structure compatible with SQL-like database (interfaces), using CSV or JSON to marshal the tables. Whereas the 0.0.x is written in 3 JSON files published each (daily) run, 0.1.x will feature 9 tables aggregating all census runs per month.

To conserve disk space, the version 0.1.0 available via a website will appear only as compressed CSV files.

The collection process is different as well, more resilient to data source failures and includes more data sources than in version 0.0.2 (those having been trialed in versions 0.0.3 and 0.0.4).

Tables and File Names

There are 9 kinds of tables that work together to document each run of the census.

Each table in a relational database, regardless of kind, contains one month's worth of census runs. Each CSV file and each JSON file will contain data for just one census run. (During development and testing, there may be more than one run per day.)

The naming convention of the tables is "DNS_CORE_CENSUS_{kind}_{year/month}_TABLE".

The naming convention of the file_names is "DNS_CORE_CENSUS_{kind}_{time}_TABLE.csv".

"kind" can be one of:

- ADDRESSES,
- AUTONOMOUS_SYSTEM_NUMBERS,
- DNSKEY_RECORDS,
- DS_RECORDS,
- IDN_PRACTICES,
- NAMESERVERS,
- ROUTE_ORIGINS,
- RRSIG_RECORDS,
- ZONES

"year/month" is written as YYYY_MM for year and month (January='01'). Note that across datafields, a hyphen is used as a separator for data elements, but for the table's file name an underscore is used as a separator.

"time" is written as YYYY_MM_DD_HHMMSS for year, month (January='01'), day, hour, minute, second. The time is expressed in UTC, it is the time the primary census data collection run began.

All fields are represented as text in CSV files.

The primary (starting point) table is the ZONES table.

Cells with multiple values delimit the values with the pipe character (|).

Conventions

Domain names are presented in lower case. Software mistakes may occur, it is best not to rely on this assumption. Values used to index the ZONES and NAMESERVERS tables will use lower case for domain names.

Much of the census information is taken directly from other sources, some is a summary of the other sources, and a small amount is derived across sources. Where there is no source information for an item, the census will insert a placeholder value. The placeholder values, beginning around 2 September 2021 runs of the census will start and end with an underscore, such as '_not applicable_', '_unused field_', '_not retrieved_', '_unset_'. A blank ("") value indicates that information was sought but the answer was an empty set of information. For example, if a zone does not have DNSKEY resource records, the list of key records will be a blank string.

All of the CSV file forms of the cells are presented in quoted UTF-8 encoded strings, each cell separated by a comma. None of the values ought to be interpreted as a null value, specifically, "NA" will represent a value (for Namibia) in the ISO 3166-1 Alpha-2 table, it should not be interpreted as "not a number" or another null value.

All table names and column names are in uppercase and use underscores as separators. This convention arises from a back end database used for the tables, unrelated to the CSV or JSON forms. There was a desire to echo the original column names in the census, but there are a few cases where the original column name was syntactically invalid for the back end database. In addition to this, all references to "country" (or other descriptions of a region) have been replaced by jurisdiction to more accurately reflect categorizations of elements. The historical term "ccTLD" is retained but is referred to as a jurisdictional-based TLD (not to be confused with geographical gTLDs).

The IP addresses used in examples are reserved for documentation per *IPv4 Address Blocks Reserved for Documentation* (RFC 5737) and *IPv6 Address Prefix Reserved for Documentation* (RFC 3849).

Autonomous System Numbers in examples are reserved for documentation per *Autonomous System (AS) Number Reservation for Documentation Use* (RFC 5398).

Domain names used in examples are reserved for documentation per *Special-Use Domain Names* (RFC 6761). By virtue of the domain names being in the Special-Use Domain Name registry, the domain names are also part of the data in the tables, but examples using the names do not appear in the tables.

There is one description where data in the table is used as an example, to exercise the specific parsing of the data's cell. The specific cell is historic (won't change) and is not identified otherwise.

Column Definitions

Regardless of whether the data is presented in CSV files or in database Tables, the columns are the same. In a CSV file, the first row contains the names of the columns as they appear in this document.

Common Columns (Appearing in Multiple Tables)

The following columns appear in multiple tables, sharing the same formatting and meaning. To ensure a consistent description, the details of these columns are here. For each subsequent table description, a note to refer to this section is added.

INDEX

A convenience handle for the table, it is a number from 0 upward for each run of the census. This column is more significant in tables that are referenced by other tables, especially when there is no other truly unique key column.

CENSUS_START

The date and time (UTC) when the census data collection began. This field's significance is in identifying all records, across all tables, that pertain to the same run of the census.

The format of the string is `YYYY-MM-DD-HHMMSS`, (Year-Month-Day-HourMinuteSecond or `'%Y-%m-%d-%H%M%S'`)

CENSUS_END

The date and time when the census data collection ended. This field's significance is in identifying all records, across all tables, that pertain to the same run of the census.

The format of the string is `YYYY-MM-DD-HHMMSS`, (Year-Month-Day-HourMinuteSecond or `'%Y-%m-%d-%H%M%S'`)

CENSUS_SOURCES

This indicates where information was found for the row in which the cell appears. Each field (usually) has three parts, separated by colons, indicating a source and providing some relevant data. Multiple sources are delimited by pipe characters.

The values in `CENSUS_SOURCES` have the following meanings:

Value	Meaning
-------	---------

AXFR:<domain>:	Information came from a zone transfer of <domain>
DNS query	Note the space between DNS and query. A result in an oversight in code, fixed starting on 2021-09-02. Only is seen in the NAMESERVERS table. Where this is seen, substitute with DNS_query:<name>:A and DNS_query:<name>:AAAA, where <name> is the value of the NAME column
DNS_query:<domain>:<qtype>	A DNS query was used, for <domain>'s <qtype> record set
FastTrack:none:	Information came from the IDN Fast Track String Evaluation table (https://www.icann.org/resources/pages/string-evaluation-completion-2014-02-19-en)
ICANNReg:none:	Information came from the ICANN Geographic Regions table (https://meetings.icann.org/en/regions)
IDNTables:none:	Information came from Repository of IDN Practices (https://www.iana.org/domains/idn-tables)
IP2ASN:_empty_:	Information came from Team Cymru's IP to ASN mapping service (https://team-cymru.com/community-services/ip-asn-mapping/)
None	Information that is hardcoded, namely the distribution sites for ICANN operated zones and NRO operated zones. These servers are not listed in an NS resource record set yet are publically accessible. Information on the ICANN distribution servers is available here: https://www.dns.icann.org/services/axfr/
RegTerm:none:	Information came from the Registry Agreement Termination Information Page (https://www.icann.org/resources/pages/gtld-r-registry-agreement-termination-2015-10-09-en)
SpecialUseRegistry:None:listed	Information came from the Special Use Domain Name registry (https://www.iana.org/assignments/special-us)

	e-domain-names/special-use-domain-names.xhtml#special-use-domain)
Sub_tld_zone:<domain>:	Information was obtained while determining whether a potential sub-domain was a sub_tld_zone for <domain>
iana_root_db:::	Information came from an XML version of the root zone database, discontinued in October 2021 (https://www.iana.org/exports/root-1.1.xml)
iana_root_db_whois:::	Information came via a web scrape of root zone database (https://www.iana.org/domains/root/db)
icann_contract_status_report_table:None:listed	Information came from a JSON formatted resource documenting ICANN gTLD contracts (https://www.icann.org/resources/registries/gtlds/v2/gtlds.json)

Table ADDRESSES

This table aggregates information about an address, one row per address per census run. This table refers to the ROUTE ORIGINS TABLE.

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

ADDRESS

An Internet address appearing in some role related to the DNS Core. For the most part, these are addresses of name servers but may be addresses in stale glue records.

IPv4 addresses appear without leading zeros, IPv6 addresses appear without leading 0's and make use of the '::' shorthand. Despite these being rendered as strings, it is recommended to store them internally using available IP address packages when trying to compare values.

CENSUS_SOURCES

See the subsection in Common Columns.

ADDRESS_FAMILY

Generally IPv4 or IPv6 as appropriate for the address. It is possible other values may appear here, if there is ever another address family.

ROUTE_ORIGINS

A pipe delimited set of a BGP prefix in CIDR notation and Autonomous System Number separated by a ":". This documents the advertised sources for the address. An example entry would be "192.0.2.0/24:64500".

The value(s) in this field can be used to look into the ROUTE_ORIGINS table for further information.

This field may be a blank string.

REGISTERED_USE_SET

A pipe-delimited list of name servers where the address is given for registration. I.e., what is seen in a registry database report (like whois or rdap) or in the zone cuts within a zone transfer.

This field may be '_not applicable_' or '_unset_'.

GLUE_USE_SET

A pipe-delimited list of source zone and name servers, separated by a :. The source zone refers to the zone publishing the glue. An example entry might be: "example.:ns1.dns.example." There are rare cases where glue might appear in more than one zone, usually in a parent zone and grand-parent zone.

This field may be '_unset_'

AUTH_USE_SET

A pipe-delimited list of name servers where the address is authoritatively used. This can be seen in DNS queries (with the authority bit set) or in the apex of a zone transfer.

This field may be '_unset_'

Table AUTONOMOUS_SYSTEM_NUMBERS

This table maps an Autonomous System Number to an operator's name as given by Team Cymru's IP to ASN mapping service.

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

ASN

An Autonomous System Number, appearing as an integer (in string form).

CENSUS_SOURCES

See the subsection in Common Columns.

ASN_NAME

The string that is given by Team Cymru's IP to ASN mapping service for the value in column ASN.

Table DNSKEY_RECORDS

This table contains the details of a DNSKEY resource record, one row per item (per census run). The INDEX is used for reference (CENSUS_START,CENSUS_END,INDEX ought to completely identify a single record). All of the information is obtained via queries for DNSKEY resource records or from a zone transfer.

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

OWNER

The domain name that owns the DNSKEY resource record. The name appears as a string in lower case, but it is best to use a specialized DNS name comparison library.

FLAGS

A numerical representation of the flags field, in base 10. '256' is an ordinary ZSK and '257' is an ordinary 'KSK'. Other values are possible, including flags indicating revocation.

PROTOCOL

This value is always 3.

ALGORITHM

The DNSSEC Security Algorithm for the DNSKEY. See IANA's registry for the possible values, <https://www.iana.org/assignments/dns-sec-alg-numbers/dns-sec-alg-numbers.xhtml#dns-sec-alg-numbers-1>. This value should be the mnemonic in that registry, but may be the numeric value if the package used to perform this mapping is not kept up to date.

KEY_BASE64

The Base64 encoding (into a string) of the "raw bits" of the public key. This field is case sensitive.

KEY_TAG

Value (as a string) of the DNSSEC defined key_tag for the DNSKEY. This value is not pre-padded with '0' to 5 digits for small values. (I.e., 100 is represented as '100'.)

Table DS_RECORDS

This table contains the details of a DS resource record, one row per item (per census run). The INDEX is used for reference (CENSUS_START,CENSUS_END,INDEX ought to completely identify a single record). All of the information is obtained via queries for DS resource records or from a zone transfer.

In early (up through November 2021) tables, there were many duplicated rows. These rows are not referenced from the ZONES table but remain to preserve the index numbering used. For this reason, there will be many rows with '_lame row_' values, indicating these once duplicated rows.

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

OWNER

The domain name that owns the DS resource record. The name appears as a string in lower case, but it is best to use a specialized DNS name comparison library.

This cell may have the value '_lame row_'.

ALGORITHM

The DNSSEC Security Algorithm for the DNSKEY referenced in the DS resource record. See the *DNS Security Algorithm Numbers* registry for the possible values (<https://www.iana.org/assignments/dns-sec-alg-numbers/dns-sec-alg-numbers.xhtml#dns-sec-alg-numbers-1>). This value should be the mnemonic in that registry, but may be the numeric value if the package used to perform this mapping is not kept up to date.

This cell may have the value '_lame row_'.

DIGEST_TYPE

The hash algorithm used for the DS resource record, mnemonics defined in the *DNSSEC Delegation Signer (DS) Resource Record (RR) Type Digest Algorithms* registry (<https://www.iana.org/assignments/ds-rr-types/ds-rr-types.xhtml#ds-rr-types-1>)

This cell may have the value '_lame row_'.

KEY_TAG

Value (as a string) of the DNSSEC defined key_tag for the DNSKEY. This value is not pre-padded with '0' to 5 digits for small values. (I.e., 100 is represented as '100'.)

This cell may have the value '_lame row_'.

DIGEST

The hexadecimal-in-string representation of the hashed value referencing a DNSKEY resource record.

This cell may have the value '_lame row_'.

Table IDN_PRACTICES

This table is a summary of the IDN tables for each TLD as registered with IANA and documented in the *Repository of IDN Practices* (<https://www.iana.org/domains/idn-tables>). This table does not identify the TLD for which the practice statement. In the ZONES table, each TLD with registered IDN practices will have a list of references to this table.

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

TOPIC

The subject of the table - typically the name of a language or written script. The value here comes from the visible portion of the IANA web page's Table column, minus the TLD name.

KIND

The values "language" and "script" indicate what kind of practice is described, from the IANA web page's Type column.

VERSION

A version tag, either an integer or dotted decimal, from the IANA web page's Version column.

DATE

A YYYY-MM-DD date string, from the IANA web page's Date column. ("%Y-%m-%d")

URL

The unseen (when viewed through a web browser) link to the IDN Practice statement. The URL is "under" the entries in the IANA web page's Table column.

Table NAMESERVERS

This table contains the details of each nameserver involved in the census.

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

NAME

The domain name of the nameserver. The name appears as a string in lower case, but it is best to use a specialized DNS name comparison library.

CENSUS_SOURCES

See the subsection in Common Columns.

REGISTERED_IPV4ADDRESS_SET

A pipe-delimited list of IPv4 addresses identified as belonging to the nameserver, usually in the form of A resource records, found in zone transfers (when authoritative) or in a registry database. Addresses have no leading 0's.

This cell may be blank or may have the values '_not applicable_' or '_unset_'.

REGISTERED_IPV6ADDRESS_SET

A pipe-delimited list of IPv6 addresses identified as belonging to the nameserver, usually in the form of A resource records, found in zone transfers (when authoritative) or in a registry database. Addresses have no leading 0's and use the '::' shorthand.

This cell may be blank or may have the values '_not applicable_' or '_unset_'.

REGISTERED_USE_SET

A pipe-delimited list of zones (syntactically domain names) that list the name server as authoritative, as seen via a registry database. (Only the root zone registry is involved.) This includes names meant to appear in cut point NS resource records.

This cell may have the values '_not applicable_' or '_unset_'.

GLUE_IPV4ADDRESS_SET

A pipe-delimited list of a source zone and IPv4 address pairs, separated by a :. The source zone refers to the zone publishing the glue. An example entry might be: "example.:192.0.2.10" There are rare cases where glue might appear in more than one zone, usually in a parent zone and grand-parent zone.

This cell may be blank or may have the value '_unset_'.

GLUE_IPV6ADDRESS_SET

A pipe-delimited list of a source zone and IPv6 address pairs, separated by a :. The source zone refers to the zone publishing the glue. An example entry might be: "example.:2001:DB8::53" There are rare cases where glue might appear in more than one zone, usually in a parent zone and grand-parent zone.

This cell may be blank or may have the value '_unset_'.

AUTH_IPV4ADDRESS_SET

A pipe-delimited list of IPv4 addresses seen in authoritative A resource records owned by the nameserver's domain name.

This cell may be blank or may have the value '_unset_'.

AUTH_IPV4ADDRESS_SET_SIGNATURES

A pipe-delimited list of index values into the RRSIG_RECORDS table for relevant DNSSEC signatures. Use this index as well as the CENSUS_START and CENSUS_END to find the intended records.

This cell may be blank or may have the value '_unset_'.

AUTH_IPV6ADDRESS_SET

A pipe-delimited list of IPv6 addresses seen in authoritative AAAA resource records owned by the nameserver's domain name.

This cell may be blank or may have the value '_unset_'.

AUTH_IPV6ADDRESS_SET_SIGNATURES

A pipe-delimited list of index values into the RRSIG_RECORDS table for relevant DNSSEC signatures. Use this index as well as the CENSUS_START and CENSUS_END to find the intended records.

This cell may be blank or may have the value '_unset_'.

CUT_USE_SET

A pipe-delimited list of zones (syntactically domain names) where the name server is listed in a NS resource record at the parent-side of a zone cut.

This cell may have the value '_unset_'.

AUTH_USE_SET

A pipe-delimited list of zones (syntactically domain names) where the name server is listed in a NS resource record owned and published at zone's apex.

This cell may have the value '_unset_'.

Table ROUTE_ORIGINS

This table contains details about route origins as obtained from the Team Cymru IP-to-ASN mapping service and the RIPE RPKI-Validator. The list of validated ROAs is obtained from this URL, <https://rpki-validator.ripe.net/api/objects/validated>. (Discontinued in October 2021 and replaced by a local instance of a validator.)

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

IDENTIFIER

A pipe delimited set of a BGP prefix in CIDR notation and Autonomous System Number separated by a "|". This documents the advertised sources for the address. An example entry would be "192.0.2.0/24:64500". (This is the same description as the ADDRESSES_TABLE column ROUTE_ORIGINS.)

This field is the BGP prefix and Origin Autonomous System Number for the record, used to match the row with a particular route origin.

CENSUS_SOURCES

See the subsection in Common Columns.

BGP_PREFIX

This is the BGP prefix, written in CIDR notation, for the origin. For example: 192.0.2.0/24.

ASN

This is the Autonomous System Number for the origin, an integer (represented a string).

ROA_STATUS

This is "True" or "False" representing whether the route origin has a validated ROA according to the RIPE RPKI validator.

Table RRSIG_RECORDS

This table contains the details of a RRSIG resource record, one row per item (per census run). The INDEX is used for reference (CENSUS_START,CENSUS_END,INDEX ought to completely identify a single record). All of the information is obtained via queries for resource records of the covered type or from a zone transfer.

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

OWNER

The domain name that owns the RRSIG resource record. The name appears as a string in lower case, but it is best to use a specialized DNS name comparison library.

TYPE_COVERED

The type covered field, for the range of valid values, see the Resource Record (RR) TYPEs registry

(<https://www.iana.org/assignments/dns-parameters/dns-parameters.xml#dns-parameters-4>) or *Handling of Unknown DNS Resource Record (RR) Types* (RFC 3597). Whether a type will be presented as an unknown DNS resource record will depend on the DNS libraries in use when the table is written.

ALGORITHM

The DNSSEC Security Algorithm of the key used to generate the signature. Values are from the *DNS Security Algorithm Numbers*. This value should be the mnemonic in that registry, but may be the numeric value if the package used to perform this mapping is not kept up to date.

LABELS

The count (numeric string) of labels involved in matching the signature. This is used to detect synthesized records (wild cards). The count does not include the root label at the end of a fully qualified domain name, per the DNSSEC definitional documents.

ORIGINAL_TTL

The original TTL of the covered resource record set, represented in seconds.

EXPIRATION_STRING

The absolute timestamp of the expiration of the validity of the signature. The numeric string is presented as it appears in the RRSIG resource record, `YYYYMMDDHHmmSS` in the GMT+00 zone zone. ('%Y%m%d%H%M%S', no separators are used.)

INCEPTION_STRING

The absolute timestamp of the inception of the validity of the signature. The numeric string is presented as it appears in the RRSIG resource record, `YYYYMMDDHHmmSS` in the GMT+00 zone zone. ('%Y%m%d%H%M%S', no separators are used.)

KEY_TAG

Value (as a string) of the DNSSEC defined `key_tag` for the DNSKEY. This value is not pre-padded with '0' to 5 digits for small values. (I.e., 100 is represented as '100'.)

SIGNER_OBJ

The domain name that owns the DNSKEY resource record to be used to verify the signature. The name appears as a string in lower case, but it is best to use a specialized DNS name comparison library. The owner of the key along with the KEY_TAG (and DNSSEC Security Algorithm) will usually identify a single key for validation but uniqueness is not guaranteed.

SIGNATURE_BASE64

The Base64 encoding of the signature's binary form. This field is case sensitive.

Table ZONES

This table contains the details of each zone involved in the census. This is the primary table of the schema.

INDEX

See the subsection in Common Columns.

CENSUS_START

See the subsection in Common Columns.

CENSUS_END

See the subsection in Common Columns.

ALABEL

The domain name of the zone. (For IDN zones, this is the punycode representation.) The name appears as a string in lower case, but it is best to use a specialized DNS name comparison library.

CENSUS_SOURCES

See the subsection in Common Columns.

ULABEL

For zones whose names are not IDN, this field is the same as the ALABEL. For IDN zones, this field is the Unicode form of the name. When placed into a file, it is encoded in UTF-8.

CENSUS_CATEGORY

This is a category assigned during the census process. The values generated are in the following table.

Value	Meaning
root	The root zone (alone)
infrastructure	ARPA zone and zones that exist to provide structure to the namespace
enum	Zones used for ENUM (enum.arpa. alone)
revMap	Reverse Map zones
test	Zones purposed for testing, such as the IDN TLD tests (note not IDN ccTLD) conducted before any IDN TLDs were introduced
special_use	A domain name (not necessarily a zone, but included as such) listed in the Special Use Domain Name registry
gTLD	TLDs that are not ccTLDs, including commercial generic TLDs, sponsored TLDs, other legacy TLDs (like mil, gov., edu.) that were not otherwise given a category
sub_gTLD	A zone below a gTLD that (heuristically) is operated as part of the gTLD
ccTLD	From Wikipedia.org: "an Internet top-level domain generally used or reserved for a country, sovereign state, or dependent territory identified with a country code."
sub_ccTLD	A zone below a ccTLD that (heuristically) is operated as part of the ccTLD

CENSUS_JURISDICTION

An ISO 3166-1 Alpha 2 code either registered or deduced for a zone. The value 'XA' is used for "world wide", which is the assumed jurisdiction for all non-ccTLD zones and subzones. For sub_ccTLD zones, the jurisdiction is deduced from the ccTLD above it.

SOA_RR_RNAME

The domain name in the RNAME field of the SOA resource record for the zone.

This cell may be blank or may have the value '_not applicable_' or '_unset_'. This may happen when a zone is defined in a table but has no active name servers.

SOA_SET_SIGNATURES

A pipe-delimited list of zero or more index values into the RRSIG_RECORDS table, which holds the details of signatures.

This cell may be blank or may have the value '_unset_'.

SPECIAL_USE_LISTED

"True" or "False" indicating whether the zone's domain name is listed in the Special-Use Domain Names registry.

This cell may have the value '_unset_'.

ROOT_DB_ATSIGNID

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml> through October 2021, later scraped from IANA web pages, if available.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_REGISTRATION_DATE

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_STATUS

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_META_CLASS

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_META_STATUS

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_META_ELIGIBILITY

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_META_ENTRY_INTO_SERVICE

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_META_ICANN_REGISTRAR_SYSTEM

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_META_ISO3166_CODE

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_META_USAGE

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_REGISTRANT_ORGANIZATION

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_REGISTRANT_ORGANIZATION_JURISDICTION

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_ADMIN_ORGANIZATION

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_ADMIN_ORGANIZATION_JURISDICTION

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_TECH_ORGANIZATION

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

ROOT_DB_TECH_ORGANIZATION_JURISDICTION

All columns prefixed with ROOT_DB_ are taken directly from fields in the root zone registry database as rendered in <https://www.iana.org/exports/root-1.1.xml>.

This cell may be blank or have one of the following values: _not applicable_, _not retrieved_, _none listed_, _unset_, or _unused field_.

GTLD_CONTRACTS_GTLD

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_ULABEL

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_SPECIFICATION13

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_APPLICATION_ID

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_CONTRACT_SIGNATURE_DATE

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_ZONE_DELEGATION_DATE

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_CONTRACT_HAS_BEEN_TERMINATED

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_REMOVAL_DATE

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_CLASS_DOMAIN_NAME_LIST

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value _not applicable_.

GTLD_CONTRACTS_3PLUSLD_REGISTRATIONS

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value `_not applicable_`.

GTLD_CONTRACTS_REGISTRY_OPERATOR

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. The remainder of the column name is based on the JSON field selector. Each capitalized letter in the JSON field selector is prepended by an underscore ('_') and all characters are capitalized.

This cell may have the value `_not applicable_`.

GTLD_CONTRACTS_REGISTRY_OPERATOR_JURISDICTION

All columns prefixed with GTLD_CONTRACTS_ are taken directly from <https://www.icann.org/resources/registries/gtlds/v2/gtlds.json>. This field corresponds to the "registryOperatorCountryCode".

This cell may have the value `_not applicable_`.

IDN_CCTLD_FAST_TRACK_STRING_IN_ENGLISH

All columns prefixed with IDN_CCTLD_FAST_TRACK are taken from the IDN ccTLD Fast Track String Evaluation Completion page.

(<https://www.icann.org/resources/pages/string-evaluation-completion-2014-02-19-en>)

Source column labeled "String in English".

This cell may have the value `_not applicable_` or `_unset_`.

IDN_CCTLD_FAST_TRACK_STRING_LANGUAGE

All columns prefixed with IDN_CCTLD_FAST_TRACK are taken from the IDN ccTLD Fast Track String Evaluation Completion page.

(<https://www.icann.org/resources/pages/string-evaluation-completion-2014-02-19-en>)

Source column labeled "Language".

This cell may have the value `_not applicable_` or `_unset_`.

IDN_CCTLD_FAST_TRACK_STRING_SCRIPT

All columns prefixed with IDN_CCTLD_FAST_TRACK are taken from the IDN ccTLD Fast Track String Evaluation Completion page.

(<https://www.icann.org/resources/pages/string-evaluation-completion-2014-02-19-en>)

Source column labeled "Script".

This cell may have the value `_not applicable_` or `_unset_`.

IDN_CCTLD_FAST_TRACK_VARIANT_STRINGS

All columns prefixed with IDN_CCTLD_FAST_TRACK are taken from the IDN ccTLD Fast Track String Evaluation Completion page.

(<https://www.icann.org/resources/pages/string-evaluation-completion-2014-02-19-en>)

Source column labeled "Desired Variant String(s)".

This cell may have the value `_not applicable_` or `_unset_`.

IDN_CCTLD_FAST_TRACK_OPERATOR_ALSO_CCTLD

All columns prefixed with IDN_CCTLD_FAST_TRACK are taken from the IDN ccTLD Fast Track String Evaluation Completion page.

(<https://www.icann.org/resources/pages/string-evaluation-completion-2014-02-19-en>)

Source column labeled "Requester Contact Details", a "True" or "False", if "Existing ccTLD manager" appears in the cell.

This cell may have the value `_not applicable_` or `_unset_`.

IDN_CCTLD_FAST_TRACK_STATUS

All columns prefixed with IDN_CCTLD_FAST_TRACK are taken from the IDN ccTLD Fast Track String Evaluation Completion page.

(<https://www.icann.org/resources/pages/string-evaluation-completion-2014-02-19-en>)

Source column labeled "Status".

This cell may have the value `_not applicable_` or `_unset_`.

AGREEMENTS_TERMINATION_ROOT_ZONE_REMOVAL

All columns prefixed with AGREEMENTS_TERMINATION_ are taken from the Registry Agreement Termination Information Page.

(<https://www.icann.org/resources/pages/gtld-registry-agreement-termination-2015-10-09-en>)

"True" or "False" for zones listed in the table, "True" if the Status column includes "TLD Delegation Revoked".

This cell may have the value `_not applicable_`.

AGREEMENTS_TERMINATION_ROOT_ZONE_REMOVAL_DATE

All columns prefixed with `AGREEMENTS_TERMINATION_` are taken from the Registry Agreement Termination Information Page.

(<https://www.icann.org/resources/pages/gtld-registry-agreement-termination-2015-10-09-en>)

A `YYYY-MM-DD` date if the Status column includes "TLD Delegation Revoked", "None" for all other zones listed in the table.

This cell may have the value `_not applicable_`.

AGREEMENTS_TERMINATION_AGREEMENT_TERMINATED

All columns prefixed with `AGREEMENTS_TERMINATION_` are taken from the Registry Agreement Termination Information Page.

(<https://www.icann.org/resources/pages/gtld-registry-agreement-termination-2015-10-09-en>)

"True" or "False" for zones listed in the table, "True" if the Status column includes "Registry Agreement Terminated".

This cell may have the value `_not applicable_`.

AGREEMENTS_TERMINATION_AGREEMENT_TERMINATED_DATE

All columns prefixed with `AGREEMENTS_TERMINATION_` are taken from the Registry Agreement Termination Information Page.

(<https://www.icann.org/resources/pages/gtld-registry-agreement-termination-2015-10-09-en>)

A `YYYY-MM-DD` date if the Status column includes "TLD Delegation Revoked", "None" for all other zones listed in the table.

This cell may have the value `_not applicable_`.

AGREEMENTS_TERMINATION_TERMINATION_IN_PROGRESS

All columns prefixed with `AGREEMENTS_TERMINATION_` are taken from the Registry Agreement Termination Information Page.

(<https://www.icann.org/resources/pages/gtld-registry-agreement-termination-2015-10-09-en>)

"True" or "False" for zones listed in the table, "True" if the Status column includes any of 'Awaiting Final Determination', 'Implementation of Final Determination' or 'Open for Comment'.

This cell may have the value `_not applicable_`.

AGREEMENTS_TERMINATION_LOG

All columns prefixed with `AGREEMENTS_TERMINATION_` are taken from the Registry Agreement Termination Information Page.

(<https://www.icann.org/resources/pages/gtld-registry-agreement-termination-2015-10-09-en>)

A pipe-delimited list of strings summarizing any termination activity for the TLD. Each summary includes the following fields separated by a colon ':':

- the date of the notice (from Date of Termination Notice) in YYYY-MM-DD format
- the type of termination (from Type of Termination)
- the events listed in the Status column with dates in YYYY-MM-DD format, enclosed in square brackets

This cell may be blank or have the value `_not applicable_`.

An example of a TLD that went through two termination processes, one ending with a withdrawal of the termination notice, the other progressing to a termination of the delegation and the agreement behind it.

The value in the cell, word wrapped:

```
2017-12-06:Termination by Registry Operator Section 4.4(b):TLD
Delegation Revoked[2018-04-06],Registry Agreement
Terminated[2018-06-04]|2017-04-12:Termination by Registry Operator
Section 4.4(b):Withdrawn[None]
```

Viewing the entries in a table.

Date	Type	Events
2017-12-06	Termination by Registry Operator Section 4.4(b)	TLD Delegation Revoked[2018-04-06] Registry Agreement Terminated[2018-06-04]
2017-04-12	Termination by Registry Operator Section 4.4(b)	Withdrawn[None]

IANA_IDN_TABLES_REGISTRY

All columns prefixed with IANA_IDN_TABLES_ are taken from the Repository of IDN Practices (<https://www.iana.org/domains/idn-tables>).

The string in this column is taken from the Registry column in the repository's web page.

This cell may have the value `_not applicable_`.

IANA_IDN_TABLES_LISTING

All columns prefixed with IANA_IDN_TABLES_ are taken from the Repository of IDN Practices (<https://www.iana.org/domains/idn-tables>).

A pipe-delimited list of index values for the IDN_PRACTICES table, each corresponding to an entry in the repository.

This cell may have the value `_not applicable_`.

UN_M49_GLOBAL_CODE

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may have the value `_not applicable_`.

UN_M49_GLOBAL_NAME

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may have the value `_not applicable_`.

UN_M49_REGION_CODE

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may be blank or may have the value _not applicable_.

UN_M49_REGION_NAME

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may be blank or may have the value _not applicable_.

UN_M49_SUB_REGION_CODE

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may be blank or may have the value _not applicable_.

UN_M49_SUB_REGION_NAME

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may be blank or may have the value _not applicable_.

UN_M49_INTERMEDIATE_REGION_CODE

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may be blank or may have the value _not applicable_.

UN_M49_INTERMEDIATE_REGION_NAME

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may be blank or may have the value _not applicable_.

UN_M49_M49_CODE

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may have the value _not applicable_.

UN_M49_ISO_ALPHA3_CODE

All columns prefixed with UN_M49_ are taken from the United Nation's Standard country or area codes for statistical use (M49). The table is accessed from <https://unstats.un.org/unsd/methodology/m49/overview/>.

For each of the columns beginning with UN_M49_, the remainder of the name indicates the source column in the UN table.

This cell may have the value _not applicable_.

ICANN_GEOGRAPHIC_REGION

This column lists the ICANN geographic region associated with a TLD (ccTLD only) from the ICANN Geographic Regions (<https://meetings.icann.org/en/regions>) table. The regions are 'AF', 'AP', 'EUR', 'LAC', 'NA' for zones listed in the table.

This cell may have the value _not applicable_.

REGISTERED_NAMESERVER_SET

A pipe-delimited list of domain names appearing as registered from reading a registry database. Domain names should appear in lower case, alabel form, use a DNS domain name library to perform comparisons. Each domain name can be used as an index into the NAMESERVERS table, using the OWNER column. This set does not arrive with a DNSSEC signature.

This cell may be blank or may have one of the values `_not applicable_` or `_unset_`.

REGISTERED_DS_SET

A pipe-delimited list of index values into the DS_RECORDS table. These are the DS resource records that appear in the parent zone's registry database. This set does not arrive with a DNSSEC signature.

This cell may be blank or may have one of the values `_not applicable_` or `_unset_`.

AUTH_DS_SET

A pipe-delimited list of index values into the DS_RECORDS table. These are the DS resource records that appear in the parent zone, obtained either by zone transfer or direct query. DS resource record sets ought to be accompanied by a DNSSEC signature.

This cell may be blank or may have one of the values `_not applicable_` or `_unset_`.

AUTH_DS_SET_SIGNATURES

A pipe-delimited list of index values into the RRSIG_RECORDS table for relevant DNSSEC signatures. Use this index as well as the CENSUS_START and CENSUS_END to find the intended records.

This cell may be blank or may have one of the values `_not applicable_` or `_unset_`.

AUTH_NAMESERVER_SET

A pipe-delimited list of domain names appearing at the zone's apex. Domain names should appear in lower case, alabel form, use a DNS domain name library to perform comparisons. Each domain name, in lowercase, can be used as an index into the NAMESERVERS table, using the OWNER column. This set may arrive with one or more DNSSEC signatures.

This cell may be blank or may have one of the values `_not applicable_` or `_unset_`.

AUTH_NAMESERVER_SET_SIGNATURES

A pipe-delimited list of index values into the RRSIG_RECORDS table for relevant DNSSEC signatures. Use this index as well as the CENSUS_START and CENSUS_END to find the intended records.

This cell may be blank or may have one of the values `_not applicable_` or `_unset_`.

AUTH_DNSKEY_SET

A pipe-delimited list of index values into the DNSKEY_RECORDS for DNSKEY records that appear at a zone's apex. The list may be empty.

This cell may be blank or may have one of the values `_not applicable_` or `_unset_`.

AUTH_DNSKEY_SET_SIGNATURES

A pipe-delimited list of index values into the RRSIG_RECORDS table for relevant DNSSEC signatures. Use this index as well as the CENSUS_START and CENSUS_END to find the intended records.

This cell may be blank or may have one of the values `_not applicable_` or `_unset_`.

CUT_NAMESERVER_SET

A pipe-delimited list of domain names appearing in referral messages from the parent zone. In some cases, this set will not be possible to determine without access to the parent zone information. (If all of the parent zone's name servers are also name servers for the child, each name server will answer from the child zone.) Domain names should appear in lower case, alabel form, use a DNS domain name library to perform comparisons. Each domain name can be used as an index into the NAMESERVERS table, using the OWNER column. This will not arrive with a DNSSEC signature.

This cell may be blank or may have one of the values `_not applicable_`, `_cannot be determined_`, or `_unset_`. When a parent zone has the same set of name servers as a child zone, the cut point name server set is not exposed via the query protocol, hence the cannot be determined value.

NON_DELEGATED_NAMES

A pipe-delimited list of domain names in the zone that are not delegated, i.e., they do not own an NS resource record set. Empty non-terminal names are included.) Domain names should appear in lower case, alabel form, use a DNS domain name library to perform comparisons.

This cell may be blank, meaning that no names were identified despite data being available, or it may be `_not determined in code_` meaning that there was no attempt made to find these names.

NON_DELEGATED_CNAMES

A pipe-delimited list of domain names in the zone, each domain naming owning a CNAME resource record set. Domain names should appear in lower case, alabel form, use a DNS domain name library to perform comparisons.

This cell may be blank, meaning that no names were identified despite data being available, or it may be `_not determined in code_` meaning that there was no attempt made to find these names.

NON_DELEGATED_DNAMES

A pipe-delimited list of domain names in the zone, each domain naming owning a DNAME resource record set. Domain names should appear in lower case, alabel form, use a DNS domain name library to perform comparisons.

This cell may be blank, meaning that no names were identified despite data being available, or it may be `_not determined in code_` meaning that there was no attempt made to find these names.

DELEGATED_SUB_TLDS

A pipe-delimited list of domain names in the zone, each domain naming owning a NS resource record. Determination of being a `sub_tld` is heuristic, currently based on sharing at least one NS resource record target with the parent zone and some other rules. Domain names should appear in lower case, alabel form, use a DNS domain name library to perform comparisons. In lowercase form, these names can be used as an index into the ZONES table, ALABEL column.

This cell may be blank, meaning that no names were identified despite data being available, or it may be `_not determined in code_` meaning that there was no attempt made to find these names.

SIZE_DATA_NS_OWNERS

All columns beginning with `SIZE_DATA_` are derived from zone files accessible by the census process. The zone file may be the result of a live zone transfer or be a reported copy corresponding to midnight beginning the day the census is run. I.e., from `CENSUS_START`, use the date portion, the zone file is supposed to be from the start of that date.

This is the count of zone cuts in the zone, the number of NS resource record sets other than the apex set.

This cell may have the value `_not determined in code_` or `_unset_`, meaning that there was no attempt made to find these names. A value of 0 is used if the zone has no active nameservers.

SIZE_DATA_DS_OWNERS

All columns beginning with `SIZE_DATA_` are derived from zone files accessible by the census process. The zone file may be the result of a live zone transfer or be a reported copy corresponding to midnight beginning the day the census is run. I.e., from `CENSUS_START`, use the date portion, the zone file is supposed to be from the start of that date.

This is the count of DS resource record sets in the zone.

This cell may have the value `_not determined in code_` or `_unset_`, meaning that there was no attempt made to find these names. A value of 0 is used if the zone has no active nameservers.

SIZE_DATA_NSEC_OWNERS

All columns beginning with `SIZE_DATA_` are derived from zone files accessible by the census process. The zone file may be the result of a live zone transfer or be a reported copy corresponding to midnight beginning the day the census is run. I.e., from `CENSUS_START`, use the date portion, the zone file is supposed to be from the start of that date.

This is the count of NSEC resource record sets in the zone. For many TLDs, those that feature only delegations, this count will be one more than the `NS_OWNERS` due to the apex.

This cell may have the value `_not determined in code_` or `_unset_`, meaning that there was no attempt made to find these names. A value of 0 is used if the zone has no active nameservers.

SIZE_DATA_NSEC3_OWNERS

All columns beginning with `SIZE_DATA_` are derived from zone files accessible by the census process. The zone file may be the result of a live zone transfer or be a reported copy corresponding to midnight beginning the day the census is run. I.e., from `CENSUS_START`, use the date portion, the zone file is supposed to be from the start of that date.

This is the count of NSEC3 resource record sets in the zone. With the possible use of opt-out, this count may be much less than the `NS_OWNERS`, may be close to the `DS_OWNERS`, or could be much higher if there are non-delegated names in the zone.

This cell may have the value `_not determined in code_` or `_unset_`, meaning that there was no attempt made to find these names. A value of 0 is used if the zone has no active nameservers.

SIZE_DATA_IDN_NS_OWNERS

All columns beginning with `SIZE_DATA_` are derived from zone files accessible by the census process. The zone file may be the result of a live zone transfer or be a reported copy

corresponding to midnight beginning the day the census is run. I.e., from CENSUS_START, use the date portion, the zone file is supposed to be from the start of that date.

This is the count of zone cuts in the zone whose alabel forms begin with "xn--", not including the apex set.

This cell may have the value `_not determined in code_` or `_unset_`, meaning that there was no attempt made to find these names. A value of 0 is used if the zone has no active nameservers.

SIZE_DATA_IDN_DS_OWNERS

All columns beginning with `SIZE_DATA_` are derived from zone files accessible by the census process. The zone file may be the result of a live zone transfer or be a reported copy corresponding to midnight beginning the day the census is run. I.e., from CENSUS_START, use the date portion, the zone file is supposed to be from the start of that date.

This is the count of DS resource record sets in the zone whose alabel forms begin with "xn--".

This cell may have the value `_not determined in code_` or `_unset_`, meaning that there was no attempt made to find these names. A value of 0 is used if the zone has no active nameservers.

Contact Information

Questions, suggested edits, requests should be sent to edward.lewis@icann.org.