



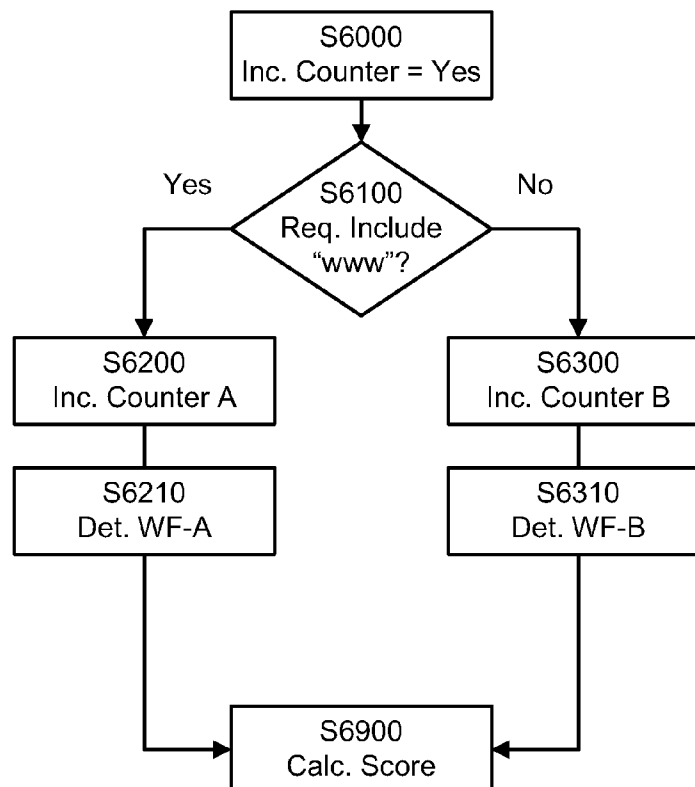
US 20100257266A1

(19) **United States**(12) **Patent Application Publication**  
**Holmes et al.**(10) **Pub. No.: US 2010/0257266 A1**(43) **Pub. Date: Oct. 7, 2010**(54) **EXISTENT DOMAIN NAME DNS TRAFFIC  
CAPTURE AND ANALYSIS**(75) Inventors: **Alexander D. Holmes**, Sterling, VA  
(US); **Andrew W. Simpson**,  
Sterling, VA (US); **Karthik**  
**Shyamsunder**, Winchester, VA  
(US); **Srinivas Sunkara**, Ashburn,  
VA (US); **Eyal Lanxner**, Kfar Saba  
(IL); **Nir Zohar**, Kfar Saba (IL);  
**Leonard Orentas**, South Riding,  
VA (US); **Matt Larson**, Dulles, VA  
(US); **Mark Kusters**, Clifton, VA  
(US); **Shevek Mankin**, (US);  
**Tamir Kamara**, Neve Yamin (IL)

Correspondence Address:

**TOWNSEND AND TOWNSEND AND CREW,  
LLP**  
**TWO EMBARCADERO CENTER, EIGHTH  
FLOOR**  
**SAN FRANCISCO, CA 94111-3834 (US)**(73) Assignee: **VeriSign, Inc.**, Mountain View, CA  
(US)(21) Appl. No.: **12/609,831**(22) Filed: **Oct. 30, 2009****Related U.S. Application Data**(60) Provisional application No. 61/167,521, filed on Apr.  
7, 2009, provisional application No. 61/167,523, filed  
on Apr. 7, 2009, provisional application No. 61/167,  
528, filed on Apr. 7, 2009.**Publication Classification**(51) **Int. Cl.**  
**G06F 15/173** (2006.01)  
**G06Q 30/00** (2006.01)  
(52) **U.S. Cl.** ..... **709/224**; 705/14.69; 709/223  
(57) **ABSTRACT**

Systems and methods for scoring a domain web traffic based on DNS traffic requests received at an authoritative name server to resolve the domain name. A request to resolve the domain name is received at an authoritative name server. A counter, such as a server counter or a hit counter, for the domain name is incremented based on the received request. A score, such as a domain traffic score or a domain rank, is calculated based upon a count of the counter. Calculating the score may also include applying a weighting factor to the counters based on information about a requesting set of resolvers and other domains/websites that may be linking and driving traffic to the domain whose traffic score is being calculated. Examples of relevant set of resolvers information may include location, traffic levels, traffic type and architecture of the set of resolvers.






Figure 1




Figure 2

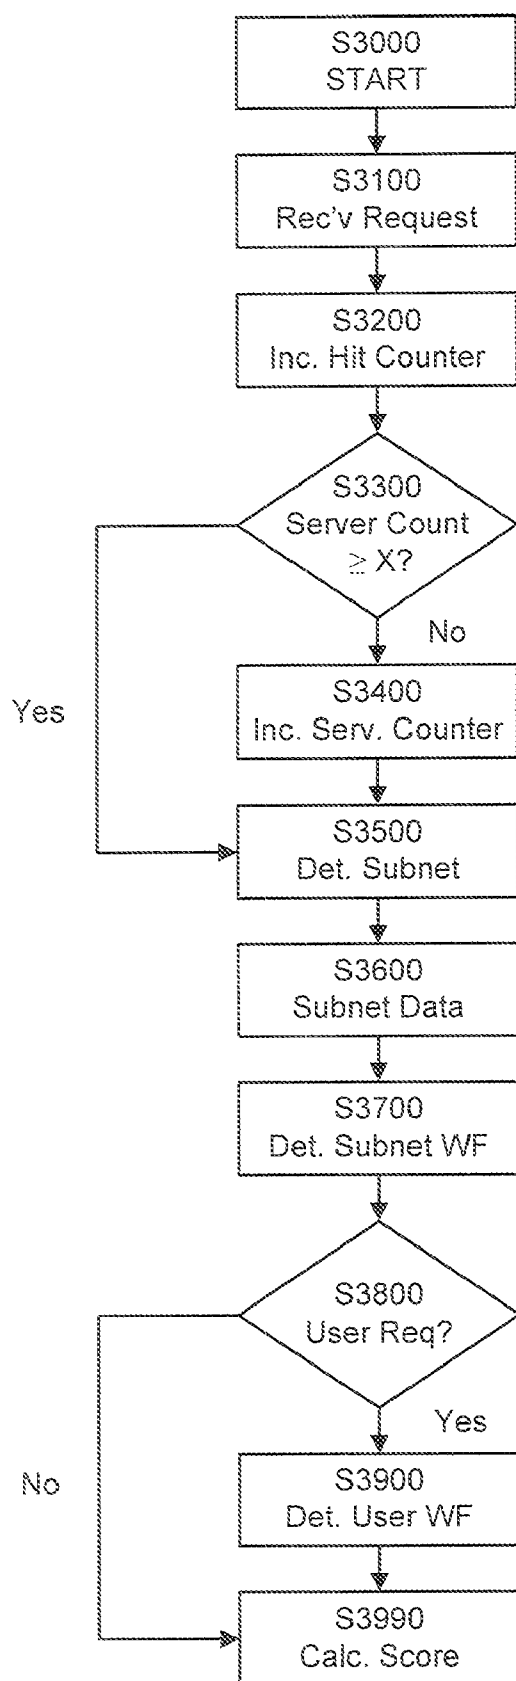
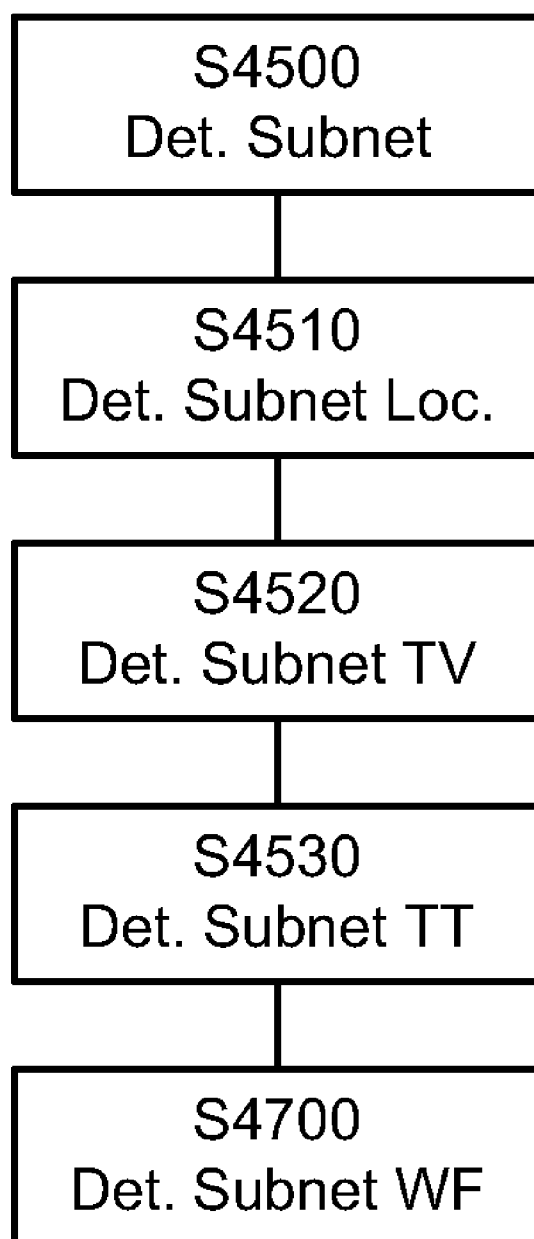
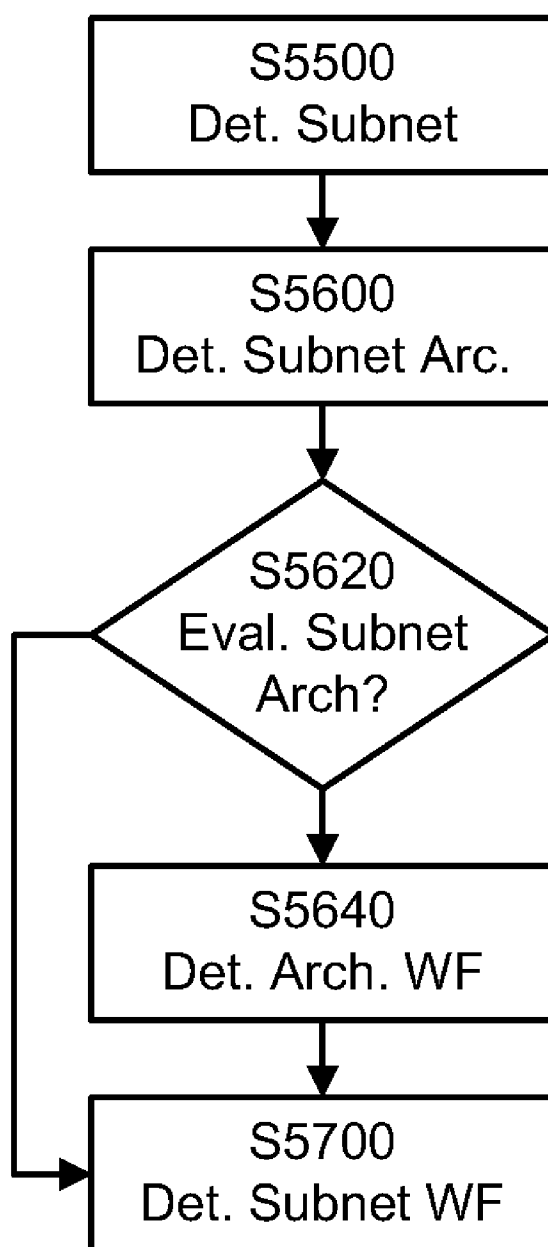
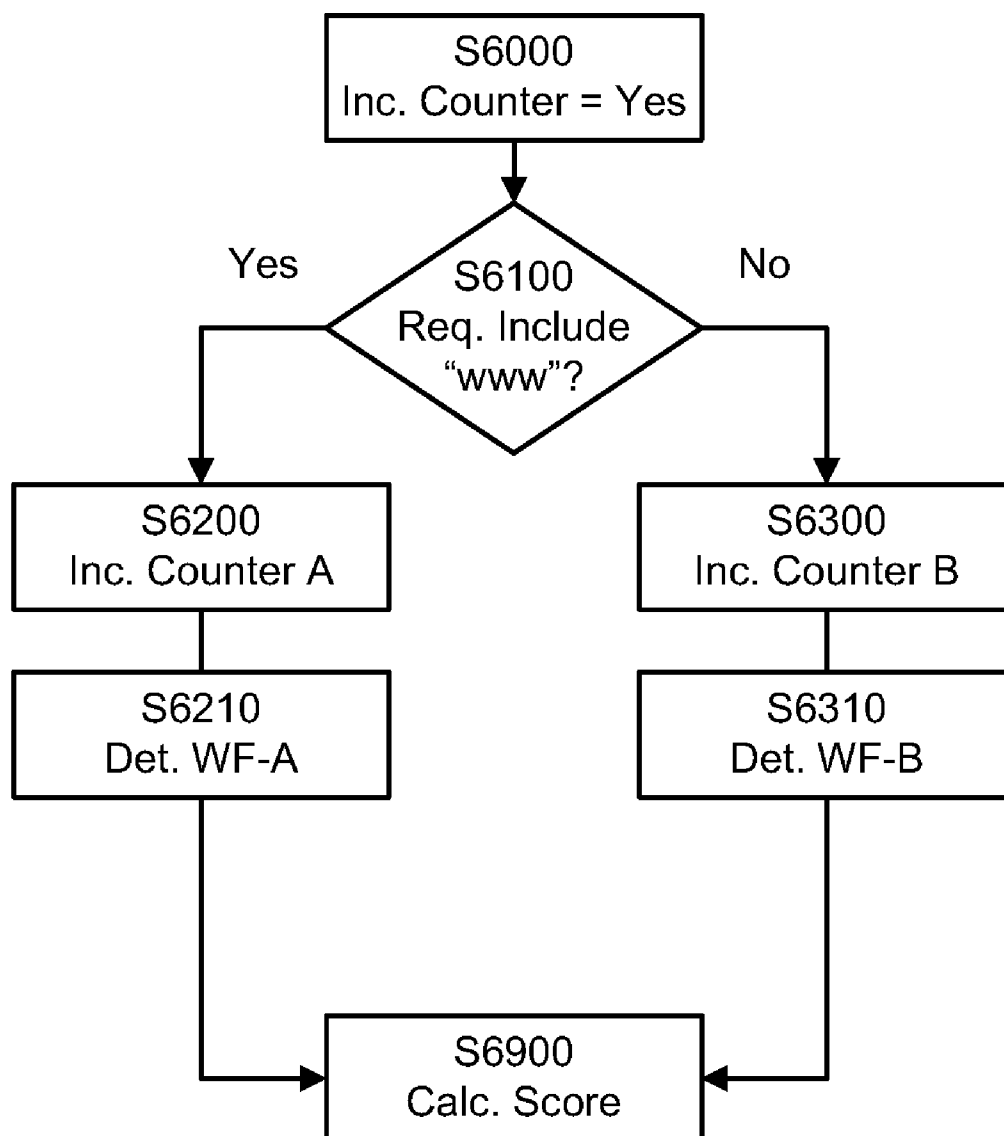


Figure 3

**Figure 4**

**Figure 5**

**Figure 6**

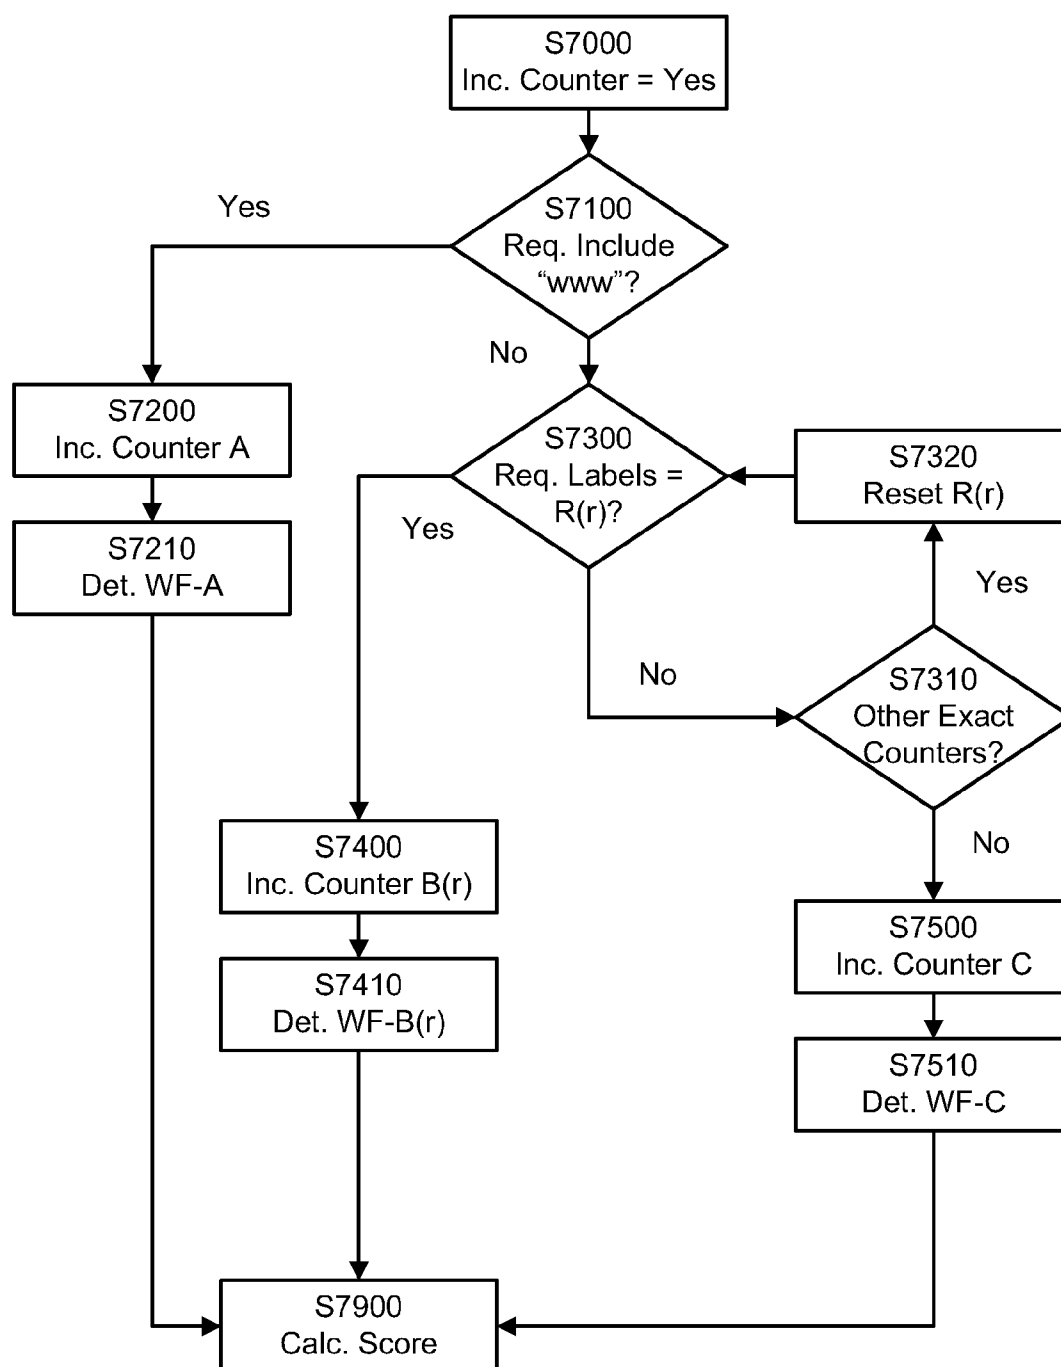



Figure 7





**Figure 8**




Figure 9




Figure 10




Figure 11




Figure 12

## EXISTENT DOMAIN NAME DNS TRAFFIC CAPTURE AND ANALYSIS

### BACKGROUND OF THE INVENTION

**[0001]** Understanding user activity on the Internet is becoming more important, and more difficult, as the Internet continues to expand. Commercial use of the internet is one area that has expanded dramatically in the last decade, and one that has a particular interest in understanding, monitoring and predicting user activity. One significant aspect regarding the commercial use of the internet is advertising. Advertisers may use factors such as traffic rankings in determining an appropriate web site or domain to advertise particular content. Likewise, web site owners may use traffic ranking to establish an appropriate fee for advertising on their web site. As known by those of skill in the art, internet advertising has also taken many different forms that may directly demonstrate the effectiveness of a particular advertisement, such as pay per click (PPC) applications. However, traffic rankings are still an important aspect of internet advertising, as well as other aspects of internet infrastructure management. For example, as the number of active web sites on the internet grows, there is an increased demand for accurate traffic ratings to inform decisions such as server management, web development, advertising focus and rates. However, there are limitations on the capabilities of conventional traffic monitoring services that typically monitor the traffic of users or web sites to calculate traffic scores.

**[0002]** A description of the ways in which the Internet is intrinsically organized can be helpful in understanding the challenges related to efficiently monitoring and rating the traffic for particular web sites on the internet.

**[0003]** The process of establishing a web site on the internet typically begins with a registrant registering a specific domain name through a registrar. The registrant is typically an individual or organization that identifies a domain name, such as "example.com". The registrant contacts a registrar to process the name registration. The registrar sends the necessary domain name service (DNS) information to a registry. A registrar may maintain a database containing additional customer information beyond that which is sent to the registry.

**[0004]** The registry receives DNS information from registrars, inserts that information into a centralized database and propagates the information on the internet so that domain names can be found by users around the world.

**[0005]** In general, the DNS is the part of the Internet infrastructure that translates human-readable domain names into the Internet Protocol (IP) numbers needed to establish TCP/IP communication over the Internet. That is, DNS allows users to refer to web sites, and other resources, using easier to remember domain names, such as "www.example.com", rather than the numeric IP addresses, such as "123.4.56.78", assigned to computers on the Internet. Each domain name is made up of a series of character strings (labels) separated by dots. The right-most label in a domain name is known as the "top-level domain" (TLD). Examples of well-known TLDs are ".com"; ".net"; ".org." etc. Each TLD supports second-level domains, listed immediately to the left of the TLD, e.g. the "example" level in "www.example.com". Each second-level domain can include a number of third-level domains located immediately to the left of the second-level domain, e.g. the "www" level in "www.example.com". There can be additional level domains as well, with virtually no limitation.

For example, a domain with additional domain levels could be "www.photos.example.com".

**[0006]** Additional non-domain information may be included in a Uniform Resource Identifier ("URI") structure that includes the domain name. For example, a "path" part is a sequence of segments (conceptually similar to directories, though not necessarily representing them) separated by a forward slash ("/"). This information may be included immediately to the right of the domain name, such as the "blog" in "www.example.com/blog", and may be used by a server or other receiving device to identify and deliver specific content or run particular code. Other examples of non-domain information may include queries and fragments, the specifics of which are understood by those of ordinary skill in the art and are not discussed in detail herein. Combinations of this information may be included in web page hyperlinks that navigate a user to another section of the same page or to another web page that may be part of the same, or a different, domain.

**[0007]** Related domain names, and content, may be organized in a hierarchical, or nested, manner, such as "www.example.com"; "www.blog.example.com"; "www.example.com/blog"; or "blog.example.com" etc, each with a different significance. Such related domains need not share similarities in the actual IP address to which the various domain names resolve to. In this regard, part of the domain name may signify a particular server which is desired, for example, "mail.example.com" and "www.example.com" may resolve to different servers, with different functions, for the same second-level domain.

**[0008]** The above registration and structural aspects of the internet are then used by end-user applications to find specific resources on the internet by using the DNS resolution process. Aspects of the DNS resolution process are discussed below to aid in an understanding of the subject matter of the present application.

**[0009]** The responsibility for operating each TLD (including maintaining a registry of the second-level domains within the TLD) is delegated to a particular domain name registry. The registry is responsible for converting domain names to IP addresses ("resolving") through DNS servers that maintain such information in large databases, and operating its top-level domain. The DNS stores IP addresses and domain names, facilitating service to addresses in TLDs, such as .com, .net, .edu, and .tv. Resolving is the process by which domain names are matched with corresponding IP numbers. Resolving is accomplished by a combination of computers and software, referred to as name servers that use the data in the DNS to determine which IP numbers correspond to a particular domain name. The following general definitions will be used herein.

**[0010]** Resolve: To translate domain name to IP address.

**[0011]** Resolver: A computer issuing a query in order to resolve a domain name.

**[0012]** Name server: A computer receiving queries and answering them directly or via resolve against other name servers.

**[0013]** Subnet: A group of IP addresses sharing octets of the IP address.

**[0014]** Internet domains can be divided to groups according to their TLD suffix (e.g. .com, .net, .co.uk . . .) with different registries responsible for each of them. A single registry may be responsible for several of these groups, such as the Veri-Sign registry which is responsible for .com and .net domains.

**[0015]** The DNS is maintained by a distributed database system, which uses the client-server model. The nodes of this database are the name servers. Each domain or subdomain has one or more authoritative DNS servers that publish information about that domain and the name servers of any domains subordinate to it. The top of the hierarchy is served by the root name servers, the servers to query when looking up (resolving) a TLD.

**[0016]** The DNS distributes the responsibility of assigning domain names and mapping those names to IP addresses by designating authoritative name servers for each domain. Authoritative name servers are assigned to be responsible for their particular domain.

**[0017]** In theory a fully qualified domain name may have several name segments, (e.g. "www.one.type.example.com.") For querying purposes, the name segment is typically interpreted by segment, from right to left. At each step along the way, a corresponding DNS server is queried to provide a pointer to the next server which it should consult.

**[0018]** Because of the huge volume of requests generated by DNS, the resolution process also allows for caching (i.e. the local recording and subsequent consultation of the results of a DNS query) for a given period of time after a successful answer. How long a resolver caches a DNS response (i.e. how long a DNS response is considered valid) is determined by a value called the time to live (TTL). The TTL is generally set by the administrator of the DNS server handling the response. The period of validity may vary from just seconds to days or even weeks.

**[0019]** Based on the DNS structure, as well as the caching function, there are two classifications typically applied to the name servers, authoritative and recursive (caching). An authoritative name server is a name server that gives original, definitive answers ("authoritative" answers) to DNS queries. Every domain name must be assigned a set of authoritative name servers that are responsible for resolving the domain name.

**[0020]** As indicated above, the DNS also uses recursive cache servers, which store DNS query results for a period of time determined TTL of the domain name record in question. Typically, such caching DNS servers also implement the recursive algorithm necessary to resolve a given name starting with the DNS root through to the authoritative name servers of the queried domain. Internet service providers (ISPs) typically provide recursive and caching name servers for their customers. In addition, many home networking routers implement DNS caches and recursors to improve efficiency in the local network.

**[0021]** DNS "stub" resolvers are also known that essentially operate as a cache-less application to resolve DNS names into IP addresses. The DNS stub resolver forwards DNS queries to the DNS server configured for the workstation (or server) and returns the DNS server's response to the requesting software. If a stub resolver queries a caching nameserver for a record that is being held by the caching server before the TTL has expired, the caching server will reply with the cached resource record rather than retrieve it from the authoritative name server again.

#### SUMMARY OF THE INVENTION

**[0022]** There are currently several approaches used to provide web traffic information. Known methods include panel-centric approaches in which selected samples of users are provided with a client agent that monitors their web traffic

usage. There are also site-centric approaches in which web traffic statistics are gathered directly from the web server for a web site. Other methods focus on the ISP, and gather web usage traffic from logs of ISP's that are providing the web data to the client. However, the known methods each suffer from drawbacks related to the point at which the data is gathered, the amount of data gathered and/or the materiality of the data that is gathered. For example, panel-centric applications are limited in their ability to obtain holistic views of activity on the Internet because they are limited to the selected users who are being monitored. Additionally, any attempt to catalogue the activities of the billions of individual Internet users, even if possible, would create massive amounts of data that would quickly reach an unmanageable level. ISP's are also limited in their ability to characterize the overall activity of the Internet based on their particular client base, geographical interests, and other factors that inherently or implicitly restrict the traffic that the particular ISP is privy to.

**[0023]** There are additional problems in determining accurate and meaningful traffic rankings for domains. In general, Internet traffic can be divided into two groups: user oriented traffic, and machine oriented traffic. User oriented traffic can be considered as activity generated by users browsing web sites. Machine oriented traffic can be considered as activity generated by automatic processes and services (e.g. antivirus updates, operation system updates, web services, etc.). For the purposes of commerce, including online advertising, identifying the user generated traffic driven into web sites is particularly important. Therefore, methods for measuring the user oriented traffic are preferable. However, this is a significant challenge because, for example, information regarding the nature of the requestor cannot be obtained from the incoming DNS queries themselves. Another challenge in distinguishing between the traffic types is caused by the architecture of the Internet DNS system, which refers queries initiated by intermediate DNS resolvers. This can obscure the nature of the end-user that initiated the request.

**[0024]** The present subject matter addresses many of the above issues through various mechanisms using DNS traffic data obtained at authoritative name servers, such as the TLD DNS servers. According to aspects of the invention, domains across the Internet can be ranked in ways that known systems do not contemplate or achieve. In embodiments, several attributes and/or measurements may be obtained from the DNS traffic seen at authoritative name servers, such as the TLD DNS server for various domains, and used to calculate traffic scores. The calculated scores may be used to rank the domains.

**[0025]** The inventors have recognized that a challenge in converting DNS data into a meaningful and useful form is the size of the DNS data. Registries may operate several resolution sites located at different parts of the world. The sites for one registry may answer over thirty four billion queries a day. This huge number of queries presents an overwhelming challenge in attempts to analyze raw data. The present application provides systems and methods for data aggregation that may reduce the overall amount of data into a more manageable size. According to these methods, it has also been found that a more accurate picture of meaningful activity on the Internet can be developed, beyond what has been achieved by conventional methods using raw traffic data.

**[0026]** Disclosed methods may include different aggregation formats including, for example, domain summarized traffic information, which may list several statistics per

domain per day, and domain detailed traffic information, which may list more detailed information about DNS queries. It has been found that a particularly useful example of DNS query information are the identities of a set of resolvers generating the queries.

**[0027]** Embodiments of the invention may include automated methods for scoring a domain in which a request to resolve a domain name is received at an authoritative name server, such as a TLD DNS. Based on the received request, a server counter may be incremented for the domain name, and a domain traffic score may be calculated based upon the server counter. In embodiments, calculating the domain traffic score may include applying a weighting factor to the server counter based on information about a requesting a set of resolvers.

**[0028]** Embodiments may also include, the server counter being limited to a predetermined number of counts for a particular requesting a set of resolvers in a predetermined period of time. For example, a counter for a set of resolvers may be limited one count per 24-hour period. Embodiments may include resetting the counter after the predetermined period of time.

**[0029]** Embodiments may include incrementing a hit counter each time the authoritative name server receives a request to resolve the domain name. In addition, calculating the domain traffic score may be based upon the hit counter.

**[0030]** Embodiments may include the weighting factor being based on various criteria such as, for example, a geographical location of the requesting set of resolvers, a network traffic level of the requesting set of resolvers, a network traffic type, and/or an architecture of the requesting set of resolvers.

**[0031]** Embodiments may include judging whether the request to resolve the domain name is initiated by a user or traffic that is likely machine generated; and determining the weighting factor based on the judgment.

**[0032]** Embodiments may include automatically calculating domain traffic scores for a plurality of domains serviced by the authoritative name server. Additionally, the plurality of domains may be automatically ranked based on the domain traffic scores. Embodiments may include calculating a ranking score based on a relative traffic score for a domain among a predetermined category of domains. This may include comparing the domain traffic score of a domain with other domain traffic scores from the predetermined category of domains. Calculated traffic and/or ranking scores may be automatically distributed as part of a fee service, or used in determining other items such as an advertising rates, etc.

**[0033]** Embodiments may include determining various secondary values based on the rank of a domain such as, for example, an advertising rate, hosting value, and/or a popularity rating of the domain. As used herein, an advertising rate is understood as a rate at which advertising on a domain is valued for sale. The advertising rate can be used, for example, in the context of offering or completing an advertising transaction and/or providing a valuation to the owner of a domain. As used herein, a hosting value is a calculated value that represents the value of domain based on the described scoring methods. A hosting value can be used, for example, in transactions such as the sale of a domain, and/or providing the owner of a domain with a valuation of the domain. A popularity rating is a calculated value representing the popularity of the domain, and may be targeted to represent popularity of various types such as, for example, geographic popularity,

user type popularity, traffic type popularity and the like. Through various combinations of the above factors, accurate rate tables used for various purposes may be developed for an extensive network of domains. This can be advantageous in providing a reliable and centralized advertising fee service, domain resale market, and the like, based on, or sponsored by, a TLD provider with information that is not otherwise available to other entities.

**[0034]** Embodiments may include counting requests that meet predetermined criteria with different counters. For example, request that meet a predetermined format may be counted by a particular counter, such as requests that include a “www” string being counted on a separate counter from requests that do not include a “www” string. In embodiments, different counters for a domain may be weighted differently such as, for example, counts of the requests including the “www” string being weighted differently in scoring calculations than a count of the requests not including the “www” string.

**[0035]** Embodiments may include distinguishing between types of addresses requested in DNS queries. For example, embodiments may separately recognize exact requests, “www” requests, and other requests. As used herein, an exact request is a request that matches predetermined criteria, such as a specific number of labels. Exact requests may be identified in embodiments as containing two labels, such as an exact 2<sup>nd</sup> level domain name requested without any prefix. For example, an exact request for the domain “example.com” may be “http://example.com/”.

**[0036]** A “www” request is a request in which the character string “www” is included in a prefix. For example, a request that reads “http://www.example.com/”, “http://www1.example.com/”, and “http://AAA-www.example.com/”, are recognized as “www” requests. As used herein, an “other” request is a request that does not meet either of the “exact” or “www” criteria. This may include other prefixes as part of the request. For example, “http://jobs.example.com/”, “http://mail.example.com/”, and “http://AAA.BBB.example.com/”, would be recognized as “other” requests, if no corresponding exact match was identified.

**[0037]** In embodiments, the following exemplary attributes about DNS queries for a domain may be obtained and used to calculate ranking:

**[0038]** WWW request hits. A count of hits on DNS records with “www” in the first label.

**[0039]** Exact request hits. A count of hits on DNS records with only two labels, e.g. only “example.com”.

**[0040]** Other request hits. A count of hits on DNS records which fall into neither category above.

**[0041]** As discussed further below, embodiments may also include applying appropriate weighting factors to the various request hits for a domain in determining traffic and/or ranking scores.

**[0042]** Embodiments may include accounting for the following attributes related to the originating IP address and/or a set of resolvers of DNS queries, and using them to calculate the ranking:

**[0043]** WWW server count. A count of unique downstream IP addresses or a set of resolvers hitting a DNS record with “www” in the first label.

**[0044]** Exact server count. A count of unique downstream IP addresses or a set of resolvers hitting a DNS record with only two labels.



**[0045]** Other server count. A count of unique downstream IP addresses or a set of resolvers hitting DNS records which fall into neither category above.

**[0046]** As discussed further below, embodiments may also include applying appropriate weighting factors to the various server counts for a domain in determining traffic and/or ranking scores.

**[0047]** Embodiments may include calculating the traffic and/or ranking scores based on a TTL of DNS records for the domain.

**[0048]** Further advantages of the present subject matter will become apparent to those of ordinary skill in the art upon reading and understanding the following detailed description of the preferred embodiments.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0049]** FIG. 1 depicts an exemplary system and related network in accordance with an embodiment of the present invention;

**[0050]** FIG. 2 depicts an exemplary system and related network in accordance with an embodiment of the present invention;

**[0051]** FIG. 3 depicts aspects of an exemplary method in accordance with an embodiment of the present invention;

**[0052]** FIG. 4 depicts aspects of an exemplary method in accordance with an embodiment of the present invention;

**[0053]** FIG. 5 depicts aspects of an exemplary method in accordance with an embodiment of the present invention;

**[0054]** FIG. 6 depicts aspects of an exemplary method in accordance with an embodiment of the present invention;

**[0055]** FIG. 7 depicts aspects of an exemplary method in accordance with an embodiment of the present invention;

**[0056]** FIG. 8 depicts aspects of an exemplary method in accordance with an embodiment of the present invention;

**[0057]** FIG. 9 depicts an exemplary system and related network in accordance with an embodiment of the present invention;

**[0058]** FIG. 10 depicts a chart reflecting a relation between domain traffic size and the percentage of the traffic going to the “Other” targets;

**[0059]** FIG. 11 depicts a chart reflecting a relation between domain traffic size and the percentage of the traffic going to the “WWW” and “Exact” targets; and

**[0060]** FIG. 12 depicts a chart reflecting numbers of domains grouped according to aspects of the invention.

#### DETAILED DESCRIPTION OF THE INVENTION

**[0061]** The following detailed description is provided with reference to exemplary embodiments for the ease of description and understanding. Applicants’ invention is not limited to the disclosed embodiments, and encompasses other variations that fall within the overall scope of description provided herein.

**[0062]** As indicated above, embodiments of the present invention may involve operations on a network, such as the Internet environment depicted in FIG. 1. A client computing device **110** may initiate a DNS request **103** to recursive name server **130**. As suggested by FIG. 1, DNS requests **103** can be originated, or transmitted to recursive name server **130**, from various sources such as mobile device **120**, wireless computing device **122**, other communication link **124**, and/or intermediate network servers **126**. The DNS request to resolve a domain name typically has *n* labels separated by periods.

These labels are usually in some form that is easier to remember than the numerical IP address for a desired Internet resource, e.g. a web page. For the ease of description, the leftmost label may be considered to be the first label and the TLD is the *n*th label. Thus, in a request for “www.example.com”, *n*=3, the first label would be “www” and the TLD would be “.com”. As described above, the recursive name server **103** may transmit the DNS request to various authoritative servers **140**, **150**, **160** as part of the recursive name resolution. Alternatively, the requested DNS information may reside in a cache of recursive name server **130**, subject to the DNS record’s TTL, in which case, the query may be answered without referring to the authoritative name servers **140**, **150**, **160**.

**[0063]** FIG. 2 depicts additional details regarding a recursive name server **220** and its interaction with authoritative name servers **230**, **240**, **250**. In FIG. 2, authoritative servers **230** are root level authoritative servers. Each of these servers contains information for particular TLDs on the internet. The root level servers can direct requests for domains within their TLD to other authoritative servers managed by that TLD registry. For example, DNS request **202** from client **210** may include a request for “www.example.com”. Recursive name server **220** may first check an internal cache for a corresponding DNS record. If one is not found, the DNS request may be forwarded at **203** to root level authoritative name servers **230**. An authoritative root level server among servers **230** that is responsible for “.com” may return DNS information for “example.com” directing the requestor to authoritative name servers **240**, in this case these servers represent the constellation of servers for a registry of “.com”. As described herein, DNS requests, e.g., **203**, **204**, **205**, can be counted at various authoritative servers, e.g., **230**, **240**, **250**. In embodiments, the counting may be done at a TLD name server such as, for example, **230**. Other authoritative name servers **240**, **250** can also count DNS hits for those records handled by the respective server. Embodiments may include combining the scores from various servers, and/or TLDs, e.g. “.com”, “.net”, “.edu”, “.uk”, to arrive at combined rankings.

**[0064]** In embodiments, determinations may be made to count a request or not count a request on a given counter. For example, certain counters, such as a “server counter”, may be limited to a single count per a predetermined period of time. As depicted in FIG. 3, a request may be received in **S3100**. A relevant hit counter may be incremented in **S3200**. In **S3300** a determination may be made whether to increment a server counter. In this example, a current count of the server counter may be compared to a predetermined value “X”. If the server count is greater than, or equal to, X, the determination is made not to increment the server counter again, and the method may proceed to **S3500**. If the server count is less than X, the method may proceed with **S3400** where the server counter is incremented. The server count may be automatically reset after the predetermined period of time. Although described in the context of a server counter, other counters may be limited in similar manner, such as, counters associated with a particular IP address, a set of resolvers, etc. As described further below, limiting the count of certain counters may be used to develop a more accurate picture of meaningful Internet activities. These limited counters may be considered alone or in combination with hit counters to provide users with a manageable and accurate representation of overall web traffic. For example, using a limited counter may reduce the unwanted

effects of artificial or unreliable traffic indicators, that may originate from a particular server, a set of resolvers, IP address, etc.

[0065] In S3500, a set of resolvers associated with the request may be determined. For ease of depiction in the figures, an exemplary set of resolvers may be referred to as a subnet. However, as discussed herein, features related to sets of resolvers are not limited to subnets per se, and may include other sets of resolvers that may represent various associations, such as, for example, address, geographic, organizational, and type, and combinations thereof. In embodiments, the determination may be made prior to determining whether to increment the server counter in S3300, such as to identify the appropriate counter to be evaluated. Data relevant to the set of resolvers may be determined and/or accessed in S3500. For example, as discussed further below, data relevant to the location, traffic level, traffic type etc. for the set of resolvers may be determined, calculated and/or accessed. In embodiments, information about a particularly identified requesting server may be used to apply an appropriate weighting factor, such as, for example, if the identified server is known to be utilized by web crawlers, or the like, that generate significant non-user traffic. A list of such servers may be stored in a database with or without appropriate weighting factors. Thus, requesting servers that appear to be, or are documented to be, data-oriented in nature can be penalized in the domain scoring. The method continues with S3700.

[0066] In S3700, a weighting factor (WF) for the set of resolvers is determined. This determination may be based on various data, as determined in S3600, and may be applied in calculating the score for the domain in S3900.

[0067] In S3800, the request may be evaluated to determine whether it represents user generated traffic. This determination may be based on various factors, such as, for example, information embedded in the request, information regarding the originating set of resolvers, the subnet architecture and the like. Exemplary methods of making this determination are described further below. If S3800 indicates that the request is likely related to user generated traffic, a user weighting factor may be determined in S3900. In embodiments, user generated traffic may be given a higher WF than other traffic, such as traffic that is of undetermined origin, or traffic that is likely machine generated. In embodiments, only user generated traffic may be considered by applying a WF=1, and all other traffic disregarded by applying a WF=0.

[0068] With reference to FIG. 4, exemplary methods may include determining a set of resolvers in S1500 associated with a received request, e.g. in similar manner to S3500 in FIG. 3. This type of information is typically available to the authoritative server with the received request in order to provide return address information for the domain address information being requested from the authoritative server. The method may continue with S4510.

[0069] In S4510, the request may be evaluated to determine a set of resolvers location. The inventors have found that the origin location, such as the country of the resolving server, has significance in representing the number of users behind its queries. This is due to various factors, including the fact that, in some countries in which the Internet is strongly developed, there are many more resolvers per Internet user than in other, less-developed, countries. These numbers are not necessarily static and can be refined on an ongoing basis to provide even more valuable information.

[0070] As shown in Table 1, data regarding the number of users per set of resolvers in different geographic regions can be used to advantageously estimate a traffic score represented by individual requests from various sets of resolvers, such as subnets.

TABLE 1

Country	Users	Total Subnets	Small Subnets	Small Subnet Ratio	Users per Subnet
Country A	34,820,000	22,482	4,320	19%	1,917
Country B	34,708,144	35,203	9,129	26%	1,331
Country C	32,700,000	30,931	5,444	18%	1,283
Country D	28,000,000	30,070	8,312	28%	1,287
Country E	26,500,000	17,958	1,197	7%	1,581

[0071] In embodiments, a number of users per set of resolvers may be used as a basis for the weight factor (WF) determined for the set of resolvers in S4700, e.g. according to the country of origin. In circumstances where these values are not well known, or subject to change, the affect of the metric may be reduced by using a non-linear function. For example, it has been determined by the inventors that a root in the 3<sup>rd</sup> order achieves good results in balancing between the range and granularity. This scale has been found to produce a max/min ratio of about 10. Information about a given country, the set of resolvers and users within the country, and/or overall countries of interest may be advantageously combined. For example, in embodiments, combinations of the geographic information above may be used in calculating the domain traffic such as by the following formula:

$$Weight_c = \sqrt[3]{\frac{InternetUsers_c}{Subnets_c - SmallSubnets_c}}$$

$$Score(d) = \frac{1}{n} \left( \frac{\sum_{c \in C} W_c(d) \cdot Weight_c}{\sum_{c \in C} UniqueW_c \cdot Weight_c} + \frac{m_e \cdot \sum_{c \in C} E_c(d) \cdot Weight_c}{\sum_{c \in C} UniqueE_c \cdot Weight_c} + \frac{m_o \cdot \sum_{c \in C} O_c(d) \cdot Weight_c}{\sum_{c \in C} UniqueO_c \cdot Weight_c} \right)$$

[0072] d—domain

[0073] C—countries of interest

[0074] W/E/O<sub>c</sub>—W/E/O subnets from country c referring to domain d

[0075] Weight<sub>c</sub>—the weight of country c

[0076] UniqueW/E/O<sub>c</sub>—unique W/E/O subnets from country c

[0077] n/m<sub>e</sub>/m<sub>o</sub>—heuristic based variables

[0078] By using such formulas in calculating a domain score, the inventors have found that a more accurate estimation of traffic volume may be achieved, particularly with domains that are more heavily frequented by non-U.S. users. That is, by separately weighting a set of resolvers according to geographic region, a more accurate representation of the global popularity of a domain may be achieved. Additionally,

in terms of advertising and value of a domain to a given customer, it may be advantageous to refine the popularity calculation to weight certain geographically based traffic more or less. For example, advertisers for a certain product or service may be interested in the relative popularity of a domain within a geographic region, or common culture, language, and the like. In those instances, embodiments of the present subject matter provide mechanisms for calculating scores that reflect the traffic that is significant to the customer, such as by adjusting the weighting factors for certain countries/regions of interest above those countries/regions that are less significant to the customer. Such methods may include receiving a criteria from a customer, such as, for example, a geographic region, calculating the scores and/or ranks for a domain based on the user-defined criteria, and reporting the results of the modified scoring to the customer. The above factors provide improved flexibility in scoring domains depending on the intended audience, e.g. a global or localized audience. Thus, aspects of the present subject matter may capitalize on the set of resolvers information received by authoritative name servers in conjunction with customer-specific information to improve and refine the scoring for domain traffic.

**[0079]** As shown in FIG. 4, as part of determining a subnet WF, embodiments may include determining a set of resolvers traffic volume (TV) in S4520. It has been found that the traffic volumes for various sets of resolvers may vary greatly across the internet, and that the traffic volume, such as the number of queries issued by the set of resolvers in a given time, can be used to estimate a WF for the set of resolvers. For example, it has been found that about 40% of the resolver IP addresses generate 10 DNS queries a day or less. This low-level TV may be used to infer a corresponding low number of users, compared to resolvers that generate a higher volume of queries per day. That is, a set of resolvers with a low number of overall queries per day may represent a low number of users, with lower level overall traffic at the resolver level, that may be assumed for purposes of weighting the relative significance of the requests received at the authoritative server.

**[0080]** Additionally, there are changing numbers and percentages of DNS resolvers that refer to a very small number of domains per day (i.e. small set of resolvers). The inventors have found it useful, in certain embodiments, to weight counters for set of resolvers that issue queries for a very small number of domains differently than a set of resolvers that issue queries for a large number of domains per day, e.g. a set of resolvers that refer to millions of domains per day. This type of sets of resolvers characteristic may be determined, for example, in S4530 as a traffic type (TT) determination. Traffic type may include other aspects of the traffic generated by a set of resolvers that may be useful in inferring the presence of user generated traffic versus machine generated traffic, which, as discussed herein, can be advantageously used in calculating the score and rank of a domain for the purposes of advertising, and the like, that are concerned with user generated traffic. Aspects of the present subject matter may capitalize on the information received by authoritative name servers in conjunction with metrics regarding requesting set of resolvers to improve and refine the scoring for domain traffic.

**[0081]** As shown in FIG. 5, embodiments of the present subject matter may also adjust a weighting factor for a set of resolvers based on characteristics of the set of resolvers architecture. This may be advantageous in circumstances such as, for example, where DNS clients are configured with primary

and secondary DNS servers. In such circumstances, different IP addresses for a plurality of servers, such as primary and secondary servers, may represent a common user base. However, each IP address and server may represent unequal levels of user activity with respect to other servers based on the relative percentage of traffic handled by the particular server. It has been found that the respective percentages for a plurality of servers of a DNS client may be found by monitoring traffic information at the authoritative name server level, and that an appropriate weighting factor may be applied based on this determination. For example, as shown in FIG. 5, a set of resolvers may be determined in S5500, e.g. in similar manner to S3500 in FIG. 3. In S5600, aspects of a set of resolvers architecture may be determined, such as determining whether there are other unique IP addresses for redundant servers for the set of resolvers. In S5620, the set of resolvers architecture may be evaluated to determine whether it meets a predetermined criteria such as, for example, is there redundancy among IP addresses for the set of resolvers. In the case of an affirmative result in S5620, the method may continue with S5640, where an architectural weighting factor may be assigned. In embodiments the architectural WF may be based on a percentage of traffic handled by the particular IP address of the set of resolvers.

**[0082]** For example, in the context of a set of resolvers with  $n$  resolvers, with a traffic percentage determined for each server, a formula may be applied in which the score for a set of resolvers may be based on discreet calculations for individual resolvers, such as individual IP addresses, within the set of resolvers that apply weighting factors based on the respective percentages to scores of individual resolvers, e.g.:

$$S_{R1} w_1 + S_{R2} w_2 + S_{Rn} w_n$$

Where:

**[0083]**  $S_{R1}$  is a score from a first resolver,

**[0084]**  $S_{R2}$  is a score from a second resolver, and

**[0085]**  $w$  is a unique weighting factor for each resolver.

**[0086]** Once the architectural WF is determined, or in the case of a negative result in S5620, the method may continue with determining an overall WF for the set of resolvers in S5700. Thus, in circumstances with separate IP addresses for a given set of resolvers, the different IP addresses may be counted with different counters, and separate weighting factors applied to the different counters.

**[0087]** With further reference to FIG. 5, an additional example of applying a set of resolvers architecture WF may include separately counting additional network entities for a given set of resolvers. In the context of set of resolvers architecture, the following entities may be separately evaluated:

**[0088]** Autonomous system number (ASN)—Globally unique number assigned to a single network or a group of networks that is managed by a single administrative entity, such as a university or business.

**[0089]** Second level domain (SLD)—The second part of the domain (e.g. “verisign” in verisign.com). This is usually associated with the ISP.

**[0090]** Organization (ORG)—The organization that makes use of the IP.

**[0091]** By examining the relationship between these entities and their corresponding set of resolvers, the inventors have found that, in circumstances, ASN may be most aggregative, followed by the SLD and ORG, which may be similar in their aggregative level. That is, an ASN entity may have a

higher number of set of resolvers represented, followed by SLD and ORG entities. In embodiments, different counters may be incremented according to the distinct ASN(s), SLD (s), and/or ORG(s) as derived from the set of resolvers. As depicted in FIG. 5, this may include determining the relevant ASN, SLD, and/or ORG in S5600-S5620, and applying an appropriate architectural WF in S5640 to a given counter for the respective set of resolvers's ASN, SLD, and/or ORG counter. It has been found by the inventors that, by using different weighting factors for the individual entities within a set of resolvers, improved predictive models may be established in calculating a domain score based on DNS traffic information. For example, traffic from a ASN may be more heavily weighted with respect to SLD and/or ORG counters in calculating the domain score. For example, an average, or other normalizing calculation, may be taken of:

$$S_{ASN} w_1 + S_{SLD} w_2 + S_{ORG} w_3 + S_{SUB} w_4$$

Where:

[0092]  $S_{ASN}$  is a score based on ASN

[0093]  $S_{SLD}$  is a score based on SLD

[0094]  $S_{ORG}$  is a score based on ORG

[0095]  $S_{SUB}$  is a score based on the subnet, and

[0096]  $w$  is a unique weighting factor for each score.

[0097] With reference to FIG. 6, exemplary methods may include determining if a counter is to be incremented with respect to a received request in S6000, e.g. in a similar manner to S3300 in FIG. 3. As detailed above, in some embodiments, a hit counter may be incremented for received DNS requests for a domain. In some instances, it may be determined that a limited counter will not be incremented. In the event that the counter in question will be incremented, the method may proceed with S6100.

[0098] In S6100, the request may be evaluated to determine whether the first label of a received domain name includes the string "www". The string may be an independent part of the requested address, as in the case of "www.example.com", or nested with other characters, such as "www1.example.com". If it is determined that the first label includes the string "www", then the method may proceed with S6200, where a counter "A" may be incremented, such as a www counter for the domain name. If the first label of the received domain name does not include the string "www", the method may proceed to S6300.

[0099] In S6300, a counter "B" may be incremented, such as a miscellaneous counter for the domain name. Although it is not required to count miscellaneous requests, i.e. requests that do not meet any of the preceding categories, a miscellaneous counter may provide advantages in evaluating the meaningfulness of detected traffic. For example, a miscellaneous counter may be used in comparison with the www counter to determine a percentage of traffic that is apparently driven by human interaction.

[0100] In S6210, and/or S6310, a weighting factor may be determined and/or assigned to the counters A and/or B, respectively. Weighting factors may be applied to all, or less than all, of the counters for a domain. Assigning a weighting factor can be based on several factors related to the nature of the counter and other information gathered with respect to the domain. Weighting factors may also be applied in a dynamic manner. For example, over the course of time, the counts of the different counters for a domain may change or information may be developed regarding the domain. This may

prompt a change in one or more of the weighting factors applied to the counters for the domain. The method may continue with S6900

[0101] In S6900, a score may be calculated for the domain, such as a traffic score. The traffic score may be based on a plurality of the counters. It should be noted that it is not required to use counts from all of the relevant counters. Scores may be calculated based on one, two or more counts, with any relevant weighting factors considered. Thus, a domain traffic score may be calculated based upon a plurality of the counters, and may include applying a weighting factor to at least one of the counters.

[0102] With reference to FIG. 7, exemplary methods may include determining if a counter is to be incremented with respect to a received request in S7000, e.g. in a similar manner to S3300 in FIG. 3. As detailed above, in some embodiments, a hit counter may be incremented for received DNS requests for a domain. In some instances, it may be determined that a limited counter will not be incremented. In the event that the counter in question will be incremented, the method may proceed with S7100.

[0103] In S7100, the request may be evaluated to determine whether the first label of a received domain name includes the string "www". The string may be an independent part of the requested address, as in the case of "www.example.com", or nested with other characters, such as "www1.example.com". If it is determined that the first label includes the string "www", then the method may proceed with S7200, where a counter "A" may be incremented, such as a www counter for the domain name. If the first label of the received domain name does not include the string "www", the method may proceed to S7300.

[0104] In S7300, the request may be evaluated to determine whether the labels satisfy a predetermined condition "R(r)". R(r) can take numerous forms including a predetermined number of labels, a specified string, or other features of DNS request. In the case of designating a number of labels, described embodiments include where R(2)=two labels. If the received domain name has only the predetermined number R(2) of labels, e.g. two labels in "example.com", then the method may proceed to S7400, where a counter B(r) is incremented, such as a two-label "exact" counter B(2) for the domain name. If the request does not have only the designed number R of labels, e.g. it has more than two labels when the predetermined number R of labels is two, then the method may proceed with S7310. Embodiments may include a plurality of B counters, e.g. multiple exact counters. These may include, for example, separate counters for various numbers for R(r), or different versions, of "exact" counters for exact portions of domain strings. As an example, an additional counter B(3) may be incremented in response to a request for "mail.example.com", in which R=3. As another example, requests with a particular string, e.g. "mail", may be counted by a counter, B("mail"), when the request contains the string, e.g. "mail5.example.com". This may be implemented by determining whether additional exact counters B(r) are being implemented in S7310. If there are other exact counters, the method may proceed to S7320 where R(r) may be reset to the new value for counter B(r). After R(r) is reset, the method may continue back to S7300 to determine if the requests satisfy the new criteria R(r). S7300, S7310 and S7320 can be reiterated until R(r) is satisfied, or there are no remaining exact counters. If there are no remaining exact counters, the method may proceed to S7500.

**[0105]** In **S7500**, a counter “C” may be incremented, such as a miscellaneous counter for the domain name. Although it is not required to count miscellaneous requests, i.e. requests that do not meet any of the preceding categories, a miscellaneous counter may provide advantages in evaluating the meaningfulness of detected traffic. For example, a miscellaneous counter may be used in comparison with the www counter and the exact counter(s) to determine a percentage of traffic that is apparently driven by human interaction.

**[0106]** In **S7210**, **S7410**, and/or **S7510**, a weighting factor may be determined and/or assigned to the counters A, B(r) and/or C, respectively. Weighting factors may be applied to all, or less than all, of the counters for a domain. For example, a weighting factor may be applied to one, two, or all counters for a domain. Assigning a weighting factor can be based on several factors related to the nature of the counter and other information gathered with respect to the domain. Weighting factors may also be applied in a dynamic manner. For example, over the course of time, the counts of the different counters for a domain may change or information may be developed regarding the domain. This may prompt a change in one or more of the weighting factors applied to the counters for the domain. The method may continue with **S7900**.

**[0107]** In **S7900**, a score may be calculated for the domain, such as a traffic score. The traffic score may be based on a plurality of the counters. It should be noted that it is not required to use counts from all of the relevant counters. Scores may be calculated based on one, two or more counts, with any relevant weighting factors considered. Thus, a domain traffic score may be calculated based upon a plurality of the counters, and may include applying a weighting factor to at least one of the counters.

**[0108]** As indicated previously, when a recursive name server gets an answer to a DNS query, it may cache it for future use. The time that each answer is saved in the cache is based on the TTL. The inventors have observed that measuring a domains’ traffic by the total number of queries targeting the domain may have inaccuracies based on the caching function and the TTL for the DNS record. Compensating for the caching function can be particularly difficult in large-scale implementations based on the fact that each domain/resolver may manage its TTL differently. For instance, a resolver may refresh its cache every 6 hours, while another resolver may refresh its cache every 12 hours. Even if both resolvers are asked for a given domain the exact number of times, roughly speaking, the first resolver would be expected to generate twice as much query traffic up the DNS chain for this domain, compared to the second, even though both resolvers are representing the same amount of user-generated traffic to the domain. The inventors have developed several features to address this issue.

**[0109]** As a first example, a weighting score may be determined based on a TTL for the DNS file. This can compensate, somewhat, for limitations of hit counters linked to authoritative name servers, that only see requests that are actually forwarded to the server. Based on the number of requests forwarded to the server, a weighting factor based on the TTL can be used to calculate an estimate of the total traffic that is being directed to the domain. In general, traffic for sites with a relatively long TTL may receive a higher weight factor than those with shorter TTLs. Weighting factors can be established based on a ratio of TTLs among the scored/ranked domains.

**[0110]** As discussed herein, disclosed methods may be directed to observing the number of distinct querying IP

addresses, set of resolvers, and the like, which has been found to introduce less noise and less dependence on the TTL and caching policy of each individual resolver. This approach has been found to be particularly useful in the field of observing traffic at the level of the TDL and other authoritative name servers, which may see enormous volumes of overall traffic. The metric of unique querying IP addresses, set of resolvers, and the like, can also prevent large offsets received by the total hit counters. For example, rather than relying exclusively on a total volume of traffic reflected by a hit counter, each distinct querying IP addresses, set of resolvers, and the like, may be counted only once per predetermined period of time, e.g. a 24-hour period.

**[0111]** In this regard, it is useful to consider two types of metrics:

**[0112]** Hits (H): The total number of queries generated for a given domain.

**[0113]** Subnets (S): The number of distinct resolver IP addresses, set of resolvers and the like generating queries for a given domain.

**[0114]** As discussed above, a given query can request one of three types of domain targets: www, Exact, Other. The inventors have observed several factors that help to meaningfully interpret the nature of the traffic based on the counts of the respective requests. This can help in understanding the extent to which each query of a particular type reflects user traffic, as opposed to non-user traffic, such as web services, etc.

**[0115]** With regards to www queries, it may be assumed that the majority of traffic referring to such targets is user-oriented. However, while www queries typically reflect user-oriented traffic in a domain, they do not always comprise all of the user-oriented traffic for the domain. It has been found that, for some domains, a portion of the user-oriented traffic resides within the “Exact” and “Other” metrics. For instance, the inventors have found that, for some popular web sites, less than 10% of the requests fall within the www category. Therefore, calculating a score based only on the www metric can result in certain web sites being ranked much lower than they should be based on the actual traffic that the web site is seeing. This has been found to be particularly relevant in evaluating the popularity of large domains in which the “Other” queries may comprise a large proportion of user-oriented traffic. For instance, a domain “example.com” may see most traffic directed to popular sub-domains, e.g. mail.example.com; games.example.com; news.example.com; etc.

**[0116]** The inventors have found that there is high correlation between the size of the domain (i.e. amount of its traffic) and the percentage of its “Other” traffic out of its overall traffic. FIG. 10 illustrates the relation between domain traffic size (in log space) and the percentage of the traffic going to the “Other” targets (sub-domains). Note that the value of each bar is the average of the percentage of all domains located in that bar.

**[0117]** This shows a clear relation between the two factors, meaning larger domains tend to have higher percentage of their total traffic in the “Other” traffic. Thus, the significance of “Other” traffic may vary depending on the overall level of traffic for a domain. This observation has been taken into account in formulating the various methodologies described herein. For example, although other traffic may normally be given a lower weight factor than exact traffic, in embodiments, the weight factor for other traffic may be increased according to relative traffic volume.

[0118] In many of the referrals to web sites by users, the “www” prefix is omitted (e.g., “example.com”). However, the ratio of user-oriented to non-user oriented traffic reflected by such requests, and hence the significance of the “exact” requests, has been uncertain.

[0119] Surprisingly, the inventors have observed that considering both of the “www” and “exact” traffic metrics yields particularly favorable results. While the “exact” metric by itself does not fulfill a clear trend or consistency, when combined with the “www” metric, a clear pattern is observed. Thus, it has been determined that the two metrics complement one another in an unexpected way, as shown in FIG. 11.

[0120] The inventors have relied on the foregoing observation to develop a streamlined and accurate approach for measuring a given domain’s performance by comparing its incoming traffic in a given period to the overall traffic in that period. In embodiments, www hits may be counted separately from all other requests. Alternatively, all three DNS target types may be included, e.g. www (W), exact (E) and other (O), and each may be assigned a different weight. As an additional feature, “set of resolvers” metrics may be used, alone, or in combination with “Hit” metrics. For example, the following represents an embodiment of the disclosed methodologies:

[0121] Calculate each domain’s W/E/O average daily ratios for the given period; in other words, for each domain, calculate the average of its daily ratio between its distinct set of resolvers metric and the overall number of distinct set of resolvers that day. This can be done for each of the target types individually; W/E/O:

$$SW\_Ratio(D, P) = \frac{1}{|P|} \cdot \sum_{Day_i \in P} \frac{SW(D, Day_i)}{SW(All, Day_i)}$$

$$SE\_Ratio(D, P) = \frac{1}{|P|} \cdot \sum_{Day_i \in P} \frac{SE(D, Day_i)}{SE(All, Day_i)}$$

$$SO\_Ratio(D, P) = \frac{1}{|P|} \cdot \sum_{Day_i \in P} \frac{SO(D, Day_i)}{SO(All, Day_i)}$$

[0122] Where D is Domain; P is Period (|P| is the size of the period in days); and SW/E/O(All, Day<sub>i</sub>) is the overall number of distinct set of resolvers querying W/E/O accordingly.

[0123] Calculate the overall score for the domain within the given period:

$$Score(D, P) = \frac{W_W \cdot SW\_Ratio(D, P) + W_E \cdot SE\_Ratio(D, P) + W_O \cdot SO\_Ratio(D, P)}{W_W + W_E + W_O}$$

[0124] Where W<sub>W/E/O</sub> is the weight of each of the target types W/E/O accordingly. In a preferred embodiment, the following approximate weights may be used in the above equation.

$$W_W=1$$

$$W_E=0.5$$

$$W_O=0.2$$

[0125] In other embodiments, the weighting factors can be described more generally. For example, W<sub>W</sub> ≥ W<sub>E</sub> ≥ W<sub>O</sub>. As

suggested previously, W<sub>O</sub> can be adjusted according to traffic volume. For example, W<sub>O</sub> may be corrected as W<sub>O</sub>=(W<sub>O</sub>K), where K is a relative factor of domain traffic volume such that K ≥ 1 represents a relatively high volume of traffic, and K ≤ 1 represents a relatively low volume of traffic.

[0126] The inventors have recognized that grouping similar scores can improve the utility of the information to various users. For example, a useful representation of groups may be created by dividing the log of the scores to 100 equal width bins where 1 is the lowest and 100 in the highest. FIG. 12 shows the number of domains in each group.

[0127] In embodiments, advertising pricing can be determined based on the group for a domain. For example, a domain in a traffic group between 90-100 may have a different advertising rate applied to it as opposed to a domain in a traffic group between 50-60. It should be noted that these groups are not limited to raw traffic score data and represent a more meaningful and valuable metric based on the methodologies described herein, as well as the broad perspective gained from the level of the TLD server. For example, the inventors have observed that the methodologies discussed herein may identify web sites and domains that are developing traffic at a lower level and in advance of conventional methods. Such results can be obtained, in part, based on the ability of the TLD to aggregate, and evaluate the significance of, traffic from across the Internet.

[0128] With reference to FIG. 8, embodiments may include taking a calculated score, such as in S3990, S6900, and/or S7900, and using that score to rank the domain. For example, a calculated traffic score for a domain may be compared against other calculated scores in S8992 in order to rank the domain among other domains on the web. Such rankings can be used in a more intuitive way for users to evaluate an appropriate pricing or value for advertising on a particular domain. The calculated ranks may be distributed as part of a pay service, or used for other purposes including system management, etc. Domain rankings can also be used in the context of commercial services, such as advertising, to calculate a value in S8994. For example, the value of advertising on a particular domain may be linked to a rank-based equation that provides automated online advertising rates. Such systems can be used to evaluate and reward web site developers that are beginning to see significant traffic, in ways that are not possible using current methods. For example, a TLD server, or other vendor with access to the disclosed information, can automatically provide offers from potential advertisers to domains that meet a particular ranking. As indicated above, the inventors have found that it is possible to identify previously unrecognized low-level traffic sites, that may be of interest to potential advertisers, when viewed from the position of the authoritative server, particularly the TLD server.

[0129] Embodiments of the present invention can include systems for implementing the described methods, as well as computer-readable storage medium coded with instructions for causing a computer to execute the described methods. For example, as shown in FIG. 9, an electronic system 900 including a processor, a memory and an electronic communication device may be configured to request DNS information via ISP server 950. The system 900 may be represent a user computer system, wireless communication devices such as 920, 970, subnetworks such as 930, 990, a server, or any other network-capable device with the requisite functional capabilities. Servers 952, 954 may operate as part of, or separate from, a DNS server associated with a registry.

**[0130]** The system **900** includes any number of processors (not shown) that are coupled to storage devices including a first storage (not shown, typically a random access memory, or “RAM”), second storage (not shown, typically a read only memory, or “ROM”). Both of these storage devices may include any suitable type of the computer-readable media described and/or mentioned above. A mass storage device (not shown) may also be used to store programs, data and the like and is typically a secondary storage medium, such as a hard disk that is slower than primary storage. It will be appreciated that the information retained within the mass storage device, may, in appropriate cases, be incorporated in standard manner as part of primary storage as virtual memory. A specific mass storage device such as a CD-ROM may also pass data uni-directionally to the processor.

**[0131]** The system **900** may also include an interface that includes one or more input/output devices such as such as video monitors, track balls, mice **904**, keyboards, microphones, touch-sensitive displays, transducer card readers, magnetic or paper tape readers, tablets, styluses, voice or handwriting recognizers, or other known input devices, including other computers **910**. The system **900** may be coupled to a computer or other electronic communication network **990**, **980** using a network connection as shown generally at **901**. The network can connect various wired, optical, electronic and other known networks to exchange information among computers **910**, servers **950**, **952**, **954**, wireless communication devices **920**, **970** and sub-networks **990**, **930**. With such a network connection, it is contemplated that the system **900** and the processor therein may receive information from the network, or may output information to the network in the course of performing the above-described method steps. The above-described devices and materials will be familiar to those of skill in the computer hardware and software arts and need not be individually or exhaustively depicted to be understood by those of skill in the art. The hardware elements described above may be configured (usually temporarily) to act as one or more modules for performing the operations described above.

**[0132]** In addition, embodiments of the present invention further include computer-readable storage media that include program instructions for performing various computer-implemented operations. The media may also include, alone or in combination with the program instructions, data files, data structures, tables, and the like. The media and program instructions may be those specially designed and constructed for the purposes of the present subject matter, or they may be of the kind available to those having skill in the computer software arts. Examples of computer-readable storage media include magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher level code that may be executed by the computer using an interpreter.

**[0133]** The invention has been described with reference to exemplary embodiments.

**[0134]** Modifications and alterations of the described embodiments may be evident to those of ordinary skill in the art upon a reading and understanding of this specification. The present invention is intended to include all such modifi-

cations and alterations insofar as they come within the scope of the appended claims, or the equivalents thereof.

What is claimed is:

1. An automated method for scoring a domain, comprising: receiving at an authoritative name server a request to resolve a domain name; incrementing a server counter for the domain name; and calculating a domain traffic score based upon the server counter, wherein, calculating the domain traffic score includes applying a weighting factor to the server counter based on information about a requesting set of resolvers.
2. The method of claim 1, wherein the server counter is limited to one count for a particular requesting set of resolvers in a predetermined period of time.
3. The method of claim 2, further comprising incrementing a hit counter each time the authoritative name server receives a request to resolve the domain name, wherein, calculating the domain traffic score is further based upon the hit counter.
4. The method of claim 1, wherein the weighting factor is based on a geographical location of the requesting set of resolvers.
5. The method of claim 1, wherein the weighting factor is based on at least one of a network traffic level and a network traffic type of the requesting set of resolvers.
6. The method of claim 1, wherein the weighting factor is based on an architecture of the requesting set of resolvers.
7. The method of claim 1, further comprising: judging whether the request to resolve the domain name is initiated by a user; and determining the weighting factor based on the judgment.
8. The method of claim 1, further comprising: automatically calculating domain traffic scores for a plurality of domains serviced by the an authoritative name server; automatically ranking the plurality of domains based on the domain traffic scores; and determining an advertising rate for at least one of the domains based on the rank of the at least one domain.
9. The method of claim 1, wherein requests that include a “www” string are counted separately from requests that do not include a “www” string.
10. The method of claim 9, wherein a count of the requests including the “www” string are weighted differently in the calculation than a count of the requests not including the “www” string.
11. A system for scoring a domain, comprising: a processor; a memory programmed with computer-executable code that causes the processor to perform a method comprising: receiving at an authoritative name server a request to resolve a domain name; incrementing a server counter for the domain name; and calculating a domain traffic score based upon the server counter, wherein, calculating the domain traffic score includes applying a weighting factor to the server counter based on information about a requesting set of resolvers.
12. The system of claim 1, wherein the server counter is limited to one count for a particular requesting set of resolvers in a predetermined period of time.

**13.** The system of claim **2**, further comprising incrementing a hit counter each time the authoritative name server receives a request to resolve the domain name,

wherein, calculating the domain traffic score is further based upon the hit counter.

**14.** The system of claim **1**, wherein the weighting factor is based on a geographical location of the requesting set of resolvers.

**15.** The system of claim **1**, wherein the weighting factor is based on at least one of a network traffic level and a network traffic type of the requesting set of resolvers.

**16.** The system of claim **1**, wherein the weighting factor is based on architecture of the requesting set of resolvers.

**17.** The system of claim **1**, further comprising:

judging whether the request to resolve the domain name is initiated by a user; and

determining the weighting factor based on the judgment.

**18.** The system of claim **1**, further comprising:

automatically calculating domain traffic scores for a plurality of domains serviced by the authoritative name server;

automatically ranking the plurality of domains based on the domain traffic scores; and

determining an advertising rate for at least one of the domains based on the rank of the at least one domain.

**19.** The system of claim **1**, wherein requests that include a “www” string are counted separately from requests that do not include a “www” string.

**20.** The system of claim **9**, wherein a count of the requests including the “www” string are weighted differently in the calculation than a count of the requests not including the “www” string.

\* \* \* \* \*