



(51) International Patent Classification:

G06F 16/958 (2019.01) G06F 16/951 (2019.01)

(21) International Application Number:

PCT/EP2022/067331

(22) International Filing Date:

24 June 2022 (24.06.2022)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/219,660	08 July 2021 (08.07.2021)	US
17/373,287	12 July 2021 (12.07.2021)	US
17/373,312	12 July 2021 (12.07.2021)	US
17/373,482	12 July 2021 (12.07.2021)	US
17/373,570	12 July 2021 (12.07.2021)	US
17/373,608	12 July 2021 (12.07.2021)	US
17/373,634	12 July 2021 (12.07.2021)	US

(71) Applicant: **METACLUSTER LT, UAB** [LT/LT]; A. Gostauto g. 40A, 03163 Vilnius (LT).(72) Inventors: **VILCINSKAS, Eivydas**; Korsako Street 15-6, 78358 Siauliai City (LT). **PETRUŠKEVICIUS, Arnas**; Kalvarijų g. 137D-27, 08221 Vilnius (LT). **STALIO-RAITIS, Giedrius**; Burbiškių 60-14, 02198 Vilnius (LT).**JURAVICIUS, Martynas**; Naugarduko Street 62-85, 03203 Vilnius (LT). **STANKEVICIUS, Rimantas**; Bitininku street 4F-23, 08310 Vilnius (LT).(74) Agent: **MEISSNER BOLTE PATENTANWÄLTE RECHTSANWÄLTE PARTNERSCHAFT MBB**; MEYER, Florian, Widenmayerstr. 47, 80538 Munich (DE).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

(54) Title: WEB SCRAPING THROUGH USE OF PROXIES, AND APPLICATIONS THEREOF




FIG. 1

(57) Abstract: Embodiments disclose a system that allows for improved generation of web requests for scraping that, because of the nature of the requests and time and manner they are sent out, appear more organic, as in human generated, than conventional automated scraping systems. The system then manages how a client request to scrape a target website is made to the site, masking the request in a manner that makes it appear to the Web server as if the request is not generated by an automated system. In this way, by appearing more organic, Web servers may be less likely to block requests from the disclosed system or may take longer to block requests from the disclosed system. By avoiding Web servers blocking requests and extending the lifetime of IP proxies before they are blocked, embodiments can use a limited IP proxy address space more efficiently.

MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report (Art. 21(3))*
- *with amended claims (Art. 19(1))*
- *in black and white; the international application as filed
contained color or greyscale and is available for download
from PATENTSCOPE*

WEB SCRAPING THROUGH USE OF PROXIES, AND APPLICATIONS THEREOF

BACKGROUND

Field

[0001] This field is generally related to web scraping.

Related Art

[0002] Web scraping (also known as screen scraping, data mining, web harvesting) is the automated gathering of data from the Internet. It is the practice of gathering data from the Internet through any means other than a human using a web browser. Web scraping is usually accomplished by executing a program that queries a web server and requests data automatically, then parses the data to extract the requested information.

[0003] To conduct web scraping, a program known as a web crawler may be used. A web crawler, sometimes called a web spider, is a program or an automated script which performs the first task, i.e. it navigates the web in an automated manner to retrieve data, such as Hypertext Transfer Markup Language (HTML) data, JSONs, XML, and binary files, of the accessed websites.

[0004] Web scraping is useful for a variety of applications. In a first example, web scraping may be used for search engine optimization. Search engine optimization (SEO) is the process of improving the quality and quantity of website traffic to a website or a web page from search engines. A web search engine, such as the Google search engine available from Google Inc. of Mountain View, California, has a particular way of ranking its results, including those that are unpaid. To raise the location of a website in search results, SEO may, for example, involve cross-linking between pages, adjusting the content of the website to include a particular keyword phrase, or updating content of the website more frequently. An automated SEO process may need to scrape search results from a search engine to determine how a website is ranked among search results.

[0005] In a second example, web scraping may be used to identify possible copyright. In that example, the scraped web content may be compared to copyrighted material to

- 2 -

automatically flag whether the web content may be infringing a copyright holder's rights. In one operation to detect copyright claims, a request may be made of a search engine, which has already gathered a great deal of content on the Internet. The scraped search results may then be compared to a copyrighted work.

[0006] In a third example, web scraping may be useful to check placement of paid advertisements on a webpage. For example, many search engines sell keywords, and when a search request includes the sold keyword, they place paid advertisements above unpaid search results on the returned page. Search engines may sell the same keyword to various companies, charging more for preferred placement. In addition, search engines may segment as sales by geographic area. Automated web scraping may be used to determine ad placement for a particular keyword or in a particular geographic area.

[0007] In a fourth example, web scraping may be useful to check prices or products listed on e-commerce websites. For example, a company may want to monitor a competitor's prices to guarantee that their prices remain competitive.

[0008] To conduct web scraping, the web request may be sent from a proxy server. The proxy server then makes the request on the web scraper's behalf, collects the response from the web server, and forwards the web page data so that the scraper can parse and interpret the page. When the proxy server forwards the requests, it generally does not alter the underlying content, but merely forwards it back to the web scraper. A proxy server changes the request's source IP address, so the web server is not provided with the geographical location of the scraper. Using the proxy server in this way can make the request appear more organic and thus ensure that the results from web scraping represent what would actually be presented were a human to make the request from that geographical location.

[0009] Proxy servers fall into various types depending on the IP address used to address a web server. A residential IP address is an address from the range specifically designated by the owning party, usually Internet service providers (ISPs), as assigned to private customers. Usually a residential proxy is an IP address linked to a physical device, for example, a mobile phone or desktop computer. However, businesswise, the blocks of residential IP addresses may be bought from the owning proxy service provider by another company directly, in bulk. Mobile IP proxies are a subset of the residential proxy

- 3 -

category. A mobile IP proxy is one with an IP address that is obtained from mobile operators. Mobile IP proxies use mobile data, as opposed to a residential proxy that uses broadband ISPs or home Wi-Fi. A datacenter IP proxy is the proxy server assigned with a datacenter IP. Datacenter IPs are IPs owned by companies, not by individuals. The datacenter proxies are typically IP addresses that are not in a natural person's home.

[0010] Exit node proxies, or simply exit nodes, are gateways where the traffic hits the Internet. There can be several proxies used to perform a user's request, but the exit node proxy is the final proxy that contacts the target and forwards the information from the target to a user device, perhaps via a previous proxy. There can be several proxies serving the user's request, forming a proxy chain, passing the request through each proxy, with the exit node being the last link in the chain that ultimately passes the request to the target.

[0011] E-commerce and search engine sites may prefer not to service web scraping requests or may try to limit web scraping requests. To that end, these sites may try to determine which of the requests it receives are automated and which requests are in response to a human web browsing request. When a web server identifies a request that the server believes to be automated, the server may block all requests coming from that proxy or requests having certain parameters from that proxy.

[0012] To identify which requests are automated, a web server may try to determine whether web requests coming from a particular IP address or subnet satisfy a pattern over time. To avoid detection, proxies may be rotated so that no single IP address makes too many requests. However, the supply of proxy IP addresses is limited. The IP address space (especially in IP version 4) in general is constrained. This limited supply is exasperated because many of the available IP addresses are labeled as data center IPs, and many target websites likely to be scraped refuse to service web requests from those IP addresses. As a result of the limited supply, the cost of obtaining a proxy IP address has increased. Taking proxy IP addresses out of circulation too quickly raises the cost of web scraping.

[0013] Systems and methods are needed for improved web scraping.

- 4 -

BRIEF SUMMARY

- [0014]** In a first embodiment, a method is provided for throttling a plurality of web scraping requests over time. In the method, during a first time period, a plurality of web scraping requests are received from a client computing device. Each of the plurality of web scraping requests specify a target website from which to capture content. Job descriptions for the plurality of web scraping requests are placed on a queue. Each job describes a web scraping job specified by a request from the plurality of requests. The job descriptions are removed from the queue incrementally over a second time period longer than the first time period. When each respective job description is removed from the queue, a web request for the target website is generated as specified in the job description. The web request is transmitted such that the web request reaches the target website via a proxy selected from a group of proxies.
- [0015]** In a second embodiment, a method is provided for tracking user activity based on an amount of data exchanged. In the method, a web scraping request from a client computing device is received. The web scraping request specifies a target website from which to capture content. Based on the web scraping request, a web request for the target website is generated. The web request is transmitted such that the web request reaches the target website via a proxy selected from a group of proxies. In response to the web request, content transmitted from the target website is received via the proxy. An amount of data in the received content is counted to determine a current traffic total for a client of the client computing device. The received content is transmitted to the client computing device.
- [0016]** In a third embodiment, a method is provided for determining which servers are available to process web scraping jobs. In the method, health of each of a plurality of database servers is repeatedly checked. Based on the health checks, whether each of a plurality of database servers are to be enabled or disabled is determined. The plurality of database servers operate independently of one another. Each database server is configured to manage data storage to at least a portion of a job database that stores the status of web scraping jobs while they are being executed. When a web scraping request is received from a client computing device, one of the database servers identified as

- 5 -

enabled is selected. And, a job description specified by the web scraping request is sent to the selected database server for storage in the job database as a pending web scraping job.

[0017] In a fourth embodiment, a method is provided for distributing application programming interface (API) requests over a plurality of data centers in a web scraping system to ensure that the API requests are routed to the data center with the requested data. In the method, an API request from a client computing device is received. The API request relates to a web scraping job that was previously requested of a web scraping system. A hostname addressing a data center of the web scraping system that serviced the web scraping job is identified. The API request is sent to the identified hostname. Finally, the API request is serviced at the identified data center.

[0018] In a fifth embodiment, a method is provided for securing a web scraping system. At an entry point to the web scraping system, credentials received with an API request from a client computing device are validated. The API request asks that the web scraping system scrape a target website. When the credentials are validated, a token indicating an identity of a client associated with the credentials is generated. The API request along with the token are transmitted to a server configured to initiate a scraping process on the web scraping system. At the server configured to initiate the web scraping system, the token is analyzed to determine whether the client is authorized to conduct the request. When the client is authorized, the web scraping system scrapes the target website.

[0019] In a sixth embodiment, a method is provided for aggregating web results scraped from a plurality of data centers. In the method, at respective first and second data centers of the plurality of data centers, requests for a target website are generated and transmitted. In response to the web requests, content from the target website is received at each of the data centers. The content received at the respective data centers is stored in an archive. Metadata for the respective content are stored in entries in index files at the respective data centers. A search request from a client is analyzed to determine whether search criteria in the search request describes the metadata in the index files across the various data centers. Based on the analysis, a file with links to the content is provided to the client.

[0020] In a seventh embodiment, a method is provided for scraping web results via a proxy. In the method, an API request is received from a client computing device. The

- 6 -

API request specifies a target website to scrape. Based on the target website, a scraper specifying how to generate a sequence of HTTP requests to access content requested by the API request is selected. Based on the selected scraper and received API request, an earlier HTTP request in the sequence of HTTP requests is generated. The earlier HTTP request is transmitted, to the target website via a proxy selected from a group of proxies. A response to the earlier HTTP request is received from the target website via the proxy. The response includes an instruction to save session data in, for example, a cookie. Based on the selected scraper and received API request, a later HTTP request in the sequence of HTTP requests is generated such that the later HTTP request includes the session data. A response to the later HTTP request is received from the target website and via the proxy. The response includes the scraped web results. Finally, the scraped web results are transmitted to the client computing device.

[0021] System and computer program product embodiments are also disclosed.

[0022] Further embodiments, features, and advantages of the invention, as well as the structure and operation of the various embodiments, are described in detail below with reference to accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0023] The accompanying drawings, which are incorporated herein and form part of the specification, illustrate the present disclosure and, together with the description, further serve to explain the principles of the disclosure and to enable a person skilled in the relevant art to make and use the disclosure.

[0024] FIG. 1 is an architecture diagram illustrating a system that allows a client to scrape web content through a proxy.

[0025] FIG. 2 illustrates an example operation of the system in Figure 1 and provides further description of how the system's components may interact.

[0026] FIG. 3 is an architecture diagram illustrating aspects of a web scraping system in greater detail.

[0027] FIG. 4 is an architecture diagram showing components of a web scraping system in greater detail.

- 7 -

- [0028] FIG. 5 illustrates a web scraping system with a throttler, according to an embodiment.
- [0029] FIG. 6 is a diagram illustrating how web requests can be throttled and trickled out over time.
- [0030] FIG. 7 is a system diagram illustrating how the throttling mechanism can be implemented in a web scraping system.
- [0031] FIG. 8 is a system diagram illustrating how traffic counting can be implemented in a web scraping system.
- [0032] FIG. 9 is a flowchart illustrating a method for counting and regulating traffic in a web scraping system.
- [0033] FIG. 10 illustrates operation of various components of a web scraping system to detect whether its various database servers should be enabled or disabled.
- [0034] FIG. 11 illustrates various components of a web scraping system to select an enabled database server to process a web scraping request.
- [0035] FIGS. 12A-B illustrate systems for load balancing web scraping requests that direct a request for results to the same subdomain that requested content from a target.
- [0036] FIGS. 13A-B illustrate methods for operation of the systems in FIGS. 12A-B.
- [0037] FIG. 14 illustrates a system for token-based authentication at a load balancer for a web scraping system.
- [0038] FIG. 15 illustrates token-based authentication for a web scraping system in greater detail.
- [0039] FIG. 16 illustrates a system that stores results from a plurality of data centers into an archive database.
- [0040] FIG. 17 illustrates a system that de-duplicates from a plurality of data centers before storing into an archive database.
- [0041] The drawing in which an element first appears is typically indicated by the leftmost digit or digits in the corresponding reference number. In the drawings, like reference numbers may indicate identical or functionally similar elements.

- 8 -

DETAILED DESCRIPTION

[0042] Embodiments disclose a system that allows for improved generation of web requests for scraping that, because of the nature of the requests and the time and manner in which they are sent out, appear more organic, as in “human generated,” than conventional automated scraping systems. In an example operation, a client may send a request to the system to scrape a particular site. Within the parameters the client provides, the system then manages how the request is made to the site, masking the request in a manner that makes it appear to the web server as if the request is not generated by an automated system. In this way, by making the request appear more organic, web servers may be less likely to block requests from the disclosed system or may take longer to block requests from the disclosed system. By avoiding web servers blocking requests and extending the lifetime of IP proxies before they are blocked, embodiments can use a limited IP proxy address space more efficiently. Aspects of a web scraping system that makes scraping requests through proxies are described with respect to FIGS. 1-4.

[0043] For various applications, clients may require automated scraping of web content. Some websites try to block requests that the websites believe to be automated. To determine whether requests are automated, these websites look at patterns of requests received over time from particular sources. For that reason, sources can be used to originate automated requests for web content only so frequently before the websites requested begin to recognize that the requests are automated. In part, to deal with this, embodiments throttle requests on a per-client basis. Throttling involves queuing requests from a client and executing requests from the queue on a periodic basis. In one example embodiment, a queue may be specific to a particular client, a target, and as to whether the request was made through a headless browser. Aspects of throttling web scraping requests are described with respect to FIGS. 5-7.

[0044] Traditionally, proxy web scraping services may have tracked user activity on by tracking a number of requests made for invoicing purposes. According to embodiments, user activity for proxy web based scraping services is tracked based on network traffic generated. In particular, when a response is received to a web scraping request, the amount of data transferred over a physical network layer may be calculated. Depending

- 9 -

on product, pricing plan and possibly other parameters, the request might be counted as well. To calculate the amount of data, an adjustment may be made to the total amount of data received to account for compression that occurs when the data is transferred. In addition, the amount of data may exclude requests and responses needed to acquire session information to request the scraped traffic. Aspects of tracking network traffic generated by web scraping requests are described with respect to FIGS. 8-9.

[0045] Requests for web scraping may be placed in a database while they are being executed. The database may hold status information pertinent to the request. To store and provision data, a plurality of database servers, sometimes called shards, may be used. When an update or query is made to the database, the update or query may be sent to the database server which stores the database for update or data queried. Sometimes, database servers can become overwhelmed or otherwise hung up. According to an embodiment, “heartbeats” are periodically sent to the database servers to determine which ones are available and which ones may have issues. A shared memory location may be used where the status of a shard could be indicated to be unavailable. Based on that information, new requests for web scraping may be sent to those database servers determined to be available and not to those determined to be unavailable. There is a shared memory location where the status of a shard could be indicated to be unavailable, but the “heartbeat” is still carried out by each of the user services. Aspects of this database server management technique are described below with respect to FIGS. 10-11.

[0046] A web scraping service may include a plurality of web scraping servers and a load balancing server to distribute load between the plurality of web scraping servers. The web scraping servers are configured to register and conduct web scraping requests. When a new request is received, the load balancing server determines which of the plurality of web scraping servers conduct the request. Each of the plurality of web scraping servers may be located in a different data center location. The new request for web scraping is directed to a data center determined by a load-balancing server. To retrieve the results of the web scraping, a client may send a subsequent request. According to an embodiment, the load balancing server may determine which of the plurality of data centers executed the original request. Then, the load balancing server may direct the subsequent request to the same data center. In this way, data can be more efficiently retrieved from the same

- 10 -

data center that serviced the web scraping request. Aspects of this direct data center access are described below with respect to FIGS. 12A-B and 13A-B.

[0047] A web scraping service may include a plurality of web scraping servers and a load balancing server. The web scraping servers are configured to register and conduct web scraping requests. When a new request is received, the load balancing server determines which of the plurality of web scraping servers conduct the request. Authentication needs to be conducted to determine whether the request is authorized. According to an embodiment, the load balancing server accesses an authentication server that assesses credentials provided with an API request to determine whether the credentials are valid. When the credentials are valid, the authentication server returns to the load balancing server, a token, indicating the user's identity and role. The load balancing server then forwards the token along with the request to the website scraping server, and the website scraping server uses the token to verify that the request is authorized. Aspects of this token-based authentication are described below with respect to FIGS. 14 and 15.

[0048] Web scraping services can keep a history of web pages previously scraped. The history may be stored by a plurality of data centers, each maintaining their own index. When a request to look up previously scraped webpages is received, the various indexes stored across multiple data centers are assessed and the results from the respective data centers are combined. In a further embodiment, web scraping results are de-duplicated across multiple data centers. This operation of an archive of web scraping results across multiple data centers is described with respect to FIGS. 16 and 17.

Automated web scraping through proxies

[0049] FIG. 1 is an architecture diagram illustrating a system 100 that allows a client to scrape web content through proxy. System 100 includes a client computing device 102, web scraping system 104, a plurality of web proxies 106A...N, and a target web server 108. Each of these components include one or more computing devices and are connected through one or more networks 110.

[0050] Client computing device 102 is a computing device that initiates requests to scrape content from the web, in particular target web server 108. As described above, client computing device 102 may seek to scrape content for various applications. For

- 11 -

example, client computing device 102 may have or interact with software to engage in search engine optimization. Client computing device 102 may be analyzing ad placement or e-commerce products or listed prices. Client computing device 102 sends a request to web scraping system 104. The request can be synchronous or asynchronous and may take a variety of formats as described in more detail with respect to FIG. 2.

[0051] Web scraping system 104 develops a request or a sequence of requests that impersonate a human using a web browser. To impersonate non-automated requests to a target website, web scraping system 104 has logic to formulate Hypertext Transfer Protocol (HTTP) requests to the target website. Still further, many of these sites require HTTP cookies from sessions generated previously. An HTTP cookie (usually just called a cookie) is a simple computer data structure made of text written by a web server in previous request-response cycles. The information stored by cookies can be used to personalize the experience when using a website. A website can use cookies to find out if someone has visited a website before and record data about what they did. When someone is using a computer to browse a website, a personalized cookie data structure can be sent from the website's server to the person's computer. The cookie is stored in the web browser on the person's computer. At some time in the future, the person may browse that website again. When the website is found, the person's browser checks whether a cookie for that website is found and available. If a cookie is found, then the data that was stored in the cookie before can be used by the website to tell the website about the person's previous activity. Some examples where cookies are used include shopping carts, automatic login, and remembering which advertisements have already been shown.

[0052] Additionally or alternatively, the second request may be generated from other data received in response to the first request, besides cookies. For example, the other data can include other types of headers, parameters, or the body of the response.

[0053] Because many websites require session information, usually stored in cookies but possibly received in other data from previously visited retrieved pages, web scraping system 104 may reproduce a series of HTTP requests and responses to scrape data from the target website. For example, to scrape search results, embodiments described herein may first request the page of the general search page where a human user would enter her

- 12 -

search terms in a text box on an HTML page. If it were a human user, when the user navigates to that page, the resulting page would likely write a cookie to the user's browser and would present an HTML page with the text box for the user to enter her search terms. Then, the user would enter the search terms in the text box and press a "submit" button on the HTML page presented in a web browser. As a result, the web browser would execute an HTTP POST or GET operation that results in a second HTTP request with the search term and any resulting cookies. According to an embodiment, the system disclosed here would reproduce both HTTP requests, using data, such as cookies, other headers, parameters or data from the body, received in response to the first request to generate the second request.

[0054] In embodiments, the system may be able to simulate a request sent from a particular geographic location. One way to do this is to select a proxy from a particular geographic location. However, one issue with this is that there may not be sufficient proxy IP addresses available in every location of interest. As proxy IP addresses become exhausted and are blocked by target websites, the supply of proxy IP addresses, in particular, geographic locations, may become limited or even exhausted.

[0055] Many sites provide another way for a browser to identify a geolocation in the HTTP request. In those sites, the IP address may not be the only source of location information. A browser may have settings and a user profile that enable a user to allow the browser to send location information. The location information is information describing the location of the device sending the HTTP request, for example, a GPS location detected by a GPS receiver on a mobile device. The browser may provide the location information in the HTTP request. In one example, the location may be embedded in header information in the HTTP request such as in a cookie.

[0056] The headers can also contain identification of the user agent. For example, they can include an identification of a type of browser being used. In this way, a web server can customize a response based on the user agent being used to access the content. For example, a site might look different on a mobile browser than on a desktop browser. Sometimes, websites may only support certain types of browsers. Browsers might have different features available to them, and when a website recognizes from the HTTP

- 13 -

header, information that the request came from a browser lacking the necessary features for the website, the site may refuse to respond with the requested page.

[0057] According to embodiments, web scraping system 104 can mimic this transmission of location information without using the actual browser. For example, the system can embed location information in headers of the HTTP requests sent to the target website. In this way, web scraping system 104 may be able to formulate a localized request, even when a proxy IP address located in the desired region may be unavailable.

[0058] Further description of how web scraping system 104 may formulate the series of requests needed to obtain the desired content is provided below in FIG. 2. Once web scraping system 104 formulates an HTTP request, it sends the request to a web proxy 106A...N.

[0059] Web proxy 106A...N are each a server that acts as an intermediary for requests from clients seeking resources from servers that provide those resources. Web proxy 106A...N thus functions on behalf of the client when requesting service, potentially masking the true origin of the request to the resource server.

[0060] In FIG. 1, at least one of web proxy 106A...N accepts a request from web scraping system 104. The request is in a proxy protocol format. The at least one web proxy 106A...N that accepts the request forwards the request to target web server 108. In forwarding the request, the web proxy 106A...N assigns its own IP address as the source IP address for packets encapsulating the HTTP request. The destination IP address for any packets encapsulating the HTTP request is the IP address of target web server 108. Because the source IP address was the web proxy 106A...N, target web server 108 may not be able to determine that the request originated from web scraping system 104. Once target web server 108 generates a response to the request, target web server 108 sends the response back to the web proxy 106A...N that forwarded the request, which in turn forwards the response to web scraping system 104.

[0061] Target web server 108 is computer software and underlying hardware that accepts requests and returns responses via HTTP. As input, target web server 108 typically takes the path in the HTTP request, any headers in the HTTP request, and sometimes a body of the HTTP request, and uses that information to generate content to be returned. The

- 14 -

content served by the HTTP protocol is often formatted as a webpage, such as using HTML and JavaScript.

[0062] FIG. 2 includes a diagram 200 that illustrates an example operation of system 100 in FIG. 1 and provides further description of how components of system 100 may interact.

[0063] As illustrated in diagram 200, to request that web content to be scraped, client computing device 102 sends a scrape request 202 to web scraping system 104. In response to scrape request 202, web scraping system 104 initiates a series of operations and transactions needed to scrape the results. In diagram 200, the transactions are labeled as messages 204A-D and 206A-D. After the results are scraped, web scraping system 104 provides the retrieved content back to client computing device 102 as scraped results 208.

[0064] In an embodiment, web scraping system 104 may cache scraped data. For example, when a request to scrape is received from a client, the system may determine whether the website targeted has been scraped within a certain time, such as within the last day. If the website has been scraped recently, e.g., within the certain time period, the system may return the previously scraped results.

[0065] Client computing device 102 interacts with web scraping system 104 in various ways. In an embodiment, a client may send scrape request 202 with the parameters describing the web scraping sought to be completed. The request and its parameters may conform to an API set forth by system 104. The parameters may include a Uniform Resource Locator (URL), Uniform Resource Identifier (URI), header information, geolocation information, and browser information, and other values listed in greater detail below with respect to FIG. 3. In response to the request, system 104 will return an acknowledgment that the request is received. The acknowledgment may include a message indicating that the scraped results will be available at a particular location. The system queues the request and, when the scraped results are retrieved, a message, also called a callback, may be sent to client 102 indicating that scraped results 208 are available. In this way, system 104 can asynchronously service a client request for the scrape data.

[0066] Alternatively or additionally, client 102 may send the request, as described above, but instead of returning an acknowledgment, system 104 may keep the connection with

- 15 -

client 102 open while the scraping is being conducted. Once the scraping is completed, the results are returned in a response to the initial request. In this way, system 104 can synchronously service a client request for the scrape data.

[0067] Alternatively or additionally, the request may service requests formatted according to a proxy protocol. The proxy protocol may be the same protocol used by proxy server 106 to accept and respond to requests. In diagram 200, such requests are labeled proxy protocol request 204A and 206A and proxy protocol responses 204D and 206D. Examples of a proxy protocol include the HTTP proxy protocol and a SOCKS protocol.

[0068] When web scraping system 104 receives scrape request 202, web scraping system 104 determines the necessary transactions needed to scrape the requested content. As mentioned above, because of state variables expected from target website 108, to scrape the requested content, multiple transactions with target website 108 may be necessary. In the example in diagram 200, two transactions are shown, but more or less may be needed. For each transaction, web scraping system 104 needs to determine all the parameters for an HTTP request and format the HTTP request accordingly.

[0069] The logic necessary to determine the series of the transactions and corresponding HTTP requests may be preprogrammed for each target website. When scraping API request 202 is received, web scraping system 104 looks at the needed logic and generates the transactions and corresponding HTTP request accordingly. In addition, the logic describing the sequence of HTTP requests to reach a target website may be stored in advance. When a customer request is received to scrape a target website, the logic necessary to reach the target website is retrieved and executed, resulting in a series of HTTP requests being sent and HTTP responses being received.

[0070] In diagram 200, the HTTP requests to target website 108 are HTTP request 204B and HTTP request 206B. As described above, web scraping server 104 may not send the requests directly to target website 108 and instead send them through at least one intermediary proxy server 106. To send the request to proxy server 106, a proxy protocol may be used.

[0071] To send a request according to an HTTP proxy protocol, the full URL may be passed, instead of just the path. Also, credentials may be required to access the proxy. All

- 16 -

the other fields for an HTTP request must also be determined. To reproduce an HTTP request, web scraping system 104 will generate all the different components of each request, including a method, path, a version of the protocol that the request wants to access, headers, and the body of the request.

[0072] An illustrative example of proxy protocol request 204A is reproduced below:

GET https://www.searchengine.com/ HTTP/1.1

Proxy-Authorization: Basic encoded-credentials

Accept: text/html

User-Agent: Mozilla/5.0

Cookie: Location=Alexandria, VA, USA;

[0073] In the above example, the HTTP method invoked is a GET command, and the version of the protocol is “HTTP/1.1.” The path is “https://www.searchengine.com/,” and because it includes a full URL as opposed to URI, it may signify to web proxy 106 that the HTTP request is for a proxy request. The body of the request is empty.

[0074] The example HTTP proxy protocol request 204A includes four headers: “Proxy-Authorization,” “Accept,” “User-Agent,” and “Cookie.” The “Proxy-Authorization” header provides authorization credentials for connecting to a proxy. The “Accept” header provides media type(s) that is/are acceptable for the response. The “User Agent” header provides a user agent string identifying the user agent. For example, the “User Agent” header may identify the type of browser and whether or not the browser is a mobile or desktop browser. The “Cookie” header is an HTTP cookie previously sent by the server with Set-Cookie (below). In this case, the server may be set up to previously have saved the location of the user. Thus, if the user had previously visited the server from Alexandria, Virginia, the server would, for example, save “Alexandria, VA, USA” as a cookie value. By sending such a cookie value with the request, web scraping system 104 can simulate the geolocation without having previously visited the location and without needing a proxy IP address located in Alexandria, Virginia.

[0075] System 104 may profile these values, the HTTP headers and the path (e.g., “https://www.searchengine.com/” v. “https://www.searchengine.com/en/”), to resemble requests that would be plausibly generated by a browser controlled by a human. In this

- 17 -

way, system 104 may generate the HTTP requests to avoid the target web server being able to detect that the requests are automatically generated from a bot.

[0076] The HTTP proxy protocol is one example of how the proxy protocol may operate. In another example, HTTP tunneling may be used, using, for example, the HTTP CONNECT command. In still another example, the proxy may use a SOCKS Internet protocol. While the HTTP proxy protocol operates at the application layer of the OSI model protocol stack, SOCKS may operate at the session layer (layer 5 of the OSI model protocol stack). Other protocols may be available forwarding data at different layers of the network protocol stack.

[0077] When proxy server 106 receives proxy protocol request 204A, it sends a corresponding HTTP request 204B to target website 108. Continuing the example above, the target website 108 may be a web server addressed by the hostname “www.searchengine.com.” In this example, HTTP request 204B may be:

GET / HTTP/1.1

Accept: text/html

User-Agent: Mozilla/5.0

Cookie: Location=Alexandria, VA, USA;

[0078] In response, target website 108 will return an HTTP response with the website located at its path “/”. The website may be customized according to the headers, for example, it may be customized according to the location information and user agent information in the headers of HTTP request 204B.

[0079] Continuing the example above, HTTP response 204C may be as follows:

HTTP/1.1 200 OK

Content-Type: text/html

Set-Cookie: Location=Alexandria, VA, USA; session-id=“1234567890abcdef”

<!DOCTYPE html>

<body>

<h1>Search Query</h1>

<form action=“/search”>

<input type=“text” id=“query” name=“query”>

<input type=“submit” value=“Submit”>

- 18 -

```
</form>
</body>
</html>
```

[0080] In that example, again, HTTP response 204C can include an indication of the version of the protocol used (here, “HTTP/1.1”), a status code indicating if the request is successful or not (here, “200”), a status message with a non-authoritative short description of the status code (here, “OK”). HTTP response 204C can also include headers and a body.

[0081] Here, the headers include a “Content-Type” header which includes a MIME (Multipurpose Internet Mail Extensions) type of this content and a “Set-Cookie” header that sets cookies to be transmitted with the next request to target website 108. In this case, the “Set-Cookie” header keeps the geolocation information previously transmitted and adds a “session ID.” Target website 108 may expect the session ID to be returned to it in further requests. In this way, target website 108 can track sessions and maintain the state across multiple request-response cycles.

[0082] The body of HTTP response 204C in the example above is a simple HTML page. If the HTML page were rendered in a browser, the browser would present a form providing a search text box field and a submit button. When the submit button is selected, the browser would submit another request to the URI “/search” carrying, as is parameter, the input in the text field.

[0083] Target website 108 sends HTTP response 204C to proxy server 106. In response to receipt of HTTP response 204C, proxy server 106 sends a proxy protocol response 204D. Proxy protocol response 204D may be formatted according to the proxy protocol as described above. In an example, the proxy protocol response 204D may be similar or identical to HTTP response 204C.

[0084] On receipt of proxy protocol response 204D, web scraping system 104 then generates the next request in the sequence or requests needed to retrieve the requested content. In diagram 200 the next request is proxy protocol request 206A. Proxy protocol request 206A is generated to simulate a request that would be generated had a user on a browser entered in the desired search terms in the search field and pressed submit.

- 19 -

[0085] Continuing the example above, scrape request 202 may instruct web scraping system 104 to scrape search results for the search term “CompanyA.” Moreover, web scraping system 104 will generate proxy protocol response 206A to include the cookie and session information provided from the previous transaction. In that case, proxy protocol response 206A may be as follows:

GET https://www.searchengine.com/search?query=CompanyA HTTP/1.1

Proxy-Authorization: Basic encoded-credentials

Accept: text/html

User-Agent: Mozilla/5.0

Cookie: Location=Alexandria, VA, USA; session-id=“1234567890abcdef”

[0086] The path specified proxy protocol response 206A includes a complete URL of the destination sought from target website 108. It includes target website 108’s hostname “www.searchengine.com.” It also includes the path within target website 108 that is sought in this case “/search?query=CompanyA.” The first part of the path “/search” is specified by the submit button in the HTML form included in HTTP response 204C. The name of the parameter submitted “query” is also specified by the text field element in the HTML form included in HTTP response 204C. As described above with respect to proxy protocol request 204A and proxy protocol response 204D, this is one example; other protocols to establish a proxy relationship may be used.

[0087] When proxy server 106 receives proxy protocol request 206A from web scraping system 104, it generates a corresponding HTTP request 206B and sends HTTP request 206B to target website 108. Proxy server 106 generates HTTP request 206B in a similar manner to what was described above with respect to HTTP request 204B. Continuing the running example, HTTP request 206B sent to target website 108 may be:

GET /search?query=CompanyA HTTP/1.1

Proxy-Authorization: Basic encoded-credentials

Accept: text/html

User-Agent: Mozilla/5.0

Cookie: Location=Alexandria, VA, USA; session-id=“1234567890abcdef”

[0088] Based on HTTP request 206B, target website 108 generates HTTP response 206C. Target website 108 may generate a page corresponding to the path requested,

- 20 -

including the parameters included in the path (here, “/search?query=CompanyA”). Commonly, target website 108 will generate such a page dynamically checking databases, perhaps running computations, and formatting results accordingly. Continuing the example above, HTTP response 206C may be:

HTTP/1.1 200 OK

Content-Type: text/html

<!DOCTYPE html>

<body>

<h1>Search Results</h1>

<p> CompanyA </p>

<p> CompanyB </p>

<p> CompanyC </p>

</body>

</html>

[0089] Similar to HTTP response 204C, HTTP response 206C includes the protocol version, a status code, a status message, headers, and a body. The body of HTTP response 206C in the example above is a simple HTML page. If the HTML page were rendered in a browser, the browser would present listing of search results, with each search result having a corresponding link. A result “CompanyA” linked to the URL “https://www.companya.com” would be presented first. This HTML page is the content sought to be scraped by client computing device 102.

[0090] Target website 108 returns HTTP response 204C to proxy server 106. As described above with respect to proxy protocol response 204D, proxy server 106 generates proxy protocol response 206D based on HTTP response 206C. Proxy protocol response 206D includes the scraped content. Proxy server 106 sends proxy protocol response 206D to web scraping system 104.

[0091] Based on and in response to proxy protocol response 206D, web scraping system 104 generates and sends scraped results 208 to client computing device 102. As described above, web scraping system 104 may send the scraped content, in this case the HTML page within HTTP response 206C, directly to client computing device 102 within the scraped results 208. Alternatively or additionally, web scraping system 104 may parse

- 21 -

relevant information from the HTML page and send the parsed information to client computing device 102 in scraped results 208.

Components and Operation of an Automated Web Scraping System

[0092] FIG. 3 is an architecture diagram illustrating aspects of web scraping system 104 in greater detail. Web scraping system 104 includes a request intake manager 302, request status manager 304, history archive 306, parser 308, scraper 310, and message queue manager 312. Each of these components is described in turn.

[0093] Request intake manager 302 receives and processes incoming requests from client computing device 102. As described above, client computing device 102 can provide requests to request intake manager 302 synchronously and asynchronously. Client computing device 102 can provide the requests using a specialized API or using a proxy protocol.

[0094] The API call may include any of the following values as parameters:

- Data source specifying the target website to retrieve, such as a particular search engine or e-commerce website;
- Domain localization specifying the domain name system localization, which may specify the top level domain name of the website to request (".com," ".fr", ".de", etc.);
- a keyword to search on the search engine or e-commerce site;
- a number indicating how many results to retrieve from the target website;
- for scraped results spanning multiple pages, a starting page number and number of pages to retrieve, which together specify the page ranges to scrape;
- Accept-Language HTTP header value to use, this header value may signal to the target website the language to return results in;
- a geographical location to mimic when making the request;
- a device type and browser to mimic when making the request;
- a flag to signal whether web scraper system 104 should parse the HTTP response; and
- a flag to signal whether web scraper system 104 should signal to the target website to AutoCorrect spelling;

- 22 -

- a flag to signal whether web scraper system 104 should signal in its cookies to use a particular results language;
- any other parameters that might be expected by a target website, such as sorting and filtering methods to be used in conducting the search; and
- user credentials, such as username/password or a token.

[0095] In addition to the above parameters, when the API request is a callback request, the API request provided to request intake manager 302 may include a network location (such as a URL) for web scraping system 104 to provide the callback message when the scraping request has been completed and results are available.

[0096] Request intake manager 302 may be configured to take parameters in multiple different formats and may be configured to accept requests using multiple different parameter names. In that embodiment, request intake manager 302 may normalize the different formats and parameter names into a common format before further processing. When request intake manager 302 accepts such a different format from a client computing device, it may remember that the client computing device used that different format and translate any results returned into a format expected by the client computing device.

[0097] When request intake manager 302 receives an API request, request intake manager 302 may first validate the request to ensure that the API request conforms to a format that is understood by web scraping system 104 and that the parameters for the API request specify web scraping can be executed by web scraping system 104. When request intake manager 302 confirms that an API request is valid, request intake manager 302 may send the API request to request status manager 304 to manage processing of the request.

[0098] For an asynchronous API request, request status manager 304 may return to request intake manager 302 a job identifier that client computing device 102 may later use to retrieve the results of the request. Request intake manager 302 may then send that job identifier to the client computing device 102 for storage and later use when the request is completed.

[0099] For a synchronous API request, request status manager 304 may hold the connection with client computing device 102 open until web scraping system 104

- 23 -

completes the scraping requests and results are available. When results are available, request status manager 304 will respond to the API request with the scraped results, closing the connection.

[0100] Once request intake manager 302 sends the API request to request status manager 304, request status manager 304 manages processing of the request. Request status manager 304 may save the job description and any return results corresponding to the job description in job database 314. As will be discussed below with respect to FIG. 4, request status manager 304 may include a plurality of servers that operate independently of one another. In particular, the plurality of servers may not know about each other and may not communicate any state to one another.

[0101] As will be described in detail below with respect to FIG. 4, request status manager 304 and request intake manager 302 work together to save a description of the job to be processed and an indication of the job status in job database 314.

[0102] Once the job description and job status are saved in job database 314, request status manager 304 sends the job description to central message queue 312. Central message queue 312 may be a message broker software application. A message broker software application is an intermediary computer program module that translates a message from a formal messaging protocol of the sender to a formal messaging protocol of the receiver. In different examples, the messaging protocols can include the Advanced Message Queuing Protocol (AMQP), Streaming Text Oriented Messaging Protocol (STOMP), and MQ Telemetry Transport (MQTT). Central message queue 312 provides communication among components of web scraping system 104, including request status manager 304, history archive 306, parser 308, and scraper 310. By mediating communication among the various components, central message queue 312 minimizes the mutual awareness that the various components need one another. This encourages decoupling of the components.

[0103] Central message queue 312 may include a first-in-first-out queue for job descriptions. Scraper 310 may monitor central message queue 312, taking jobs off the central message queue 312 and conducting the scraping from target website 108 accordingly. As mentioned above, some jobs may represent a synchronous API request while others may be asynchronous. Scraper 310 may prioritize the synchronous jobs

- 24 -

above the asynchronous jobs to help ensure that results are determined before web scraping system 104's connection with the client computing device 102 times out.

[0104] As described above, based on the job description retrieved from central message queue 312, scraper 310 may formulate a series of HTTP requests from target website 108 to ultimately retrieve results as specified in the request, such as a desired HTML page.

[0105] Once scraper 310 has retrieved the results, scraper 310 may place the retrieved HTML along with an identification of the job on to central message queue 312. Request status manager 304 may retrieve the results from central message queue 312 and placement job database 314. Alternatively or additionally, scraper 310 may send the results directly back to request status manager 304, which places the results in job database 314 and sends a link to the results as will be described below with respect to FIG. 4.

[0106] If the API request was strictly requesting HTML, request intake manager 302 may then return the scraped HTML to the user. As described above, if the request from client computing device 102 is synchronous, request intake manager 302 may send the scraped HTML to client computing device 102 in response to the request. Alternatively, if the request is asynchronous, request intake manager 302 may send a notification to client computing device 102 that the results are available and client computing device 102 may retrieve the scraped HTML from job database 314 via request intake manager 302.

[0107] As mentioned above, alternatively or additionally to the HTML results, client computing device 102 may request parsed results. Parsing of the scraped HTML may be executed by parser 308. Parser 308 may analyze the scraped HTML file and may extract relevant fields from the HTML file. To analyze the HTML file, parser 308 may use a known format or patterns within the HTML file (such as the Document Object Model) to identify where the relevant fields are located. With the relevant fields extracted, parser 308 may insert the extracted fields into a new data structure, such as a file. In an example, the new file may be a JavaScript Object Notation (JSON) format, which is a standard data interchange format. The resulting file with the parsed data may be stored in job database 314 and returned to client computing device 102.

[0108] Retrieved results may also be stored to history archive 306. History archive 306 is an archival, or cold database service. History archive 306 stores the scraped data for

- 25 -

longer than job database 314. It is not meant to represent current content from a target website, instead representing historical content. While job database 314 may erase results after a certain time period, such as 24 hours, history archive 306 may store results indefinitely. In one embodiment, history archive 306 may store parsed scraped data but not HTML data because HTML data has structure and formatting that may not be relevant to a client. When the parsed data is stored, the job description may be also stored and used as metadata in an index to allow the parsed data to be searched. In the event that a client makes an identical request twice, the results may only be stored in history archive 306 if the results from the first request are older than a certain age, such as one month.

[0109] FIG. 4 is an architecture diagram showing components of web scraping system 104 in greater detail.

[0110] As mentioned above, request status manager 304 includes a plurality of database servers 404A...N that operate independently of each other. In particular, the plurality of servers may not know about each other and may not communicate any state to one another.

[0111] In an embodiment, each of the plurality of servers 404A-N may be referred to as a database shard. A database shard, or simply a shard, is a horizontal partition of data in a database or search engine. Each shard is held on a separate database server instance, to spread load. Some data within a database may remain present in all shards, but some appears only in a single shard. Each shard acts as the single source for this subset of data. Horizontal partitioning is a database design principle whereby rows of a database table are held separately, rather than being split into columns (which is what normalization and vertical partitioning do, to differing extents). Each partition forms part of a shard, which may in turn be located on a separate database server or physical location.

[0112] Horizontal partitioning may be advantageous because tables are divided and distributed across multiple servers. Since the tables are divided and distributed into multiple servers, the total number of rows in each table in each database is reduced. This reduces index size, which may improve search or query performance of the database. Also, a database shard can be placed on separate hardware, and multiple shards can be placed on multiple machines. This enables a distribution of the database over a large number of machines, improving performance.

- 26 -

- [0113] In another embodiment, instead of independent database servers, database clustering may be used to spread functionality across multiple servers or machines. With database clustering, the various servers must be aware of each other. With database clustering, the database may not be horizontally partitioned across the multiple servers. However, with clustering, the state may need to be synchronized between the various servers. For example, two servers may not be able to update the same row at the same time. While this may have advantages in avoiding the need to partition rows of the database, synchronizing state and avoiding collisions, it may slow performance.
- [0114] Because the plurality of servers operate independently of one another, request intake manager 302 may need to decide which of the plurality of servers to use and remember which server is processing each job to ensure that information relating to the job, such as scraped results, is sent to the proper server. That is, the database instance that is responsible for managing the API request. Moreover, to realize performance advantages, request status manager 304 may be able to determine how the data is segmented across the multiple shards, thus avoiding the need to query every shard to determine a jobs status or retrieve job information, such as scraped results.
- [0115] In one embodiment, a job identifier encodes an identification of the shard that is servicing a web scraping request. The job identifier is sent to a client computing device when the request is created. And, when the client computing device seeks to retrieve status information or results for a previous request, the client computing device may use the job identifier it received for that previous request. Request status manager 304 may be able to decode the job identifier to determine which shard is handling the request.
- [0116] In-memory job table 452 may be used to store and retrieve status information. The in-memory job table 452 may be stored in a volatile (but also non-transitory) memory. Because the memory is volatile, data may be lost if power is disconnected. But volatile memory may, in general, have faster retrieval times than persistent memory. By storing the mapping table in memory, embodiments may provide faster access times. To store job information, each scraping job needs a unique job identification number. In particular, the job identification number needs to be unique across all the database shards.
- [0117] To access in-memory job table 452 (both reading and writing), request intake manager 302 may send a request to in-memory access server 450. In-memory access

- 27 -

server 450 may act as a database server that controls and allows access (and perhaps performs other memory management functions) for in-memory job table 452.

[0118] Similar to how request status manager 304 includes multiple servers operating independently of one another, request intake manager 302 may also include multiple servers operating independently of one another. In FIG. 4, these multiple servers are illustrated as API receivers 402A...N. Each of API receivers 402A...N may perform the functions described above for request intake manager 302, including receiving an API request from client computing device 102, validating the API request, sending information describing the scraping job to request status manager 304, and (for a call back request) returning to client computing device 102 data necessary to retrieve the scraped results when they are available.

[0119] Having multiple API receivers 402A...N allows the receivers to spread service API requests in parallel and spread computing load across multiple computing resources and machines. Yet, having multiple API receivers 402A...N operating independently without sharing state or being aware that other API receivers exist may add complexity in coordinating functionality.

[0120] For example, because both the request intake manager 302 and request status manager 304 include multiple servers operating independently of one another, they may not be aware of job identification numbers that have already been assigned. This can lead to difficulties in determining a unique job identification number associated with a new scraping job.

[0121] To determine a job identification number, an API receiver 402A...N may first determine to which of the plurality of database servers 404A...N to send the scraping job. That determination is described in greater detail with respect to FIGS. 10-11. API receivers 402A...N may then retrieve a next value from a counter maintained by each respective database servers 404A...N. The counter may reset at every time interval, such as every millisecond. API receiver 402A...N may concatenate an identification of database servers 404A...N with a timestamp and the next value retrieved from in-memory access server 450. In this way, API receiver 402A...N may generate a unique job identifier.

- 28 -

[0122] In addition to generating the unique job identifier, on receipt of an API request from client computing device 102, API receivers 402A...N may validate all the parameters received from client computing device 102, convert them into a format for internal processing by request status manager 304, and send the converted and validated parameters along with the unique job identifier (together referred to as a job description) to the selected database server 404A...N.

[0123] In particular, the API receivers 402A...N may send the job description to a message queue 454 implemented within the selected database server 404A...N. Once the job description is placed into the appropriate message queue 454, the API receiver may (in the case of an asynchronous request) return the job description to client computing device 102 along with other links on how to retrieve status and how to retrieve results. At that point, for an asynchronous request, API receivers 402A...N will close the connection with client computing device 102. For a synchronous request, the connection with computing device 102 will stay open until results are available.

[0124] Message queue 454 may be a first-in-first-out queue. It may be implemented as a message broker as described above with respect to central message queue 312. Database micro service 456 may be monitoring message queue 454. Database micro service 456 may periodically or intermittently retrieve a next queued job description from message queue 454 and place the job description and an indication of the job's status (such as "pending") in job database 458. In addition, either by database micro service 456 or another component, the job description and status may be cached at in-memory job table 452.

[0125] According to an embodiment, a client computing device 102 may want to inquire of system 104 a job's status. In some applications, client computing device 102 may frequently send system 104 new scraping jobs and frequently ask system 104 of the status of previously sent scraping jobs while they are still in progress. The status requests can quickly add up.

[0126] To deal with such requests, when a client tries to retrieve status of a job, the request from client computing device 102 may be sent to in-memory access server 450 (via request intake manager 302), which retrieves the status information from in-memory job table 452. By caching the status information in job table 452 and using the cached

- 29 -

information to service client requests for status, embodiments may relieve job database 314 from performing that task, improving overall performance. Also, the status information may be retrieved more quickly from in-memory job table 452 than from job database 314, improving latency and processing the client request for a scraping job's current status.

[0127] Once the job description and job status are saved in job database 314, request status manager 304 sends the job description to central message queue 312. In particular, database micro service 456 places the job description again on message queue 454 implemented in the previously selected database server 404A-N. From message queue 454, the job description is forwarded to central message queue 312.

[0128] As described above, scraper 310 retrieves and executes the job from central message queue 312. Scraper 310 retrieves the results and enters them into a message queue 468 that is local to scraper 310. Scraper 310 includes a plurality of target scrapers 466A...N. Each of the target scrapers 466A...N is configured to scrape a particular target website. For example, target scraper 466A may include instructions to generate a series of HTTP requests to scrape search results from a particular search engine; target scraper 466B may include instructions to generate a series of HTTP requests to scrape search results from a particular e-commerce site; etc. One or more of the target scrapers 466A...N may be generic, meaning that they are configured to scrape any website, but are not particularly adapted to the sessions expected of a particular site. Based on the job description (and the customer API request upon which it is based), one of the target scrapers 466A...N is selected to generate and transmit the HTTP requests to a proxy server 106.

[0129] As discussed above, to generate the HTTP requests so that they appear organic (that is, appear to be generated from a user's browser and not from a bot), various determinations must be made. First, the proxy to be used must be selected. This selection may be done by proxy rotator 486. The selected target scraper 466A...N may request an appropriate proxy from proxy rotator 486. Second, the selected target scraper 466A...N must determine the appropriate headers to use in the HTTP request, including any cookies needed. Third, selected target scraper 466A...N must generate the path for the HTTP request itself.

- 30 -

[0130] Proxy rotator 486 determines the appropriate proxy to which to send the request. Proxy rotator 486 may track which proxies are available to use and which have been blocked. For example, proxy rotator 486 may track what responses a target website has previously returned to the proxy in response to previous requests from scraper 310. To track this information, scraper 310 may report to rotator 486 the status of the request through the proxy, as rotator 486 may not monitor the connections between the proxies and scrapers. If a response was an error or timeout, proxy rotator 486 may identify that proxy is blocked by that target website. Moreover, proxy rotator 486 may track which proxies have been used recently or heavily. When a proxy has been used recently or heavily, proxy rotator 486 may select a different proxy to allow the recently or heavily used proxy to cool down, extending its life before a target website might block it.

[0131] As mentioned above for FIG. 2, the selected target scraper 466A...N may determine many of the headers based on the user parameters from the API, which are saved in the job description that the selected target scraper 466A...N receives from central message queue 312. These headers include the user agent to specify the request, such as a particular desktop or mobile browser. If the API request does not specify a particular user agent, the selected target scraper 466A...N may select a user agent to specify from a predetermined list. As discussed above with respect to FIG. 2, cookies may require that the selected target scraper 466A...N make multiple successive requests to a target website to get the needed session information.

[0132] Finally, the selected target scraper 466A...N may determine a path of the HTTP request. Depending on how the client provided its scraping API request, the client may have specified the full URL to scrape or the client may have merely provided parameters for scraping. If the client provided the full URL, the full URL path may be used. Alternatively, if the client merely provided parameters, the parameters may be incorporated into a path in the format expected by a target website. As mentioned above, a client can specify a geolocation to spoof in the request. The geolocation may specify certain parameters that the selected target scraper 466A...N includes in the path. Even if the client specifies the full URL, the selected target scraper 466A...N may need to modify the URL or URI to impersonate that location.

- 31 -

- [0133] With the HTTP request(s) generated, the selected target scraper 466A...N sends the request to a proxy and retrieves the results from the target website via the proxy as discussed above with respect to FIG. 2. The resulting page typically includes HTML. The HTML may include links to other objects, such as images and widgets to display and interact with things like geographic maps (perhaps retrieved from a third party web service). In addition, the HTML may include JavaScript that has some functionality requiring execution to render. In some cases, a client may be interested in aspects of the page not represented in the HTML. In this case, the selected target scraper 466A...N may interact with a headless web browser (not shown) that has the necessary functionality to execute the JavaScript and retrieve any objects linked within the HTML. In this way, the headless web browser can develop a full rendering of the scraped webpage, or at least retrieve the information that would be needed to develop the full rendering.
- [0134] When the selected target scraper 466A...N sends the request to a proxy, an error could occur making the scraper unable to retrieve the results. For example, the selected target scraper 466A...N may recognize the response received from the proxy as incomplete, blocked with errors (such as an error status code in the HTTP response), or as having other problems. Depending on the type of error, the selected target scraper 466A...N may try to retrieve results again. In that case, the selected target scraper 466A...N will place the job description as a new message back on central message queue 312. When scraper 310 extracts that message from central message queue 312, scraper 310 will try to generate a new HTTP request. In generating the new HTTP request, proxy rotator 486 will determine a new proxy to which to send the new HTTP request, increasing the likelihood that the target website will service the new request and return the desired results.
- [0135] When the results are returned, target scraper places the results on message queue 468. From there, the results are forwarded to message queue 454. Database microservice 456 retrieves the results from message queue 454. Then, database microservice 456 stores the results in job database 314.
- [0136] As mentioned above, the message forwarded from message queue 454 to database microservice 456 may not include the actual HTML scraped. Instead, scraper 466 places the HTML in an HTML cache 462. HTML cache 462 may be a distributed memory-

- 32 -

caching system. Instead of storing the HTML in any message queue, the HTML may be stored in HTML cache 462 and the message queued and transmitted to database microservice 456 may be a link to HTML cache 462. When database microservice 456 receives a message indicating that results have been retrieved, database microservice 456 uses the link to retrieve the results from HTML cache 462.

[0137] Bypassing message queues in this way may be advantageous because having a large amount of HTML within a message may slow down operation of a message queue, such as message queue 454. This is because performance of message broker software applications tends to be sensitive to the size of the messages being sent.

[0138] With the results returned to database microservice 456, the job status entry of job database 314 needs to be updated to reflect the job's "completed" status. The job status entry of in-memory job table 452 may now be out-of-date and, as a consequence, the entry for the job in in-memory job table 452 may be removed, for example, by database microservice 456. With the entry removed, if a client is requesting the job status, the next time a status request is received, the results may be retrieved directly from job database 314.

[0139] Additionally or alternatively, database microservice 456 may send a message with a job identifier and with an indication that the corresponding job has been completed to central message queue 312. That message may be picked up by notifier 460 which attempts to contact the client, perhaps at the URL they provided with the initial request, with a callback message indicating that results are available for the scraping request. The callback message may include a URL where the results can be retrieved.

[0140] As mentioned above, in addition to HTML, client computing device 102 may request parsed data, which can be generated by parser 308. In that embodiment, before database microservice 456 changes the job status to complete in job database 314 and sends a message to notifier 460 to send the callback message, database microservice 456 will send a message with the HTML via message queue 454 to parser 308. In an alternative embodiment, instead of sending the message with the HTML via message queue 454 to parser 308, the message will be sent with a link to the HTML file in HTML cache 462. When the message is received at parser 308, parser 308 can retrieve the HTML file from HTML cache 462 using the link within the message. Avoiding sending a

- 33 -

large HTML file via message brokers, including message queue 454, may have speed benefits as described above. Alternatively or additionally, parser 308 can consume a message linking to results in HTML cache 462 directly from central message queue 312, avoiding message queue 454.

[0141] Parser 308 conducts the parsing, generates the parsed data file and returns the parsed data file back to database microservice 456, which stores the parsed data file in job database 314. Only then, after the parsed data is stored in job database 314, does database microservice 456 update the job completion status and job database 314 and send a message to notifier 460 to send the callback message.

[0142] Also, as mentioned above, client computing device 102 can request that results be written to a particular external database belonging to the client, the external database may be a cloud storage database such as an S3 database available from Amazon.com, Inc. of Seattle Washington. In that embodiment, database microservice 456 may send the results to external database writer 464, which sends the results to the external database that the client had specified. In that case, the callback is sent after database writer 464 completes its job. The completion confirmation is sent via central message queue 312 and forwarded to message queue 454 where it is consumed by database microservice 456 for transmission of the callback message to the client.

Throttling Web Scraping Requests

[0143] FIG. 5 illustrates a web scraping system 104 with a throttler 502. In some cases, client computing device 102 may send a large number of requests for web scraping to web scraping system 104 in a short amount of time. The client may not want to stagger requests or may not know how to stagger requests into web scraping system 104. Instead, the client may prefer to send all the scraping job requests to web scraping system 104 at substantially the same time.

[0144] As mentioned above, scraper 310 may have a limited number of resources, including available proxies that can execute the requests. Those proxies become depleted with time or with enough concurrent requests as target websites begin to recognize that the requests originating from a proxy are automated and start to block incoming requests

- 34 -

from the proxy. If too many proxies are depleted, scraper 310 may be unable to scrape data, effectively disabling web scraping system 104.

[0145] While the client may prefer to send a large number of the scraping job requests to web scraping system 104 within a short time interval, the client may not care whether the results are returned immediately or over the course of an extended time period.

[0146] Throttler 502 enables web scraping system 104 to accept a large number of requests, but stagger when they are executed. Operation of throttler 502 is described in greater detail with respect to FIGS. 6 and 7.

[0147] FIG. 6 shows a timeline 600 illustrating how web requests can be throttled and trickled out over time. Timeline 600 illustrates how a plurality of web scraping requests 610A-E are received during a time period 602. As described above, each of the plurality of web scraping requests specifies a target website from which to capture content. And based on the incoming requests, job descriptions describing the web scraping job described by the request are generated. The descriptions are placed on a queue.

[0148] The job descriptions are removed from the queue over a time period 604. Time period 604 is longer than time period 602. The job descriptions may be removed incrementally over a repeating, fixed interval 606. The job description corresponding to API request 610A is removed at time 612A; the job description corresponding to API request 610B is removed at time 612B; the job description corresponding to API request 610C is removed at time 612C; the job description corresponding to API request 610D is removed at time 612D; and the job description corresponding to API request 610E is removed at time 612E. When each respective job description is removed from the queue, a web request for the target website as specified in the job description is generated. The generated web request is transmitted to the target website for processing.

[0149] Alternatively or additionally, the interval 606 may be variable. Alternatively or additionally, the requests may be removed non-sequentially (that is, out of order). For example, there can be prioritized queuing. In another example, a scraper can remove a job description lower ranked in the queue while another scraper is still processing a higher ranked job description.

[0150] As described above, the generated web request may be transmitted to the target website such that the web request reaches the target website via a proxy selected from a

- 35 -

group of proxies. According to an embodiment, the interval 606 may be selected to avoid the group of proxies from becoming blocked by the target website. In response to the web request, the content transmitted from the target website is received via the selected proxy. And the received content is transmitted to the client computing device.

[0151] More details on how the requests can be throttled are illustrated in FIG. 7. FIG. 7 is a system diagram 700 illustrating how the throttling mechanism can be implemented within web scraping system 104.

[0152] As described above, a request intake manager 302 is configured to receive a plurality of web scraping requests from a client computing device 102 (not shown in FIG. 7). Those API requests are sent to request status manager 304. Request status manager 304 may evaluate each API request and compare the request to parameters stored in throttle configuration database 714.

[0153] Throttle configuration database 714 stores information about which requests should be throttled. This may be set up on a per-client basis based on the level of service to which they have subscribed. For example, a client purchasing a higher level of service may experience less or no throttling, web scraping system 104 accepting that proxies may be depleted more quickly when servicing that customer's requests. Conversely, a client purchasing a lower level of service may opt for more throttling as they may not need results as quickly or may prefer to send a large number of requests within a short time period without requiring them all to be serviced immediately.

[0154] The parameters stored in throttle configuration database 714 may include aspects of the API request that are needed to define whether or not the API request should be throttled and how the API request should be throttled. For example, throttle configuration database 714 may have an entry for a particular client, target, and whether or not the API request is requesting that the web request be done from a headless browser.

[0155] Request status manager 304 may compare the API request to the parameters stored in throttle configuration database 714 and, based on that, may send the corresponding job description and a command 710 to a queue selector 716. Command 710 may indicate whether or not a web request described by the respective job description should be throttled.

- 36 -

- [0156] Based on the command 710, queue selector 716 sends the respective job description to either throttler 502 or central message queue 312. When command 710 indicates that the request should not be throttled, queue selector 716 may send the job description to central message queue 312, where it is picked up and processed by scraper 310 as described above.
- [0157] When command 710 indicates the request should be throttled, queue selector 716 sends the job description to throttler 502. Throttler 502 includes a plurality of throttle queues 702A...N. Throttler 502 may select from the plurality of throttle queues 702A...N depending on aspects of the job description. In an embodiment, each job description may include a client identifier indicating a client from which the request originates. Queue selector 716 may select the queue for the job description from a plurality of queues 702A...N based on the client identifier. Additionally or alternatively, each job description may include an identifier of the target website. Queue selector 716 may select the queue for the job description from a plurality of queues 702A...N based on the target website.
- [0158] Additionally or alternatively, queue selector 716 may select the queue for the job description from a plurality of queues 702A...N based on whether the job description specifies that web scraping be made from a headless browser. As described above, when the job description requests that web scraping be made with the headless browser, further requests are made to retrieve content that would be needed to render the target web page. For example, when a web page HTML is received from the target website, the content is analyzed to determine web addresses for additional content needed to render the web page and the additional content from the web addresses.
- [0159] Jobs are removed from the respective throttle queues 702A...N by a job remover 704. Job remover 704 is configured to remove job descriptions incrementally. Each of the throttle queues 702A...N may have an associated time window that job remover 704 waits between job descriptions. The associated time window may be stored in throttle configuration database 714 and may correspond to a level of service to which the client has subscribed. In addition, throttle configuration database 714 may have a number of job descriptions that should be removed from the respective throttle queue 702A...N during each time increment. For example, according to the specification in throttle

- 37 -

configuration database 714, job remover 704 may remove five job descriptions from throttle queue 702A every 20 seconds.

[0160] When job remover 704 removes a job description from throttle queues 702A...N, job remover 704 provides the job description to queue selector 716 along with a command 712 indicating that throttling is no longer necessary and that the job should be processed. Based on command 712, queue selector 716 forwards the job description to central message queue 312 where it is processed by scraper 310 as described above.

[0161] As mentioned above, throttling functionality is specified in throttle configuration database 714 according to a level of service requested by the customer. Based on the level of service, throttling may no longer be needed. When web scraping system 104 determines that web scraping requests from the client computing device no longer need to be throttled, system 104 may first determine when the queue is empty. Only when the queue is determined to be empty does web scraping system 104 delete the queue.

Traffic Counting for Proxy Web Scraping

[0162] Traditionally, proxy web scraping services may have tracked user activity on a per request basis for invoicing purposes. According to embodiments, user activity for proxy web based scraping services is tracked on a traffic basis. In particular, when a response is received to a web scraping traffic total, the amount of data transferred over a physical network layer may be calculated. To calculate the amount of data, an adjustment may be made to the total amount of data received to account for compression that occurs when the data is transferred. In addition, the amount of data may exclude requests and responses needed to acquire session information to request the scraped traffic.

[0163] FIG. 8 is a system diagram illustrating how traffic counting can be implemented in a web scraping system. FIG. 8 illustrates that web scraping system 104 includes a traffic counter 802, traffic count database 806 and invoice generator 808.

[0164] Traffic counter 802 is configured to track the amount of traffic received and/or sent by scraper 310 and store the data in traffic count database 806. Traffic counter 802 may count data from content and may count data from requests.

[0165] When web scraping system 104 receives new client requests from a client computing device 102, request intake manager 302 has a limit checker 804 that checks

- 38 -

whether the current amount of data that a client requests to scrape in a particular time exceeds a maximum threshold, which can be set based on a service level of the client. If it has exceeded a maximum threshold, request intake manager 302 may refuse to service the request.

[0166] Invoice generator 808 generates an invoice based on traffic counts in traffic count database 806. In addition, invoice generator 808 can serve as an interface to enable a client to get the amount of remaining data in his periodic (e.g. monthly) usage plan. For example, invoice generator 808 can receive a request from a client, determine an amount of data remaining as a difference between the current traffic total and a maximum allowable for the client, and return the amount of data remaining to the client.

[0167] FIG. 9 is a flowchart illustrating a method 900 for counting and regulating traffic in a web scraping system 104.

[0168] Method 900 begins with receipt of a web scraping request at step 902. As described above this receipt may occur at request intake manager 302 and the web scraping request may be from a client computing device, such as client computing device 102. The web scraping request is formatted according to an API understood by web scraping system 104 and may target a website. The API request may also specify the client making the request.

[0169] At step 904, limit checker 804 determines whether a current traffic total representing the amount of traffic tracked for particular client during the time window (such as monthly) exceeds a maximum allowable for the client. For example, limit checker 804 may check traffic count database 806 to determine the maximum allowable for the client. This check may occur in response to a request, repeatedly or intermittently. The maximum allowable may be set based on the level of service to which the client subscribes. Moreover, as will be described below, the maximum allowed may be specific, not just to a client but also to a scraper that the web scraping request received from the client computing device invokes. In that situation, traffic count database 806 tracks not just the current traffic level for the client but also the current traffic level for a particular scraper used by the client. Limit checker 804 determines whether the current traffic level for the scraper used by the client exceeds the maximum threshold.

- 39 -

- [0170] When traffic count database 806 determines that the current traffic level exceeds a maximum threshold for the time period, the operation proceeds to step 906, where request intake manager 302 refuses to process the request. When traffic count database 806 determines that the current traffic level does not exceed the maximum threshold for the time period, the operation proceeds to step 908.
- [0171] At step 908, scraper 310 selects a scraper from a plurality of scrapers 466A...N based on the target website such that the selected scraper includes instructions on how to generate the web request to extract data from the target website. As described above, each of the plurality of scrapers 466A...N include instructions on how to generate one or more HTTP requests from the target websites to be scraped. The instructions specified by the scraper are executed in steps 910 and 912.
- [0172] At step 910, any preliminary requests specified by the scraper selected in step 908 are conducted. As described above with respect to FIG. 2, to reach the ultimate target website, preliminary requests and responses may need to be conducted to get cookies expected by the target website. These requests are made from a proxy selected from a group of proxies. In different embodiments, these requests and responses may be counted towards the traffic total for a client and scraper. Alternatively, they may be excluded.
- [0173] At step 912, the HTTP web request is made to capture the target website that the user sought to scrape. The request is generated as specified in the scraper and transmitted to the target website via the proxy. Via the proxy, content is received from the target website. The size of the request and response is counted towards the total for a client of the client computing device, the size of this content is counted toward the traffic total.
- [0174] At step 914, additional content may be requested if the web API request is for a headless browser. As described above, and in that situation, an additional request may be made to retrieve content that would be needed to render the scraped page. In an embodiment, these additional requests and responses may also be counted toward the traffic total. Alternatively, they may be excluded.
- [0175] As described above, many websites are trying to block automated requests. At step 916, a determination may be made whether the request made in steps 910-914 was successful. If it was not successful, then the data counts for that request and response may not be excluded from the total. A new proxy is selected at step 918, and the operation

- 40 -

passes back to step 910 to repeat steps 910-916 until the request is successfully processed. In this way, when the target website is determined to have refused to serve the web request from the proxy, web scraping system 104 repeatedly attempts to re-send the web request to the target website via a different proxy. When the request is successful at step 916, the operation passes to step 920.

[0176] At step 920, the amount of data in the requests and responses in steps 910-916 as compressed for transmission over a network is determined. The determination may be made based on the compression level that would be used to transmit at a lower layer, such as the network, data link, or physical layer of the network.

[0177] To conduct that determination, an estimated compression level may be determined. In particular, a type of data represented by the received content may be determined. As described above with respect to FIG. 2, this type of data may be described, for example, in the HTTP headers. Based on the type of data, a compression factor representing an amount of compression expected when the type of data is transmitted over a network is determined. That determination may be made, for example, by looking up in a table a benchmark compression ratio for data of that type. Based on that compression factor, the amount of data in the received content as compressed for transmission is determined.

[0178] At step 922, the compressed count is included in a total amount of data for the client/scrapper. Then, when new requests come in, the process 900 is repeated. If the total amount of data has exceeded the threshold at 904, the request is refused at 906.

Shards Disabling

[0179] As described above, a plurality of database servers 404A... N, for example shards, are used to access job database 314. Sometimes, at least one of the plurality of database servers 404A... N should be disabled, avoiding further jobs sent to the disabled server. For example, a database server may stop performing at all (for example, the process may be hung up or have crashed), it may stop performing well enough, or it may be overloaded and should not receive any more jobs until it frees up. How database servers 404A... N can be disabled is described with respect to FIGS. 10-11.

- 41 -

[0180] FIG. 10 illustrates operation of various components of web scraping system 104 to detect whether respective database servers 404A...N should be enabled or disabled.

[0181] As described above, each respective database server 404A...N has a message queue 454. As described above, message queue 454 is a message broker that queues job descriptions for consumption by database microservice 456. When database microservice 456 consumes a message from message queue 454, database microservice 456 places the job description in job database 314 with a “pending” status, initiating processing of the job.

[0182] Request intake manager 302 repeatedly checks the health each of a plurality of database servers 404A...N to determine whether they are available to accept new jobs. In various embodiments, the health checks may be periodic or intermittent. Based on the health checks, request intake manager 302 determines whether each of a plurality of database servers are to be enabled or disabled in a table 1008. Table 1008 may be stored in a non-persistent memory to increase access speed.

[0183] To check the health of a respective server from the database servers 404A...N, request intake manager 302 can repeatedly send various messages to check whether the respective database server is operational in different embodiments. Message 1002 checks the connection between request intake manager 302 (which is a server that receives web scraping requests from client computing devices) and message queue 454 (the respective database server’s message broker). If the connection is nonresponsive or returns an error message, request intake manager 302 may mark the respective database server as disabled in table 1008.

[0184] Message 1004 checks a number of job descriptions (and possibly other messages used in the scraping process) queued in message queue 454. If the number of job descriptions (and possibly other messages) exceeds a threshold, request intake manager 302 may mark the respective database server as disabled in table 1008.

[0185] Message 1006 checks a connection to jobs database 314. Request intake manager 302 may try to connect to jobs database 314. If the connection is nonresponsive or returns an error message, request intake manager 302 may mark the respective database server as disabled in table 1008. If request intake manager 302 can connect to jobs database 314, Request intake manager 302 can infer that database microservice can as well.

- 42 -

- [0186] In addition to a binary enable or disable status, request intake manager 302 may, depending on the results of messages 1002-1006, mark features of database servers 404A...N as enabled or disabled in table 1008. For example, request intake manager 302 may label a database server as read-only, when it is overworked or hung-up, but is capable of providing read access. And, request intake manager 302 may label a database server as shutdown, when even read-only access is not available.
- [0187] FIG. 11 shows how request intake manager 302 uses database server status table 1008 to determine which database servers 404A...N to use to service an API request from client computing device 102.
- [0188] Request intake manager 302 includes a database server selector 1100 that, when a web scraping request is received from a client computing device, selects one of the database servers identified as enabled in the table. After the selection is made, request intake manager 302 sends a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job.
- [0189] Request intake manager 302 can also evaluate table 1008 to determine whether a number of database servers that are disabled in the plurality of database servers exceeds a threshold. When the number of database servers that are disabled exceeds the threshold, request intake manager 302 may alert an administrator.

Direct Data Center Request Relating to a Web Scraping Job

- [0190] Web scraping system 104 may be distributed over a large number of computing devices. The computing devices may be geographically disparate. As described above, the job database 314 may be horizontally partitioned over a large number of database shards. At least in part for this reason, data, such as web scraping results and possibly job status information, may be available at one computing device within web scraping system 104 at one geographic location, but not another. To improve efficiency, it is advantageous to ensure that a request for the results is directed to the correct computing device at the correct geographic location.
- [0191] FIGS. 12A-B illustrate systems for load balancing web scraping requests that direct a request for results to the same subdomain that requested results from a target. FIG. 12A shows a system 1200 with tiered load balancers across different locations.

- 43 -

[0192] As illustrated in Figure 12A, system 1200 which includes a data center balancer 1202 that can distribute API requests from client computing device 102 to one of several data centers 1206A...N. In one example, data center balancer 1202 may operate by sending back to client computing device 102 an HTTP redirect to one of several destinations each addressable by a hostname. In another example, data center balancer 1202 may act as a proxy, forwarding requests from client computing device 102 to a respective data center.

[0193] As illustrated in FIG. 12A, the possible destinations are located at different data centers 1206A...N. Each of the data centers 1206A...N is a building, a dedicated space within a building, or a group of buildings used to house computer systems and associated components, such as telecommunications and storage systems. Each may include redundant or backup components and infrastructure for power supply, data communication connections, environmental controls (e.g., air conditioning, fire suppression), and various security devices.

[0194] In system 1200, different data centers 1206A...N include one or more load balancers and one or more instances of request intake manager 302, which, as described above, is a software server that accepts incoming API requests from client computing device 102. In the example in FIG. 12, data center 1206A includes load balancer 1208-1A...N and request intake manager 302-1A...N; data center 1206B includes load balancer 1208-2A...N and request intake manager 302-2A...N; and data center 1206N includes load balancer 1208-M, A...N and request intake manager 302-M, A...N.

[0195] Each of the load balancers and request intake managers are addressable by a hostname. A hostname is a label that is assigned to a device connected to a computer network and that is used to identify the device in various forms of electronic communication. Internet hostnames, like those shown in FIG. 12A, are composed of a sequence of labels concatenated with dots. The hostname is resolvable into an Internet protocol (IP) address (that is, an OSI layer 3 network address) using the domain name system. The domain name system is a hierarchical naming space, where the rightmost label conveys the top-level domain and the hierarchy descends right to left, each label to the left specifying a subdomain of the domain to the right. In an embodiment, it is not

- 44 -

necessary to differentiate between intake managers in the same DC, so they need not have direct hostnames.

[0196] In the example in FIG. 12, the host names of all the host load balancers and request intake manager's have the top level domain "net." All are within a subdomain "example.net" of "net." Each has a unique subdomain within "example.net" representing the particular network address for the load balancer or request intake manager service. The unique subdomains might follow a regular naming pattern.

[0197] For the respective data centers 1206A-N, load balancers 1208-1A...N, load balancers 1208-2A...N, and load balancers 1208-M, A...N determine how to distribute requests among the various request intake managers 302-1A...N, request intake managers 302-2A...N, and request intake managers 302-M, A...N in the corresponding data centers 1206A-N. The group of request intake managers and load balancers within each data center request intake managers and load balancers which may be referred to as a "leg."

[0198] Each load balancer may work as a proxy distributing requests among corresponding request intake managers in a particular fashion, such as round-robin. They may, for example, use OSI layer 4 (transport layer) load-balancing forwarding all port 80 TCP requests onto the backend, which may constitute the corresponding group of request intake managers. In this way, the load balancers distribute requests among corresponding request intake managers.

[0199] In an example operation, client computing device 102 may make a new API request to web scraping system 104. For usability purposes, data center balancer 1202 may represent the main entry point into web scraping system 104. For example, all new requests may be sent to the hostname "api.example.io," which addresses data center balancer 1202.

[0200] With a new web scraping request, which data center, whether 1206A...N and request intake manager 302-1-M, A-N or another, ultimately services the request may not matter. Data center balancer 1202 may select which leg and which load balancer within the leg in a random or round-robin fashion. Alternatively, data center balancer 1202 may select a data center 1206A...N out of geographical convenience. Data center balancer 1202 redirects the API request for web scraping to one of the load balancers, which

- 45 -

selects one of the corresponding request intake managers and accesses the proxy to forward the request on to the corresponding request intake manager. From there, request intake manager 302 validates the request, and causes web scraping system 104 to store a job description in job database 314, initiating the scraping operation, as described above. Once the job is initiated, request intake manager 302 may return, perhaps through load balancer 1208, a job identifier.

[0201] As mentioned above, the web scraping request may be asynchronous. In the case of an asynchronous request, client computing device 102 may make a second API request to retrieve the results of the previously requested web scraping job. In addition, as described above, client computing device 102 may make additional API requests to request the status or other information related to the previously requested web scraping job. To identify the relevant job, the API request may include the job description identifier to which it relates.

[0202] In that circumstance, at least in one embodiment, client computing device 102 may send the API request to data center balancer 1202. Data center balancer 1202 has a job look up table 1204. Job look up table 1204 may match a job identification number to a corresponding database shard and may match a corresponding database shard to a data center responsible for that shard in that job. Then, data center balancer 1202 may direct the API request to the hostname of a load balancer of a data center with the data needed to respond to the API request.

[0203] FIG. 12B illustrates an alternative embodiment. As mentioned above, when request intake manager 302 accepts the initial request, it may return to client computing device 102 a job identifier for the job. In addition, client computing device 102 can return a hostname or complete URL of the request intake manager 302 corresponding to the shard that will receive the results. When the results are available, web scraping system 104 may notify client computing device 102. Alternatively or additionally, that notification message could include the hostname or complete URL of the request intake manager 302 corresponding to the shard that has the results. By supplying the hostname in this manner, client computing device 102 can retrieve the results and request status information directly from the data center that has the data requested.

- 46 -

- [0204] However, sometimes, client computing device 102 will nonetheless make the request directly to the common entry point offered by data center balancer 1202. In that situation, when the request reaches a data center 1206A-N, software at the data center, such as a load balancer or request intake manager, determines whether the request has reached the correct leg. If it has not, then it identifies the correct leg and identifies a hostname that corresponds to that leg. With the hostname identified, the software returns a redirect command to client computing device 102 to go to the correct leg that has the data the client requested.
- [0205] To identify the correct leg, each request intake manager 302 (or load balancers 1208) may need to know which data center is responsible for the shard that has the data that the user requested. For example, the request intake manager 302 (or load balancers 1208) may need to map the job identifier with a shard identifier and match the shard identifier with the hostname of a particular leg, such as is described above with respect to job lookup 1204 in FIG. 12A.
- [0206] FIGS. 13A-B illustrate methods for operation of the systems in FIGS. 12A-B. FIG. 13A illustrates a method 1300 for processing a new web scraping request and FIG. 13B illustrates a method 1350 for retrieving status or results of a prior web scraping job.
- [0207] Starting with method 1300, at 1302, a request to scrape content is received at the main entry point, such as data center balancer 1202 described with respect to FIGS. 12A and B. The request may be an asynchronous API request to request that the web scraping system scrape the target website as described above and may specify the target website.
- [0208] At 1304, the web scraping request is generated and executed. The web request for the target website is generated as specified in the API request. The web request is transmitted such that the web request reaches the target website via a proxy selected from a group of proxies. And, in response to the web request, scraped content is received from the target website via the proxy.
- [0209] At 1306, web scraping system 104 sends a message to client computing device 102 acknowledging the request. The message may include a job identifier and may specify an address or hostname (such as a sub-domain) for the client computing device 102 to retrieve the scraped content.

- 47 -

- [0210] Turning to method 1350, at 1352, an API request is received from client computing device 102. The API request relates to a web scraping job that was previously requested of a web scraping system. For example, the API request may request results of the web scraping job or a state of the web scraping job.
- [0211] At 1354, a hostname addressing a host at a data center of a plurality of data centers executing the web scraping system is identified. The data center identified is the one that is servicing or has serviced the web scraping job.
- [0212] At 1356, the API request is directed to the identified hostname. Steps 1352-1356 may be executed by data center balancer 1202 as the main entry point for web scraping system 104, as described above with respect to FIG. 12A. Alternatively or additionally, steps 1352-1356 may be executed by request intake manager 302, as described above with respect to FIG. 12B. To direct the API request to the identified hostname, web scraping system 104 may return an instruction to client computing device 102 to redirect to the identified hostname, for example using an HTTP redirect response. Alternatively, the respective request intake manager 302 may act as a proxy to forward the API request to the identified hostname.
- [0213] Finally, at 1358, the API request is serviced at the identified data center, servicing the API request. In this way, the requested results or status is returned to client computing device 102.

Token-Based Authentication for a Proxy Web Scraping Service

- [0214] An API call from client computing device 102 may include credentials that enable a client to authenticate herself. Web scraping service 104 may subscribe clients at different service levels so it is important to authenticate the client when receiving API requests. Traditionally, web scraping systems may have had to re-authenticate the credentials at several points in the scraping process.
- [0215] According to an embodiment, the credentials are authenticated at an entry point of the system where a token is generated that is passed along to various servers within web scraping system 104. The various servers within web scraping system 104 may use the token as a way to check the client's identity. With the client's identity confirmed, the

- 48 -

various servers within web scraping system 104 can provide functions according to a service level to which the client has subscribed.

[0216] FIG. 14 illustrates a system 1400 for token-based authentication at a load balancer for a web scraping system.

[0217] In system 1400, request intake manager 302 includes three sets of API receivers, each set of API receivers being a set of servers configured to accept a different type of web scraping request. URL API receiver 1402A...N accepts a web scraping request including a fully formatted URL. API receiver 1402A...N may accept API requests as HTTP proxy as described above. When such an HTTP proxy request is received at API receiver 1402A...N, the respective API receiver 1402A...N reformats the request according to an API with known parameters and sends the reformatted request to at least one of synchronous API receivers 1404A...N.

[0218] Synchronous API receiver 1404A...N is configured to accept a synchronous, that is, realtime, web scraping request. When a synchronous API receiver 1404A...N receives a web scraping request, it leaves the TCP/HTTP connection with the client computing device 102 open until the scraped results are available. Once web scraping system 104 has completed the scraping request, the respective synchronous API receiver 1404A...N returns the scraped results to client computing device 102. To conduct the web scraping request, synchronous API receiver 1404A...N may send the API request to at least one of asynchronous API receivers 1406A...N.

[0219] Asynchronous API receiver 1406A...N may be the entry point into the remainder of web scraping system 104 as described above. For example, as described above, asynchronous API receiver 1406A...N may generate a job description and send the job description to request status manager 304, which initiates the scraping job. When asynchronous API receiver 1406A...N receives an API request from client computing device 102 or from synchronous API receiver 1404A...N, asynchronous API receiver 1406A...N closes the connection to the receiver or device before the web scraping is complete, perhaps replying with an acknowledgment. Once the web scraping is complete and results are available at web scraping system 104, the client computing device 102 or synchronous API receiver 1404A...N receives a callback notification that results are

- 49 -

available. In this way, an API request is passed between a plurality of servers, each configured to perform a function of web scraping system 104.

[0220] Each set of API receivers is accessible via a respective load balancers 1208A-C, and each of the respective load balancers 1208A-C may represent an entry point into web scraping system 104 for client computing device 102. As described above, each of load balancers 1208A-C distribute requests among the plurality of associated receivers to utilize and distribute requests across the multiple, parallel servers.

[0221] In the example in FIG. 14, a client desiring to send an API request using the full URL of the destination, such as using the Web proxy protocol as described above, may send the API request to a hostname “url.example.io.” That request would reach load balancer 1208A. Load balancer 1208A selects one of URL API receivers 1402A...N to handle the request and forwards the API request on to the selected receiver. The selected API receiver 1402A...N reformats the proxy request into a different API format and sends on to one of synchronous API receivers 1404A...N, which forwards the requests on to one of asynchronous API receivers 1406A...N. When the request is completed and results are scraped, the selected synchronous API receiver 1404A...N is notified and returns scraped results to the selected URL API receiver 1402A...N. The selected URL API receiver 1402A...N formats the response as a proxy response and sends the response to client computing device 102.

[0222] Continuing the example operation in FIG. 14, a client desiring to send an API request using a synchronous API call may send the API request to a hostname “sync.example.io.” That request would reach load balancer 1208B. Load balancer 1208B selects one of synchronous API receivers 1404A...N to handle the request and forwards the API request on to the selected receiver. The selected synchronous API receivers 1404A...N keeps the connection open while forwarding the API request to asynchronous API receiver 1406A...N, which sends the job for scraping. When the results are available, the selected synchronous API receiver 1404A...N is notified and returns the results to client computing device 102, closing the connection.

[0223] Finally, continuing the example operation in FIG. 14, a client desiring to send an API request using an asynchronous API call may send the API request to a hostname “api.example.io.” That request would reach load balancer 1208C. Load balancer 1208C

- 50 -

selects one of asynchronous API receivers 1406A...N to handle the request and forwards the API request on to the selected receiver. The selected asynchronous API receivers 1404A...N close a connection between the web scraping system 104 and the client computing device while the web scraping system scrapes the target website.

[0224] According to an embodiment, each load balancer 1208A-C includes a respective Load balance service 1404A-C and authentication service 1412A-C. Load balance service 1404A-C performs the backend selection and proxy forwarding or redirection for the load balancer as described above. On receipt of a request, Load balance service 1404A-C communicates with the respective authentication service 1412A-C to authenticate the user.

[0225] In particular, each of authentication service 1412A-C is configured to validate credentials received with an API request from client computing device 102. When the authentication service 1412A-C cannot validate the credentials, the respective load balancer 1208A-C refuses to service received API request. When the authentication service 1412A-C validates the credentials, the authentication service 1412A-C generates a token indicating an identity of a client associated with the credentials. Load balancer 1208A-C transmits the API request along with the token to a server configured to initiate a scraping process on the web scraping request.

[0226] The token may be passed along with the API request to the various API receivers and other servers and modules providing functions within web scraping system 104. These various servers analyze the token to determine whether the client is authorized to conduct the function performed by the respective server, and when the client is authorized to conduct the function, perform the function, such as handling a particular type of API or conducting a particular type of scraping request. To conduct the analysis, a server may check the level of service of the client indicated within the token. Based on the level of service, the server may determine whether the function provided by the server is authorized. Moreover, when an API request is made to retrieve previously scraped results or to check status of a previously made request, a server may use the token to determine whether the previously-made request was made by the same client.

[0227] In addition to specifying the client, the token may specify a role for the client. Subsequent servers may use the role within the token to determine whether a particular

- 51 -

function of web scraping system 104 is authorized. In an example, an “admin” role may allow the request made by any client, while a “user” role may allow the request to retrieve only the data requested by the authenticated client.

[0228] FIG. 15 shows a system 1500 illustrating the operation of token-based authentication for a web scraping system 104 in greater detail. Load balance service 1410 makes a request 1502 of authentication service 1412. Request 1502 may, for example, be an HTTP POST request. The body of the post request may have credentials, and a path used by the POST request and may contain the type of credentials being used. Credentials, for example, may be a username and password. In another example, the credentials can be a token.

[0229] Authentication service 1412 compares the credentials to what is stored in a credentials database 1506. Credentials database 1506 may have multiple layers in a memory hierarchy to speed up lookup. Based on that comparison, authentication service 1412 generates and returns authentication token 1504.

[0230] In an example, token 1504 may be a JSON Web Token (JWT). Token 1504 includes a header 1510, payload 1512, and digital signature 1514. Header 1510 may indicate the algorithm and type used for the JSON. Payload 1512 may have data regarding the authentication. This can include the client’s identity, services to which the client has subscribed, roles of the client, and expiration date/time for the token, and whether the authentication was successful.

[0231] Digital signature 1514 uses cryptographic methods to guarantee that the payload of the client has not been tampered with. For example, generating digital signature 1514 may involve generating a one-way hash and encrypting it with a private key of the authentication service 1412. Then, when digital signature 1514 is verified, the public key is used to decrypt the hash and compare it with a hash of the payload 1512. In this way, digital signature 1514 can be used to verify that the payload has not been tampered with by anyone who does not have possession of the private key.

Cold Database for Archiving Results of Web Scraping Across Data Centers

[0232] As mentioned above, web scraping system 104 may include a history archive 306 that stores the scraped data for longer periods. History archive 306 allows a user to track

- 52 -

changes to a target website over time. To retrieve data, a client may request a dump from a history archive 306 of the previously scraped results that meet certain criteria. In particular, the client's search criteria may describe metadata of results sought to be retrieved. The metadata may be compared to an index that describes the results. Based on the index, links to the results may be returned to the client.

[0233] However, as described above, functionality of web scraping system 104 may be spread across multiple data centers. Each data center may have its own index. According to an embodiment, results from the various indexes are combined to generate search results corresponding to particular search criteria. Aspects of this feature are illustrated in FIG. 16.

[0234] FIG. 16 illustrates a system 1600 that stores results from a plurality of data centers into an archive database.

[0235] System 1600 includes an access server 1630 and an archive database 1612. In an example, archive database 1612 may store parsed data scraped by web scraping system 104. As an example, the parsed data retrieved for each web scraping request may be stored as a JSON file and written to archive database 1612. The web scraping system 104 stores the scraped data into archive database 1612, where it is available for future analysis.

[0236] To access data in archive database 1612, client computing device 102 may send a request to access server 1630. The request may address a location of the desired JSON file within archive database 1612. Access server 1630 retrieves data from archive database 1612 via one or more cloud interfaces 1604A-B. Cloud interfaces 1604A-B may have the credential and interface information needed to send and retrieve data from archive database 1612.

[0237] Archive database 1612 may be a cloud storage service. It may be an Infrastructure as a service (IaaS), it may utilize encryption for security and data compression for bandwidth, storage, and latency optimization. Similarly, dump database 1610 may also be a cloud storage database.

[0238] Each respective data center 1206A-B includes a scraper (not shown in FIG. 16), parser 308A-B, and history archive 306A-B. Each history archive 306A-B includes a writer 1602A-B, cloud interface 1604A-B, index 1606A-B, and dumper 1608A-B.

- 53 -

- [0239] As described above with respect to FIGS. 3 and 4, the scraper is configured to generate a web request for a target website according to a job description generated based on an API request. The scraper transmits the generated web request to the target website via one or more proxies. In response to the web request, the scraper receives a content from the target website. Also as described above, parser 308A-B are each configured to parse the response from the target website. The parsed content may be represented, for example, as a data structure in a JSON format. Parser 308A-B sends the parsed content to the respective writer 1602A-B.
- [0240] Once received from parser 308A-B, the respective writer 1602A-B sends the parsed content to the respective cloud interface 1604A-B to be written to archive database 1612. In addition, metadata about the scraped content is saved to index 1606A-B. The metadata may include, for example, any of the information in the corresponding API request or job description and a time stamp when the scraping actually occurred.
- [0241] Information from a target website may only be stored in archive database 1612 at a particular frequency. In other words, if content is retrieved from a target website (with particular parameters, as set out in the job description) and was recently stored in archive database 1612, writer 1602A-B may decline to write it again. In this way, storage is conserved on archive database 1612.
- [0242] To determine whether a writer 1602A-B has previously archived content from the target website, writer 1602A-B may compare metadata parameters with entries already existing in its respective index 1606A-B. Because the respective index 1606A-B may only have metadata of content stored from a particular data center, additional synchronization and de-duplication may need to occur as described below with respect to FIG. 17.
- [0243] Having the index distributed across multiple data centers means that search criteria must be distributed across multiple servers to search all the indexes. If client computing device 102 wishes to retrieve a plurality of scraped results meeting certain criteria, client computing device 102 sends the search criteria to web scraping system 104. The search criteria is distributed among a plurality of data centers 1206A-B and is received by a respective dumper 1608A-B. The respective dumper 1608A-B searches the respective index 1606A-B at that data center and determines what entries in the index file

- 54 -

satisfy the search criteria. In particular, each entry may have metadata describing the search results and the metadata for each entry which may be compared against the search criteria. The dumper 1608A-B sends the resulting list, including links to the location within archive database 1612 where the scraped results may be found, to a dump database 1610, where they are combined.

[0244] As shown in FIG. 16, client computing device 102 reads the combined results as 1622 from dump database 1610. Then, client computing device 102 sends each individual link to access server 1630, which uses the address of the scraped results to retrieve the appropriate results from archive database 1612 and return the retrieved results to client computing device 102.

[0245] FIG. 17 illustrates a system 1700 that de-duplicates data scraped from a plurality of data centers before storing into an archive database.

[0246] As mentioned above, when parsed, scraped content is stored in the archive, writer 1602A-B writes metadata describing the respective content into respective index file 1606A-B. At that time, writer 1602A-B will send the metadata to other data centers via a data feed 1702.

[0247] Data feed 1702 may be a software bus that uses stream-processing to distribute real-time data feeds. It may use a binary TCP-based protocol and may rely on “message set” obstruction which groups messages together to reduce the overhead of the network round trip.

[0248] Each respective data center includes a respective feed receiver 1704A-B. The feed receiver is configured to accept data feed 1702, streaming metadata from other data centers and storing it into the respective prior scraped file 1706A-B. Each metadata entry may be stored with an expiration date. The expiration date may be set based on the time period to preclude future scrapes from being archived, such as one month. When the expiration date occurs, the entry may be removed from the prior scraped file 1706A-B.

[0249] To determine whether another data center has previously retrieved the content from the target website and stored it in archive database 1612, the respective writer 1602A-B may compare metadata parameters with entries existing in its prior scraped file 1706A-B. In this way, duplications across data centers within archive database 1612 are avoided.

Conclusion

- [0250]** Each of the modules, servers and other components described above (including client computing device 102, web scraping system 104, Web proxy 106A... N, target Web server 108, scraper 310, parser 308, request intake manager 302, request status manager 304, central message queue 312, history archive 306, proxy rotator 486, target scraper 460 6A-N, message queue 454, notifier 460, external database writer 464, API receiver 402A-N, in memory access server 450, database servers 404A-N, message queue 454, micro service 456, throttler 502, throttle queue 702A... N, job remover 704, queue selector 716, invoice generator 808A, traffic counter 802, limit checker 804, database server selector 1100, load balancer 1208, data center balancer 1204, URL API receiver 1402A-N, synchronous API receiver 1404A-N, asynchronous API receiver 1406A-N, load balance service 1410A-C, authentication service 1412A-C, writer 1602A-B, cloud interface 1604A-B, dumper 1608A-B, and access server 1630) may be implemented on software executed on one or more computing devices or different computing devices.
- [0251]** A computing device may include one or more processors (also called central processing units, or CPUs). The processor may be connected to a communication infrastructure or bus. The computer device may also include user input/output device(s), such as monitors, keyboards, pointing devices, etc., which may communicate with communication infrastructure through user input/output interface(s).
- [0252]** One or more of the processors may be a graphics processing units (GPU). In an embodiment, a GPU may be a processor that is a specialized electronic circuit designed to process mathematically intensive applications. The GPU may have a parallel structure that is efficient for parallel processing of large blocks of data, such as mathematically intensive data common to computer graphics applications, images, videos, etc.
- [0253]** The computer device may also include a main or primary memory 408, such as random access memory (RAM). Main memory 408 may include one or more levels of cache. Main memory 408 may have stored therein control logic (i.e., computer software) and/or data.
- [0254]** The computer device may also include one or more secondary storage devices or memory. The secondary memory may include, for example, a hard disk drive, flash storage and/or a removable storage device or drive.

- 56 -

- [0255] The computing device may further include a communication or network interface. The communication interface may allow the computer system 400 to communicate and interact with any combination of external devices, external networks, external entities, etc. For example, the communication interface may allow the computer system to access external devices via network 100, which may be wired and/or wireless (or a combination thereof), and which may include any combination of LANs, WANs, the Internet, etc
- [0256] The computing device may also be any of a rack computer, server blade, personal digital assistant (PDA), desktop workstation, laptop or notebook computer, netbook, tablet, smartphone, smartwatch or other wearable, appliance, part of the Internet-of-Things, and/or embedded system, to name a few non-limiting examples, or any combination thereof.
- [0257] The computer device may access or host any applications and/or data through any delivery paradigm, including but not limited to remote or distributed cloud computing solutions; local or on-premises software ("on-premise" cloud-based solutions); "as a service" models (e.g., content as a service (CaaS), digital content as a service (DCaaS), software as a service (SaaS), managed software as a service (MSaaS), platform as a service (PaaS), desktop as a service (DaaS), framework as a service (FaaS), backend as a service (BaaS), mobile backend as a service (MBaaS), infrastructure as a service (IaaS), etc.); and/or a hybrid model including any combination of the foregoing examples or other services or delivery paradigms.
- [0258] Any applicable data structures, file formats, and schemas in the computing devices may be derived from standards including but not limited to JavaScript Object Notation (JSON), Extensible Markup Language (XML), Yet Another Markup Language (YAML), Extensible Hypertext Markup Language (XHTML), Wireless Markup Language (WML), MessagePack, XML User Interface Language (XUL), or any other functionally similar representations alone or in combination. Alternatively, proprietary data structures, formats or schemas may be used, either exclusively or in combination with known or open standards. Any of the databases or files described above (including job database 314 In-memory job table 452, HTML cache 462, throttle configuration database 714, traffic count database 806, database server status table 1008, credentials

- 57 -

database 1506, archive database 1612, and dump database 1610) may be stored in any format, structure, or schema in any type of memory and in a computing device.

[0259] In some embodiments, a tangible, non-transitory apparatus or article of manufacture comprising a tangible, non-transitory computer-usable or readable medium having control logic (software) stored thereon may also be referred to herein as a computer program product or program storage device. This includes, but is not limited to, main memory, secondary memory, and removable storage units, as well as tangible articles of manufacture embodying any combination of the foregoing. Such control logic may cause such data processing devices to operate as described herein.

[0260] A website is a collection of web pages containing related contents identified by a common domain name and published on at least one web server. A domain name is a series of alphanumeric strings separated by periods, serving as an address for a computer network connection and identifying the owner of the address. Domain names consist of two main elements – the website's name and the domain extension (e.g., .com). Typically, websites are dedicated to a particular type of content or service. A website can contain hyperlinks to several web pages, enabling a visitor to navigate between web pages. Web pages are documents containing specific collections of resources that are displayed in a web browser. A web page's fundamental element is one or more text files written in Hypertext Markup Language (HTML). Each web page in a website is identified by a distinct URL (Uniform Resource Locator). There are many varieties of websites, each providing a particular type of content or service.

[0261] Identifiers, such as "(a)," "(b)," "(i)," "(ii)," etc., are sometimes used for different elements or steps. These identifiers are used for clarity and do not necessarily designate an order for the elements or steps.

[0262] The present invention has been described above with the aid of functional building blocks illustrating the implementation of specified functions and relationships thereof. The boundaries of these functional building blocks have been arbitrarily defined herein for the convenience of the description. Alternate boundaries can be defined so long as the specified functions and relationships thereof are appropriately performed.

[0263] The foregoing description of the specific embodiments will so fully reveal the general nature of the invention that others can, by applying knowledge within the skill of

- 58 -

the art, readily modify and/or adapt for various applications such as specific embodiments, without undue experimentation, and without departing from the general concept of the present invention. Therefore, such adaptations and modifications are intended to be within the meaning and range of equivalents of the disclosed embodiments based on the teaching and guidance presented herein. It is to be understood that the phraseology or terminology herein is for the purpose of description and not of limitation, such that the terminology or phraseology of the present specification is to be interpreted by the skilled artisan in light of the teachings and guidance.

[0264] The breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

[0265] The disclosure presents a computer-implemented method for executing a plurality of web scraping requests over time, comprising:

(a) during a first time period, receiving a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;

(b) placing job descriptions for the plurality of web scraping requests on a queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;

(c) incrementally over a second time period longer than the first time period, removing the job descriptions from the queue;

when each respective job description is removed from the queue:

(d) generating a web request for the target website as specified in the job description; and

(e) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.

[0266] The method is presented wherein the removing (c) occurs in increments selected to avoid the group of proxies from becoming blocked by the target website.

[0267] The method is presented further comprising:

(f) in response to the web request, receiving content transmitted from the target website; and

- 59 -

(g) transmitting the received content to the client computing device.

[0268] The method is presented wherein the plurality of web scraping requests each include a client identifier indicating a client from which the web scraping request originates, further comprising:

(f) selecting the queue from a plurality of queues based on the client identifier.

[0269] The method is presented wherein the plurality of web scraping requests each further include an identifier of the target website, further comprising:

(g) selecting the queue from the plurality of queues based on the identifier of the target website.

[0270] The method is presented further comprising selecting the queue from the plurality of queues based on whether the job description specifies that web scraping be made from a headless browser.

[0271] The method is presented further comprising, when the job description requests that web scraping be made from the headless browser:

(h) in response to the web request, receiving content specifying a web page, the content transmitted from the target website;

(i) analyzing the content to determine web addresses for additional content needed to render the web page; and

(j) requesting the additional content from the web addresses.

[0272] The method is presented wherein incrementally removing (c) comprises repeated removing a job description from the queue at a time window.

[0273] The method is presented wherein incrementally removing (c) comprises repeated removing a plurality of job descriptions from the queue at the time window.

[0274] The method is presented wherein an amount of job descriptions in the plurality of job descriptions are specified based on a client that sent the plurality of web scraping requests.

[0275] The method is presented further comprising:

(f) determining that web scraping requests from the client computing device no longer need to be throttled;

(g) determining when the queue is empty; and

- 60 -

(h) when the queue is determined to be empty, deleting the queue.

[0276] The method is presented further comprising:

(f) determining which of the plurality of web scraping requests are synchronous requests and which are asynchronous requests; and

(g) prioritizing web scraping requests determined in (f) to be synchronous requests over web scraping requests determined in (f) to be asynchronous requests to help ensure that results from synchronous requests are returned before a connection with the client computing device times out.

[0277] A system is presented for executing a plurality of web scraping requests over time, comprising:

at least one processor;

a memory configured to store a queue;

a request intake manager configured to, during a first time period, receive a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;

a throttler configured to place job descriptions for the plurality of web scraping requests on the queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;

a job remover configured to, incrementally over a second time period longer than the first time period, remove the job descriptions from the queue;

a scraper configured to, when each respective job description is removed from the queue, generate a web request for the target website as specified in the job description, and transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.

[0278] The system is presented wherein the job remover is configured to remove in increments selected to avoid the group of proxies from becoming blocked by the target website.

[0279] The system is presented wherein the scraper is configured to, in response to the web request, receive content transmitted from the target website, and the request intake manager is configured to notify the client that the received content is available for transmission to the client computing device;

- 61 -

[0280] The system is presented further comprising a memory configured to store a plurality of queues, wherein the plurality of web scraping requests each include a client identifier indicating a client from which the client originates, further comprising:

a queue selector that selects the queue from a plurality of queues based on the client identifier.

[0281] The system is presented wherein the plurality of web scraping requests each further include an identifier of the target website, the queue selector is configured to select the queue from the plurality of queues based on the identifier of the target website.

[0282] The system is presented further comprising the queue selector configured to select the queue from the plurality of queues based on whether the job description specifies that web scraping be made from a headless browser.

[0283] The system is presented further comprising:

a throttle configuration database that stores a regular time window associated with the queue, wherein the job remover repeatedly removes a job description from the queue at the time window.

[0284] A non-transitory computer-readable device having instructions stored thereon is presented that, when executed by at least one computing device, cause the at least one computing device to perform operations comprising:

(a) during a first time period, receiving a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;

(b) placing job descriptions for the plurality of web scraping requests on a queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;

(c) incrementally over a second time period longer than the first time period, removing the job descriptions from the queue;

when each respective job description is removed from the queue:

(d) generating a web request for the target website as specified in the job description; and

(e) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.

- 62 -

[0285] A computer-implemented method for tracking user activity is presented, comprising:

- (a) receiving a web scraping request from a client computing device, the web scraping request specifying a target website to capture content from;
- (b) based on the web scraping request, generating a web request for the target website;
- (c) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies;
- (d) in response to the web request, receiving, via the proxy, content transmitted from the target website;
- (e) counting an amount of data in the received content to determine a current traffic total for a client of the client computing device; and
- (f) transmitting the received content to the client computing device.

[0286] The method is presented, further comprising:

- (g) based on the current traffic total, generating an invoice for a client corresponding to the client computing device.

[0287] The method is presented, wherein the current traffic total is for a time period, further comprising:

- (g) receiving an additional web scraping request from the client computing device;
- (h) determining whether the current traffic total exceeds a maximum allowable for a client corresponding to the client computing device; and
- (i) when the current traffic total is determined to exceed the maximum allowable in (h), refusing to service the additional web scraping request.

[0288] The method is presented, wherein the current traffic total is for a time period, further comprising:

- (g) receiving an additional web scraping request from the client computing device;
- (h) determining whether the current traffic total exceeds a maximum allowable for a client corresponding to the client computing device; and

- 63 -

(i) when the current traffic total is determined to exceed the maximum allowable in (h), terminating the additional web scraping request.

[0289] The method is presented, further comprising:

(g) determining whether the target website has refused to serve the web request from the proxy, wherein steps (b)-(f) are conducted when the target website is determined in (g) not to have refused to serve the web request from the proxy.

[0290] The method is presented, further comprising:

(i) when the target website is determined in (g) to have refused to serve the web request from the proxy, retrying to send the web request to the target website via a different proxy.

[0291] The method is presented, further comprising:

(g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate the web request to extract data from the target website,

wherein the generating (b) comprises generating the web request according to the instructions in the selected scraper, and

wherein the counting (e) comprises counting the amount of data in the received content to determine a current traffic total retrieved by the scraper for the client.

[0292] The method is presented, wherein the web request is a second web request, and the received content is a second content, further comprising:

(g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate a first web request and the second web request;

(f) generating the first web request for the target website according to the instructions;

(g) transmitting the first web request such that the web request reaches the target website via the proxy; and

(h) in response to the first web request, receiving, via the proxy, a first content including a data transmitted from the target website via the proxy,

wherein the generating (b) comprises generating, based on the data, the second web request according to the instructions in the selected scraper.

- 64 -

- [0293] The method is presented, wherein the counting (e) comprises excluding an amount of data in the first content to determine the current traffic total retrieved by the scraper for the client.
- [0294] The method is presented, wherein the counting (e) comprises determining the amount of data in the received content as compressed for transmission.
- [0295] The method is presented, wherein the counting (e) further comprises:
- (i) determining a type of data represented by the received content;
 - (ii) based on the type of data, determining a compression factor representing an amount of compression expected when the type of data is transmitted over a network; and
 - (ii) based on the compression factor, determining the amount of data in the received content as compressed for transmission.
- [0296] The method is presented, further comprising:
- (g) analyzing the content to determine web addresses for additional content needed to render a web page; and
 - (h) retrieving the additional content from the web addresses,
- wherein the counting (e) comprises including an amount of data in the additional content in the current traffic total for a client of the client computing device.
- [0297] The method is presented, further comprising:
- (g) receiving a request from a client corresponding to the client computing device for an amount of data remaining;
 - (h) determining the amount of data remaining as a difference between the current traffic total and a maximum allowable for the client; and
 - (i) returning the amount of data remaining to the client.
- [0298] A non-transitory computer-readable device having instructions stored thereon is presented that, when executed by at least one computing device, cause the at least one computing device to perform operations, comprising:
- (a) receiving a web scraping request from a client computing device, the web scraping request specifying a target website to capture content from;
 - (b) based on the web scraping request, generating a web request for the target website;

- 65 -

(c) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies;

(d) in response to the web request, receiving, via the proxy, content transmitted from the target website;

(e) counting an amount of data in the received content to determine a current traffic total for a client of the client computing device; and

(f) transmitting the received content to the client computing device.

[0299] The device is presented, the operations further comprising:

(g) determining whether the target website has refused to serve the web request from the proxy, wherein steps (b)-(f) are conducted when the target website is determined in (g) not to have refused to serve the web request from the proxy.

[0300] The device is presented, the operations further comprising:

(h) when the target website is determined in (g) to have refused to serve the web request from the proxy, retrying to send the web request to the target website via a different proxy.

[0301] The device is presented, wherein the web request is a second web request, and the received content is a second content, further comprising:

(g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate a first web request and the second web request;

(h) generating the first web request for the target website according to the instructions;

(i) transmitting the first web request such that the web request reaches the target website via the proxy; and

(j) in response to the first web request, receiving, via the proxy, a first content including a data transmitted from the target website via the proxy,

wherein the generating (b) comprises generating, based on the data, the second web request according to the instructions in the selected scraper

wherein the counting (e) comprises excluding an amount of data in the first content to determine the current traffic total retrieved by the scraper for the client.

- 66 -

[0302] The device is presented, wherein the counting (e) comprises determining the amount of data in the received content as compressed for transmission.

[0303] The device is presented, wherein the counting (e) further comprises:

- (i) determining a type of data represented by the received content;
- (ii) based on the type of data, determining a compression factor representing an amount of compression expected when the type of data is transmitted over a network; and
- (ii) based on the compression factor, determining the amount of data in the received content as compressed for transmission.

[0304] The device is presented, further comprising:

- (g) analyzing the content to determine web addresses for additional content needed to render a web page; and
- (h) retrieving the additional content from the web addresses, wherein the counting (e) comprises including an amount of data in the additional content in the current traffic total for a client of the client computing device.

[0305] A computer-implemented method for determining which servers are available to process web scraping jobs is presented, comprising:

- repeatedly checking health of each of a plurality of database servers;
- based on the health checks, determine whether each of a plurality of database servers are to be enabled or disabled in a table, the plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;
- when a web scraping request is received from a client computing device:
 - selecting one of the database servers identified as enabled in the table; and
 - sending a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job.

[0306] The method is presented, wherein each of the repeated checking comprises, for each of the plurality of database servers, connecting to the portion of the job database for the respective database server.

- 67 -

- [0307] The method is presented, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.
- [0308] The method is presented, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a number of messages queued within the respective database server's message broker.
- [0309] The method is presented, wherein each of the plurality of database servers is a shard managing storage in a horizontal partition of the jobs database.
- [0310] The method is presented, wherein each of the plurality of database servers do not synchronize states to one another.
- [0311] The method is presented, wherein the plurality of database servers are executed by a plurality of different computing devices.
- [0312] The method is presented, further comprising:
determining whether a number of database servers that are disabled in the plurality of database servers exceeds a threshold; and
when the number of database servers that are disabled exceeds the threshold, alerting an administrator.
- [0313] A non-transitory computer-readable device having instructions stored thereon is presented that, when executed by at least one computing device, cause the at least one computing device to perform operations, comprising:
repeatedly checking health each of a plurality of database servers;
based on the health checks, determining whether each of a plurality of database servers are to be enabled or disabled in a table, the plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;
when a web scraping request is received from a client computing device:

- 68 -

selecting one of the database servers identified as enabled in the table; and
sending a job description specified by the web scraping request to the
selected database server for storage in the job database as a pending web scraping job.

[0314] The device is presented, wherein each of the repeatedly checking comprises, for each of the plurality of database servers, connecting to the portion of the job database for the respective database server.

[0315] The device is presented, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.

[0316] The device is presented, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a number of messages queued within the respective database server's message broker.

[0317] The device is presented, wherein each of the plurality of database servers is a shard managing storage in a horizontal partition of the jobs database.

[0318] The device is presented, wherein each of the plurality of database servers do not synchronize states to one another.

[0319] The device is presented, wherein the plurality of database servers are executed by a plurality of different computing devices.

[0320] The device is presented, further comprising:
determining whether a number of database servers that are disabled in the plurality of database servers exceeds a threshold; and
when the number of database servers that are disabled exceeds the threshold, alerting an administrator.

[0321] A system for determining which servers are available to process web scraping jobs is presented, comprising:
a processor;

- 69 -

a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;

a memory that stores the job database;

a plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of the job database;

a database monitor configured to repeatedly check health of each of the plurality of database servers and, based on the results of the health checks, determine whether each of the plurality of database servers are to be enabled or disabled in a table;

a database server selector configured to, when a web scraping request is received from a client computing device, select one of the database servers identified as enabled in the table; and

a request intake manager configured to send a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job.

[0322] The system is presented, wherein the database monitor is configured to, for each of the plurality of database servers, check a connection between the request intake manager and the jobs database.

[0323] The system is presented, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the job database, and the database monitor is configured to, for each of the plurality of database servers, check a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.

[0324] The system is presented, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and the database monitor is configured to, for each of the plurality of database servers, check a messages queued within the respective database server's message broker.

[0325] A computer-implemented method for distributing web scraping requests over a plurality of data centers is presented, comprising:

(a) receiving an API request from a client computing device, the API request relating to a web scraping job that was previously requested of a web scraping system;

- 70 -

(b) identifying a hostname addressing a host at a data center of a plurality of data centers executing the web scraping system, the data center servicing the web scraping job;

(c) directing the API request to the identified hostname; and

(d) at the identified data center, servicing the API request.

[0326] The method is presented, wherein the API request is a second API request, further comprising, at the identified data center:

receiving a first API request to request that the web scraping system scrape content from a target website via a proxy, wherein the first API request is received prior to the second API request;

generating a web request for the target website as specified in the first API request;

transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies; and

in response to the web request, receiving scraped content from the target website via the proxy.

[0327] The method is presented, wherein the second API request is to retrieve state of the web scraping job specified by the API request.

[0328] The method is presented, wherein the second API request is to retrieve the received scraped content.

[0329] The method is presented, further comprising:

sending a message to the client computing device, the message specifying a sub-domain for the client computing device to retrieve the scraped content.

[0330] The method is presented, wherein the hostname addresses a load balancer that balances requests among a plurality of servers that accept web scraping requests.

[0331] The method is presented, wherein steps (a)-(c) are executed at a data center balancer that selects which data center of the plurality of data centers to service the web scraping job.

[0332] The method is presented, wherein the data center is at a first data center wherein steps (a)-(c) are executed at a second data center different from the first data center.

- 71 -

- [0333] The method is presented, wherein the directing (c) comprises sending an instruction to the client computing device directing the client computing device to resend the API request to the sub-domain.
- [0334] The method is presented, wherein the directing (c) comprises forwarding the API request to the sub-domain for the first data center processing the request, further comprising:
- at the second data center, forwarding the scraped content from the first data center to the client computing device.
- [0335] The method is presented, wherein the plurality of data centers executing the web scraping system each operate to accept job requests independently of one another, wherein status and results from the API request are stored on one of a plurality of independent, horizontally partitioned database shards.
- [0336] A non-transitory computer-readable device having instructions stored thereon is presented that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:
- (a) receiving an API request from a client computing device, the API request relating to a web scraping job that was previously requested of a web scraping system;
 - (b) identifying a hostname addressing a host at a data center of a plurality of data centers executing the web scraping system, the data center servicing the web scraping job; and
 - (c) directing the API request to the identified hostname for servicing at the identified data center.
- [0337] The device is presented, wherein the API request is a second API request, further comprising, at the identified data:
- receiving a first API request to request that the web scraping system scrape content from a target website via a proxy, wherein the first API request is received prior to the second API request;
 - generating a web request for the target website as specified in the first API request;
 - transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies; and

- 72 -

in response to the web request, receiving scraped content from the target website via the proxy.

[0338] The device is presented, wherein the second API request is to retrieve the received scraped content.

[0339] The device is presented, the operations further comprising:
sending a message to the client computing device, the message specifying the sub-domain for the client computing device to retrieve the scraped content.

[0340] The device is presented, wherein the plurality of data centers executing the web scraping system each operate to accept job requests independently of one another, wherein status and results from the API request are stored on one of a plurality of independent, horizontally partitioned database shards.

[0341] The device is presented, wherein the hostname addresses a load balancer that balances requests among a plurality of servers that accept web scraping requests.

[0342] A system for distributing web scraping requests over a plurality of data centers, comprising:

at least one processor and memory;

a data center balancer configured to (i) receive a first API request to scrape content from a target website and a second API request to retrieve the content once scraped, (ii) direct the first API request to a data center selected from a plurality of data centers, and (iii) direct the second API request to the data center; and

a request intake manager configured to (i) in response to the first API request, initiate web scraping from the target website and to (ii) in response to the second API request, return content scraped from the target website to a client computing device.

[0343] A system for distributing web scraping requests over a plurality of data centers, comprising:

at least one processor and memory;

a data center balancer configured to (i) receive a first API request to scrape content from a target website and a second API request to retrieve the content once scraped, (ii) direct the first API request to a first data center selected from a plurality of data centers, and (iii) direct the second API request to a second data center selected from the plurality of data centers;

- 73 -

a first request intake manager executed at the first data center, the first request intake manager configured to (i) in response to the first API request, initiate web scraping from the target website and to (ii) in response to the second API request, return content scraped from the target website to a client computing device; and

a second request intake manager executed at the second data center, the second request intake manager configured to in response to the second API request, (i) determine which of the plurality of data centers processed the first API request, and to (ii) direct the second API request to the determined data center.

[0344] The system is presented further comprising:

a first load balancer executed at the first data center, the first load balancer configured to, in response to the first API request: (i) select the first request intake manager from a first plurality of first request intake managers executed at the first data center, and (ii) direct the first API request to the selected first request intake manager; and

a second load balancer executed at the second data center, the second load balancer configured to, in response to the second API request: (i) select the second request intake manager from a second plurality of first request intake managers executed at the second data center, and (ii) direct the second API request to the selected second request intake manager;

[0345] A computer-implemented method for securing a web scraping system is presented, comprising:

at an entry point to the web scraping system, performing the following:

(a) validating credentials received with an API request from a client computing device, the API request asking that the web scraping system scrape content from a target website;

(b) when the credentials are validated, generating a token indicating an identity of a client associated with the credentials;

(c) transmitting the API request along with the token to a server configured to initiate a scraping process on the web scraping system;

at the server configured to initiate the web scraping system;

- 74 -

(d) analyzing the token to determine whether the client is authorized to conduct the request; and

(e) when the client is authorized, causing the web scraping system to scrape the target website.

[0346] The method is presented, further comprising:

(f) passing the API request between a plurality of servers, each configured to perform a function of the web scraping system, the server configured to initiate the web scraping system being included in the plurality of servers;

at each of the respective servers:

(g) analyzing the token to determine whether the client is authorized to conduct the function performed by the respective server; and

(e) when the client is authorized to conduct the function, performing the function.

[0347] The method is presented, wherein the plurality of servers includes a server configured to service API requests formatted as a web proxy request.

[0348] The method is presented, wherein the plurality of servers includes a server configured to service synchronous API requests, leaving a connection between the web scraping system and the client computing device open while the web scraping system scrapes the target website.

[0349] The method is presented, wherein the plurality of servers includes a server configured to service asynchronous API requests, closing a connection between the web scraping system and the client computing device before the web scraping system scrapes the target website.

[0350] The method is presented, wherein the generating (b) comprises generating the token to include a role of the client.

[0351] The method is presented, wherein the generating (b) comprises generating the token to include a digital signature that cryptographically guarantees that the identity of the client has not been tampered with.

[0352] The method is presented, wherein the API request is a first API request, and the token is a first token, further comprising:

- 75 -

(f) validating credentials received with a second API request, the second API request asking to retrieve content that the web scraping system has previously scraped from the target website;

(g) when the credentials are validated, generating a second token indicating an identity of a client associated with the credentials received with the second API request;

(h) determining whether the first and second tokens indicate that the first and second API requests came from the client; and

(i) when the first and second tokens indicate that the first and second API requests came from the client, returning the scraped content in response to the second API request.

[0353] The method is presented, further comprising, when the first and second tokens do not indicate that the first and second requests came from the client, refusing to return the scraped content.

[0354] The method is presented, wherein the entry point is a load balancer that selects the server from a plurality of parallel servers.

[0355] A non-transitory computer-readable device having instructions stored thereon is presented that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:

at an entry point to a web scraping system, performing the following:

(a) validating credentials received with an API request from a client computing device, the API request asking that the web scraping system scrape a target website;

(b) when the credentials are validated, generating a token indicating an identity of a client associated with the credentials;

(c) transmitting the API request along with the token to a server configured to initiate a scraping process on the web scraping system;

at the server configured to initiate the web scraping system:

(d) analyzing the token to determine whether the client is authorized to conduct the request; and

(e) when the client is authorized, causing the web scraping system to scrape the target website.

- 76 -

[0356] The device is presented, the operations further comprising:

- (f) passing the API request between a plurality of servers, each configured to perform a function of the web scraping system, the server configured to initiate the web scraping system being included in the plurality of servers;
- at each of the respective servers:
- (g) analyzing the token to determine whether the client is authorized to conduct the function performed by the respective server; and
- (e) when the client is authorized to conduct the function, performing the function.

[0357] The device is presented, wherein the plurality of servers includes a server configured to service API requests formatted as a web proxy request.

[0358] The device is presented, wherein the plurality of servers includes a server configured to service synchronous API requests, leaving a connection between the web scraping system and the client computing device open while the web scraping system scrapes the target website.

[0359] The device is presented, wherein the plurality of servers includes a server configured to service asynchronous API requests, closing a connection between the web scraping system and the client computing device before the web scraping system scrapes the target website.

[0360] The device is presented, wherein the generating (b) comprises generating the token to include a role of the client.

[0361] The device is presented, wherein the generating (b) comprises generating the token to include a digital signature that cryptographically guarantees that the identity of the client has not been tampered with.

[0362] The device is presented, wherein the API request is a first API request, and the token is a first token, the operations further comprising:

- (f) validating credentials received with a second API request, the second API request asking to retrieve content that the web scraping system has previously scraped from the target website;
- (g) when the credentials are validated, generating a second token indicating an identity of a client associated with the credentials received with the second API request;

- 77 -

(h) determining whether the first and second tokens indicate that the first and second requests came from the client; and

(e) when the first and second tokens indicate that the first and second requests came from the client, returning the scraped content in response to the second API request.

[0363] The device is presented, the operations further comprising, when the first and second tokens do not indicate that the first and second requests came from the client, refusing to return the scraped content.

[0364] The device is presented, wherein the entry point is a load balancer that selects the server from a plurality of parallel servers.

[0365] A method for aggregating web results scraped from a plurality of data centers is presented, comprising:

at a first data center of the plurality of data centers, comprising:

(a) generating a first web request for a target website;
(b) transmitting the first web request to the target website;
(c) in response to the first web request, receiving a first content from the target website;

(d) storing the first content in an archive;
(e) storing a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;

(f) analyzing a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;

at a second data center of the plurality of data centers:

(g) generating a second web request for the target website;
(h) transmitting the second web request to the target website;
(i) in response to the second web request, receiving a second content from the target website;

(j) storing the second content in the archive;
(k) storing a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;

- 78 -

(l) analyzing the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and

(m) in response to the analyzing (f) and (l), providing to the client a file with the first and second links.

[0366] The method is presented, wherein the providing (m) comprises:
storing the file in a database separate from the archive; and
providing the client read access to the database.

[0367] The method is presented, wherein the file includes the metadata of the first and second entries.

[0368] The method is presented, further comprising:

(n) parsing a response from the target website to the first web request to determine the first content; and

(o) parsing a response from the target website to the second web request to determine the second content.

[0369] The method is presented, further comprising:

(n) determining whether the steps (a)-(e) have occurred during a recent time period, wherein steps (j)-(k) occur only when the steps (a)-(e) have not occurred during the recent time period.

[0370] The method is presented, wherein the determining (n) comprises, at the second data center:

(i) receiving, from the first data center, a data feed with additions to the first index file;

(ii) storing the additions in a prior scraped file; and

(ii) comparing the metadata describing the second content to the prior scraped file.

[0371] The method is presented, further comprising removing metadata from the prior scraped file at an expiration date selected based on the recent time period.

[0372] A non-transitory computer-readable device having instructions stored thereon is presented that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:

at a first data center of the plurality of data centers, comprising:

- 79 -

- (a) generating a first web request for a target website;
- (b) transmitting the first web request to the target website;
- (c) in response to the first web request, receiving a first content from the target website;
- (d) storing the first content in an archive;
- (e) storing a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;
- (f) analyzing a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;
at a second data center of the plurality of data centers:
 - (g) generating a second web request for the target website;
 - (h) transmitting the second web request to the target website;
 - (i) in response to the second web request, receiving a second content from the target website;
 - (j) storing the second content in the archive;
 - (k) storing a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;
 - (l) analyzing the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and
 - (m) in response to the analyzing (f) and (l), providing to the client a file with the first and second links.

[0373] The device is presented, wherein the providing (m) comprises:
storing the file in a database separate from the archive; and
providing the client read access to the database.

[0374] The device is presented, wherein the file includes the metadata of the first and second entries.

[0375] The device is presented, the operations further comprising:

- (n) parsing a response from the target website to the first web request to determine the first content; and

- 80 -

(o) parsing a response from the target website to the second web request to determine the second content.

[0376] The device is presented, the operations further comprising:

(n) determining whether the steps (a)-(e) have occurred during a recent time period, wherein steps (j)-(k) occur only when the steps (a)-(e) have not occurred during the recent time period.

[0377] The device is presented, wherein determining (n) comprises, at the second data center:

(i) receiving, from the first data center, a data feed with additions to the first index file;

(ii) storing the additions in a prior scraped file; and

(ii) comparing the metadata describing the second content to the prior scraped file.

[0378] The device is presented, the operations further comprising removing metadata from the prior scraped file at an expiration date selected based on the recent time period.

[0379] A system for aggregating web results scraped from a plurality of data centers is presented, comprising:

a processor;

at least one memory;

an archive database stored in the at least one memory;

a first index file stored in at least one memory at a first data center;

a first scraper at the first data center, the first scraper configured to (i) generate a first web request for a target website, (ii) transmit the first web request to the target website, and (iii) in response to the first web request, receive a first content from the target website;

a first writer at the first data center, the first writer configured to (i) store the first content in an archive and (ii) store a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;

- 81 -

a first dumper at the first data center, the first dumper configured to analyze a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;

a second index file stored in the at least one memory at a second data center;

a second scraper at the second data center, the second scraper configured to (i) generate a second web request for the target website, (ii) transmit the second web request to the target website, and (iii) in response to the second web request, receiving a second content from the target website;

a second writer at the second data center, the second writer configured to (i) store the second content in an archive and (ii) store a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;

a second dumper at the second data center, the second dumper configured to analyze the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and

a dump database configured to provide to the client a file with the first and second links.

[0380] The system is presented, wherein the file includes the metadata of the first and second entries.

[0381] The system is presented, the operations further comprising:

a first parser at the first data center, the first parser configured to parse a response from the target website to the first web request to determine the first content; and

a second parser at the second data center, the second parser configured to parse a response from the target website to the second web request to determine the second content.

[0382] The system is presented, wherein the second writer is configured to determine whether the first writer has recorded the target website in the archive during a recent time period, wherein the second writer is configured to record the target website in the archive only when the first writer has not during the recent time period.

[0383] The system is presented, wherein further comprising, at the second data center:
a prior scraped file stored in the at least one memory;

- 82 -

a feed receiver configured to receive, from the first data center, a data feed with additions to the first index file and to storing the additions in the prior scraped file, wherein the writer is configured to compare the metadata describing the second content to the prior scraped file.

[0384] The system is presented, wherein metadata is removed from the prior scraped file at an expiration date selected based on the recent time period.

[0385] A method for scraping web results via a proxy is presented, comprising:

- (a) receiving an API request from a client computing device, the API request specifying a target website to scrape;
- (b) based on the target website, selecting a scraper specifying how to generate a sequence of HTTP requests to access content requested by the API request;
- (c) based on the selected scraper and received API request, generating an earlier HTTP request in the sequence of HTTP requests;
- (d) transmitting, to the target website and via a proxy selected from a group of proxies, the earlier HTTP request;
- (e) receiving, from the target website and via the proxy selected from the group of proxies, a response to the earlier HTTP request, the response including session data;
- (f) based on the selected scraper and received API request, generating a later HTTP request in the sequence of HTTP requests such that the later HTTP request includes the session data;
- (g) receiving, from the target website and via the proxy selected from the group of proxies, from the a response to the later HTTP request, the response including the scraped web results; and
- (h) transmitting to the client computing device the scraped web results.

[0386] A method as in any of the above claims is presented, further comprising:
according to the API request and the selected scraper, generating a path to use in the earlier HTTP request.

[0387] A method as in any of the above methods is presented, wherein the API request specifies a geographic location, further comprising:

- 83 -

selecting the proxy based on the group of proxies such that an IP address of the selected proxy is located in the geographic location.

[0388] A method as in any of the above claims, wherein the API request specifies a geographic location, further comprising:

determining at least one header value to use in the earlier and later HTTP requests such that the at least one header identifies to the target website that the earlier and later HTTP requests originated from the geographic location.

[0389] A method as in any of the above methods is presented, further comprising:

determining at least one header value to use in the earlier and later HTTP requests to mimic what a browser at the proxy would generate.

[0390] A method as in any of the above claims, further comprising:

determining whether the target website as requested in the API request has been recently scraped from the target website and the recent scraped results stored in a cache;

when the recent scraped results are determined not to be stored in the cache, executing steps (b)-(h); and

when the recent scraped results are determined to be stored in the cache, transmitting the recently scraped results as a response to the API request.

[0391] A method as in any of the above methods is presented, wherein the API request is an HTTP request from the client computing device, wherein the transmitting (h) comprises transmitting the scraped web results as an HTTP response to the HTTP request from the client computing device.

[0392] A method as in any of the above methods is presented, when the receiving (g) is complete transmitting, to the client computing device, a callback message.

[0393] A method as in any of the above claims, wherein the API request is an HTTP proxy protocol request, and wherein the transmitting (h) comprises transmitting the scraped results as a response to the HTTP proxy protocol request.

[0394] A method as in any of the above methods is presented, further comprising: retrieving, by a headless browser, additional content that would be needed to render a web page specified by the scraped results.

[0395] A method as in any of the above methods is presented, further comprising:

- 84 -

selecting the proxy from the group of proxies to prevent any proxy from the group of proxies from being identified as a bot.

[0396] A method as in any of the above methods is presented, the API request specifying a user-agent, further comprising:

determining at least one header value to use in the earlier and later HTTP requests to mimic the user agent.

[0397] A method as in any of the above methods is presented, further comprising:

receiving both synchronous and asynchronous API requests to scrape web data;
and

prioritizing processing of the synchronous API requests above processing of the asynchronous API requests.

[0398] A method as in any of the above methods is presented, wherein the response to the earlier HTTP request includes an HTML page with a form to submit information to a URI, wherein the later HTTP request includes a path with the URI.

[0399] A method as in any of the above methods is presented, wherein the API request includes a value, wherein the response to the earlier HTTP request includes an HTML page with a form with a parameter name, wherein the later HTTP request indicates that the parameter name corresponds to the value.

[0400] A method as in any of the above methods is presented, wherein the API request includes a value matched to one of several parameter names, further comprising normalizing the API request to a common format.

[0401] A method as in any of the above methods is presented, further comprising generating a job identifier based on a server servicing the API request.

[0402] A method as in any of the above methods is presented, further comprising storing status of a web scraping job specified in the API request in an in-memory table.

[0403] A method as in any of the above methods is presented, further comprising storing status of a web scraping job in a horizontally partitioned database.

[0404] A system for scraping web results via a proxy is presented, comprising:
a means for receiving an API request from a client computing device, the API request specifying a target website to scrape;

- 85 -

a means for, based on the target website, selecting a scraper specifying how to generate a sequence of HTTP requests to access content requested by the API request;

a means for, based on the selected scraper and received API request, generating an earlier HTTP request in the sequence of HTTP requests;

a means for transmitting, to the target website and via a proxy selected from a group of proxies, the earlier HTTP request;

a means for receiving, from the target website and via the proxy selected from the group of proxies, a response to the earlier HTTP request, the response including a session data;

a means for, based on the selected scraper and received API request, generating a later HTTP request in the sequence of HTTP requests such that the later HTTP request includes the session data;

a means for receiving, from the target website and via the proxy selected from the group of proxies, from the a response to the later HTTP request, the response including the scraped web results; and

a means for transmitting to the client computing device the scraped web results.

- 86 -

WHAT IS CLAIMED IS:

1. A computer-implemented method for executing a plurality of web scraping requests over time, comprising:
 - (a) during a first time period, receiving a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;
 - (b) placing job descriptions for the plurality of web scraping requests on a queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;
 - (c) incrementally over a second time period longer than the first time period, removing the job descriptions from the queue;
when each respective job description is removed from the queue:
 - (d) generating a web request for the target website as specified in the job description; and
 - (e) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.
2. The method of claim 1, wherein the removing (c) occurs in increments selected to avoid the group of proxies from becoming blocked by the target website.
3. The method according to any one of the preceding claims, in particular according to claim 1, further comprising:
 - (f) in response to the web request, receiving content transmitted from the target website; and
 - (g) transmitting the received content to the client computing device.
4. The method according to any one of the preceding claims, in particular according to claim 1, wherein the plurality of web scraping requests each include a client identifier indicating a client from which the web scraping request originates, further comprising:

- 87 -

- (f) selecting the queue from a plurality of queues based on the client identifier.
- 5. The method of claim 4, wherein the plurality of web scraping requests each further include an identifier of the target website, further comprising:
 - (g) selecting the queue from the plurality of queues based on the identifier of the target website.
- 6. The method according to any one of the claims 4 and 5, in particular according to claim 5, further comprising selecting the queue from the plurality of queues based on whether the job description specifies that web scraping be made from a headless browser.
- 7. The method according to any one of the claims 4 to 6, in particular according to claim 6, further comprising, when the job description requests that web scraping be made from the headless browser:
 - (h) in response to the web request, receiving content specifying a web page, the content transmitted from the target website;
 - (i) analyzing the content to determine web addresses for additional content needed to render the web page; and
 - (j) requesting the additional content from the web addresses.
- 8. The method according to any one of the preceding claims, in particular according to claim 1, wherein incrementally removing (c) comprises repeated removing a job description from the queue at a time window.
- 9. The method of claim 8, wherein incrementally removing (c) comprises repeated removing a plurality of job descriptions from the queue at the time window.

- 88 -

10. The method of claim 9, wherein an amount of job descriptions in the plurality of job descriptions are specified based on a client that sent the plurality of web scraping requests.
11. The method according to any one of the preceding claims, in particular according to claim 1, further comprising:
 - (f) determining that web scraping requests from the client computing device no longer need to be throttled;
 - (g) determining when the queue is empty; and
 - (h) when the queue is determined to be empty, deleting the queue.
12. The method according to any one of the preceding claims, in particular according to claim 1, further comprising:
 - (f) determining which of the plurality of web scraping requests are synchronous requests and which are asynchronous requests; and
 - (g) prioritizing web scraping requests determined in (f) to be synchronous requests over web scraping requests determined in (f) to be asynchronous requests to help ensure that results from synchronous requests are returned before a connection with the client computing device times out.
13. A system for executing a plurality of web scraping requests over time, comprising:
 - at least one processor;
 - a memory configured to store a queue;
 - a request intake manager configured to, during a first time period, receive a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;
 - a throttler configured to place job descriptions for the plurality of web scraping requests on the queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;

- 89 -

a job remover configured to, incrementally over a second time period longer than the first time period, remove the job descriptions from the queue;

a scraper configured to, when each respective job description is removed from the queue, generate a web request for the target website as specified in the job description, and transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.

14. The system of claim 13, wherein the job remover is configured to remove in increments selected to avoid the group of proxies from becoming blocked by the target website.
15. The system according to any one of the claims 13 and 14, in particular according to claim 13, wherein the scraper is configured to, in response to the web request, receive content transmitted from the target website, and the request intake manager is configured to notify the client that the received content is available for transmission to the client computing device;
16. The system according to any one of the claims 13 to 15, in particular according to claim 13, further comprising a memory configured to store a plurality of queues, wherein the plurality of web scraping requests each include a client identifier indicating a client from which the client originates, further comprising:
 - a queue selector that selects the queue from a plurality of queues based on the client identifier.
17. The system of claim 16, wherein the plurality of web scraping requests each further include an identifier of the target website, the queue selector is configured to select the queue from the plurality of queues based on the identifier of the target website.
18. The system according to any one of the claims 16 and 17, in particular according to claim 17, further comprising the queue selector configured to select the queue from the plurality

- 90 -

of queues based on whether the job description specifies that web scraping be made from a headless browser.

19. The system according to any one of the claims 13 to 18, in particular according to claim 13, further comprising:

a throttle configuration database that stores a regular time window associated with the queue, wherein the job remover repeatedly removes a job description from the queue at the time window.

20. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations comprising:

(a) during a first time period, receiving a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;

(b) placing job descriptions for the plurality of web scraping requests on a queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;

(c) incrementally over a second time period longer than the first time period, removing the job descriptions from the queue;

when each respective job description is removed from the queue:

(d) generating a web request for the target website as specified in the job description; and

(e) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.

21. A computer-implemented method for tracking user activity, comprising:

(a) receiving a web scraping request from a client computing device, the web scraping request specifying a target website to capture content from;

- 91 -

- (b) based on the web scraping request, generating a web request for the target website;
 - (c) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies;
 - (d) in response to the web request, receiving, via the proxy, content transmitted from the target website;
 - (e) counting an amount of data in the received content to determine a current traffic total for a client of the client computing device; and
 - (f) transmitting the received content to the client computing device.
- 22. The method of claim 21, further comprising:
 - (g) based on the current traffic total, generating an invoice for a client corresponding to the client computing device.
- 23. The method according to any one of the claims 21 and 22, in particular according to claim 21, wherein the current traffic total is for a time period, further comprising:
 - (g) receiving an additional web scraping request from the client computing device;
 - (h) determining whether the current traffic total exceeds a maximum allowable for a client corresponding to the client computing device; and
 - (i) when the current traffic total is determined to exceed the maximum allowable in (h), refusing to service the additional web scraping request.
- 24. The method according to any one of the claims 21 to 23, in particular according to claim 21, wherein the current traffic total is for a time period, further comprising:
 - (g) receiving an additional web scraping request from the client computing device;
 - (h) determining whether the current traffic total exceeds a maximum allowable for a client corresponding to the client computing device; and

- 92 -

- (i) when the current traffic total is determined to exceed the maximum allowable in (h), terminating the additional web scraping request.
- 25. The method according to any one of the claims 21 to 24, in particular according to claim 21, further comprising:
 - (g) determining whether the target website has refused to serve the web request from the proxy, wherein steps (b)-(f) are conducted when the target website is determined in (g) not to have refused to serve the web request from the proxy.
- 26. The method of claim 25, further comprising:
 - (i) when the target website is determined in (g) to have refused to serve the web request from the proxy, retrying to send the web request to the target website via a different proxy.
- 27. The method according to any one of the claims 21 to 26, in particular according to claim 21, further comprising:
 - (g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate the web request to extract data from the target website,
 - wherein the generating (b) comprises generating the web request according to the instructions in the selected scraper, and
 - wherein the counting (e) comprises counting the amount of data in the received content to determine a current traffic total retrieved by the scraper for the client.
- 28. The method according to any one of the claims 21 to 27, in particular according to claim 21, wherein the web request is a second web request, and the received content is a second content, further comprising:

- 93 -

(g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate a first web request and the second web request;

(f) generating the first web request for the target website according to the instructions;

(g) transmitting the first web request such that the web request reaches the target website via the proxy; and

(h) in response to the first web request, receiving, via the proxy, a first content including a data transmitted from the target website via the proxy,

wherein the generating (b) comprises generating, based on the data, the second web request according to the instructions in the selected scraper.

29. The method of claim 28, wherein the counting (e) comprises excluding an amount of data in the first content to determine the current traffic total retrieved by the scraper for the client.
30. The method according to any one of the claims 21 to 29, in particular according to claim 21, wherein the counting (e) comprises determining the amount of data in the received content as compressed for transmission.
31. The method of claim 29, wherein the counting (e) further comprises:
- (i) determining a type of data represented by the received content;
 - (ii) based on the type of data, determining a compression factor representing an amount of compression expected when the type of data is transmitted over a network; and
 - (ii) based on the compression factor, determining the amount of data in the received content as compressed for transmission.
32. The method according to any one of the claims 21 to 31, in particular according to claim 21, further comprising:

- 94 -

- (g) analyzing the content to determine web addresses for additional content needed to render a web page; and
 - (h) retrieving the additional content from the web addresses,

wherein the counting (e) comprises including an amount of data in the additional content in the current traffic total for a client of the client computing device.
- 33. The method according to any one of the claims 21 to 32, in particular according to claim 21, further comprising:
 - (g) receiving a request from a client corresponding to the client computing device for an amount of data remaining;
 - (h) determining the amount of data remaining as a difference between the current traffic total and a maximum allowable for the client; and
 - (i) returning the amount of data remaining to the client.
- 34. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, comprising:
 - (a) receiving a web scraping request from a client computing device, the web scraping request specifying a target website to capture content from;
 - (b) based on the web scraping request, generating a web request for the target website;
 - (c) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies;
 - (d) in response to the web request, receiving, via the proxy, content transmitted from the target website;
 - (e) counting an amount of data in the received content to determine a current traffic total for a client of the client computing device; and
 - (f) transmitting the received content to the client computing device.
- 35. The device of claim 34, the operations further comprising:

- 95 -

(g) determining whether the target website has refused to serve the web request from the proxy, wherein steps (b)-(f) are conducted when the target website is determined in (g) not to have refused to serve the web request from the proxy.

36. The device of claim 35, the operations further comprising:

(h) when the target website is determined in (g) to have refused to serve the web request from the proxy, retrying to send the web request to the target website via a different proxy.

37. The device according to any of the claims 34 to 36, in particular according to claim 34, wherein the web request is a second web request, and the received content is a second content, further comprising:

(g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate a first web request and the second web request;

(h) generating the first web request for the target website according to the instructions;

(i) transmitting the first web request such that the web request reaches the target website via the proxy; and

(j) in response to the first web request, receiving, via the proxy, a first content including a data transmitted from the target website via the proxy,

wherein the generating (b) comprises generating, based on the data, the second web request according to the instructions in the selected scraper

wherein the counting (e) comprises excluding an amount of data in the first content to determine the current traffic total retrieved by the scraper for the client.

38. The device according to any of the claims 34 to 37, in particular according to claim 34, wherein the counting (e) comprises determining the amount of data in the received content as compressed for transmission.

- 96 -

39. The device according to any of the claims 34 to 38, in particular according to claim 34, wherein the counting (e) further comprises:
- (i) determining a type of data represented by the received content;
 - (ii) based on the type of data, determining a compression factor representing an amount of compression expected when the type of data is transmitted over a network;
- and
- (ii) based on the compression factor, determining the amount of data in the received content as compressed for transmission.
40. The device according to any of the claims 34 to 39, in particular according to claim 34, further comprising:
- (g) analyzing the content to determine web addresses for additional content needed to render a web page; and
 - (h) retrieving the additional content from the web addresses,
- wherein the counting (e) comprises including an amount of data in the additional content in the current traffic total for a client of the client computing device.
41. A computer-implemented method for determining which servers are available to process web scraping jobs, comprising:
- repeatedly checking health of each of a plurality of database servers;
 - based on the health checks, determine whether each of a plurality of database servers are to be enabled or disabled in a table, the plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;
 - when a web scraping request is received from a client computing device:
 - selecting one of the database servers identified as enabled in the table; and
 - sending a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job.

- 97 -

42. The method of claim 41, wherein each of the repeated checking comprises, for each of the plurality of database servers, connecting to the portion of the job database for the respective database server.
43. The method according to any of the claims 41 and 42, in particular according to claim 41, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.
44. The method according to any of the claims 41 to 43, in particular according to claim 41, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a number of messages queued within the respective database server's message broker.
45. The method according to any of the claims 41 to 44, in particular according to claim 41, wherein each of the plurality of database servers is a shard managing storage in a horizontal partition of the jobs database.
46. The method according to any of the claims 41 to 45, in particular according to claim 41, wherein each of the plurality of database servers do not synchronize states to one another.
47. The method according to any of the claims 41 to 46, in particular according to claim 41, wherein the plurality of database servers are executed by a plurality of different computing devices.
48. The method according to any of the claims 41 to 47, in particular according to claim 41, further comprising:

- 98 -

determining whether a number of database servers that are disabled in the plurality of database servers exceeds a threshold; and

when the number of database servers that are disabled exceeds the threshold, alerting an administrator.

49. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, comprising:
- repeatedly checking health each of a plurality of database servers;
 - based on the health checks, determining whether each of a plurality of database servers are to be enabled or disabled in a table, the plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;
 - when a web scraping request is received from a client computing device:
 - selecting one of the database servers identified as enabled in the table; and
 - sending a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job.
50. The device of claim 49, wherein each of the repeatedly checking comprises, for each of the plurality of database servers, connecting to the portion of the job database for the respective database server.
51. The device according to any one of the claims 49 and 50, in particular according to claim 49, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.

- 99 -

52. The device according to any one of the claims 49 to 51, in particular according to claim 49, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a number of messages queued within the respective database server's message broker.
53. The device according to any one of the claims 49 to 52, in particular according to claim 49, wherein each of the plurality of database servers is a shard managing storage in a horizontal partition of the jobs database.
54. The device according to any one of the claims 49 and 53, in particular according to claim 49, wherein each of the plurality of database servers do not synchronize states to one another.
55. The device according to any one of the claims 49 and 54, in particular according to claim 49, wherein the plurality of database servers are executed by a plurality of different computing devices.
56. The device according to any one of the claims 49 and 55, in particular according to claim 49, further comprising:
- determining whether a number of database servers that are disabled in the plurality of database servers exceeds a threshold; and
 - when the number of database servers that are disabled exceeds the threshold, alerting an administrator.
57. A system for determining which servers are available to process web scraping jobs, comprising:
- a processor;
 - a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;

- 100 -

a memory that stores the job database;

a plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of the job database;

a database monitor configured to repeatedly check health of each of the plurality of database servers and, based on the results of the health checks, determine whether each of the plurality of database servers are to be enabled or disabled in a table;

a database server selector configured to, when a web scraping request is received from a client computing device, select one of the database servers identified as enabled in the table; and

a request intake manager configured to send a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job.

58. The system of claim 57, wherein the database monitor is configured to, for each of the plurality of database servers, check a connection between the request intake manager and the jobs database.
59. The system according to any one of the claims 57 and 58, in particular according to claim 57, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the job database, and the database monitor is configured to, for each of the plurality of database servers, check a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.
60. The system according to any one of the claims 57 to 59, in particular according to claim 57, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and the database monitor is configured to, for each of the plurality of database servers, check a messages queued within the respective database server's message broker.

- 101 -

61. A computer-implemented method for distributing web scraping requests over a plurality of data centers, comprising:
- (a) receiving an API request from a client computing device, the API request relating to a web scraping job that was previously requested of a web scraping system;
 - (b) identifying a hostname addressing a host at a data center of a plurality of data centers executing the web scraping system, the data center servicing the web scraping job;
 - (c) directing the API request to the identified hostname; and
 - (d) at the identified data center, servicing the API request.
62. The method of claim 61, wherein the API request is a second API request, further comprising, at the identified data center:
- receiving a first API request to request that the web scraping system scrape content from a target website via a proxy, wherein the first API request is received prior to the second API request;
 - generating a web request for the target website as specified in the first API request;
 - transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies; and
 - in response to the web request, receiving scraped content from the target website via the proxy.
63. The method according to any one of the claims 61 and 62, in particular according to claim 62, wherein the second API request is to retrieve state of the web scraping job specified by the API request.
64. The method of claim 63, wherein the second API request is to retrieve the received scraped content.

- 102 -

65. The method according to any one of the claims 61 to 64, in particular according to claim 63, further comprising:
- sending a message to the client computing device, the message specifying a sub-domain for the client computing device to retrieve the scraped content.
66. The method according to any one of the claims 61 to 65, in particular according to claim 61, wherein the hostname addresses a load balancer that balances requests among a plurality of servers that accept web scraping requests.
67. The method according to any one of the claims 61 to 66, in particular according to claim 61, wherein steps (a)-(c) are executed at a data center balancer that selects which data center of the plurality of data centers to service the web scraping job.
68. The method according to any one of the claims 61 to 67, in particular according to claim 61, wherein the data center is at a first data center wherein steps (a)-(c) are executed at a second data center different from the first data center.
69. The method of claim 68, wherein the directing (c) comprises sending an instruction to the client computing device directing the client computing device to resend the API request to the sub-domain.
70. The method of claim 68, wherein the directing (c) comprises forwarding the API request to the sub-domain for the first data center processing the request, further comprising:
- at the second data center, forwarding the scraped content from the first data center to the client computing device.
71. The method according to any one of the claims 61 and 70, in particular according to claim 61, wherein the plurality of data centers executing the web scraping system each operate to accept job requests independently of one another, wherein status and results

- 103 -

from the API request are stored on one of a plurality of independent, horizontally partitioned database shards.

72. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:
- (a) receiving an API request from a client computing device, the API request relating to a web scraping job that was previously requested of a web scraping system;
 - (b) identifying a hostname addressing a host at a data center of a plurality of data centers executing the web scraping system, the data center servicing the web scraping job; and
 - (c) directing the API request to the identified hostname for servicing at the identified data center.
73. The device of claim 72, wherein the API request is a second API request, further comprising, at the identified data:
- receiving a first API request to request that the web scraping system scrape content from a target website via a proxy, wherein the first API request is received prior to the second API request;
 - generating a web request for the target website as specified in the first API request;
 - transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies; and
 - in response to the web request, receiving scraped content from the target website via the proxy.
74. The device according to any of the claims 72 and 73, in particular according to claim 72, wherein the second API request is to retrieve the received scraped content.

- 104 -

75. The device according to any of the claims 72 to 74, in particular according to claim 72, the operations further comprising:
- sending a message to the client computing device, the message specifying the sub-domain for the client computing device to retrieve the scraped content.
76. The device according to any of the claims 72 to 75, in particular according to claim 72, wherein the plurality of data centers executing the web scraping system each operate to accept job requests independently of one another, wherein status and results from the API request are stored on one of a plurality of independent, horizontally partitioned database shards.
77. The device according to any of the claims 72 to 76, in particular according to claim 72, wherein the hostname addresses a load balancer that balances requests among a plurality of servers that accept web scraping requests.
78. A system for distributing web scraping requests over a plurality of data centers, comprising:
- at least one processor and memory;
 - a data center balancer configured to (i) receive a first API request to scrape content from a target website and a second API request to retrieve the content once scraped, (ii) direct the first API request to a data center selected from a plurality of data centers, and (iii) direct the second API request to the data center; and
 - a request intake manager configured to (i) in response to the first API request, initiate web scraping from the target website and to (ii) in response to the second API request, return content scraped from the target website to a client computing device.
79. A system for distributing web scraping requests over a plurality of data centers, comprising:
- at least one processor and memory;

- 105 -

a data center balancer configured to (i) receive a first API request to scrape content from a target website and a second API request to retrieve the content once scraped, (ii) direct the first API request to a first data center selected from a plurality of data centers, and (iii) direct the second API request to a second data center selected from the plurality of data centers;

a first request intake manager executed at the first data center, the first request intake manager configured to (i) in response to the first API request, initiate web scraping from the target website and to (ii) in response to the second API request, return content scraped from the target website to a client computing device; and

a second request intake manager executed at the second data center, the second request intake manager configured to in response to the second API request, (i) determine which of the plurality of data centers processed the first API request, and to (ii) direct the second API request to the determined data center.

80. The system of claim, 79 further comprising:

a first load balancer executed at the first data center, the first load balancer configured to, in response to the first API request: (i) select the first request intake manager from a first plurality of first request intake managers executed at the first data center, and (ii) direct the first API request to the selected first request intake manager; and

a second load balancer executed at the second data center, the second load balancer configured to, in response to the second API request: (i) select the second request intake manager from a second plurality of first request intake managers executed at the second data center, and (ii) direct the second API request to the selected second request intake manager;

81. A computer-implemented method for securing a web scraping system, comprising:

at an entry point to the web scraping system, performing the following:

- 106 -

- (a) validating credentials received with an API request from a client computing device, the API request asking that the web scraping system scrape content from a target website;
- (b) when the credentials are validated, generating a token indicating an identity of a client associated with the credentials;
- (c) transmitting the API request along with the token to a server configured to initiate a scraping process on the web scraping system;
at the server configured to initiate the web scraping system:
- (d) analyzing the token to determine whether the client is authorized to conduct the request; and
- (e) when the client is authorized, causing the web scraping system to scrape the target website.

82. The method of claim 81, further comprising:

- (f) passing the API request between a plurality of servers, each configured to perform a function of the web scraping system, the server configured to initiate the web scraping system being included in the plurality of servers;
at each of the respective servers:
- (g) analyzing the token to determine whether the client is authorized to conduct the function performed by the respective server; and
- (e) when the client is authorized to conduct the function, performing the function.

83. The method of claim 82, wherein the plurality of servers includes a server configured to service API requests formatted as a web proxy request.

84. The method according to any one of the claims 82 and 83, in particular according to claim 82, wherein the plurality of servers includes a server configured to service synchronous API requests, leaving a connection between the web scraping system and

- 107 -

the client computing device open while the web scraping system scrapes the target website.

85. The method according to any one of the claims 82 to 84, in particular according to claim 82, wherein the plurality of servers includes a server configured to service asynchronous API requests, closing a connection between the web scraping system and the client computing device before the web scraping system scrapes the target website.
86. The method according to any one of the claims 82 to 85, in particular according to claim 82, wherein the generating (b) comprises generating the token to include a role of the client.
87. The method according to any one of the claims 82 to 86, in particular according to claim 82, wherein the generating (b) comprises generating the token to include a digital signature that cryptographically guarantees that the identity of the client has not been tampered with.
88. The method according to any one of the claims 82 to 87, in particular according to claim 81, wherein the API request is a first API request, and the token is a first token, further comprising:
 - (f) validating credentials received with a second API request, the second API request asking to retrieve content that the web scraping system has previously scraped from the target website;
 - (g) when the credentials are validated, generating a second token indicating an identity of a client associated with the credentials received with the second API request;
 - (h) determining whether the first and second tokens indicate that the first and second API requests came from the client; and
 - (i) when the first and second tokens indicate that the first and second API requests came from the client, returning the scraped content in response to the second API request.

- 108 -

89. The method of claim 88, further comprising, when the first and second tokens do not indicate that the first and second requests came from the client, refusing to return the scraped content.
90. The method according to any one of the claims 82 to 89, in particular according to claim 81, wherein the entry point is a load balancer that selects the server from a plurality of parallel servers.
91. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:
- at an entry point to a web scraping system, performing the following:
 - (a) validating credentials received with an API request from a client computing device, the API request asking that the web scraping system scrape a target website;
 - (b) when the credentials are validated, generating a token indicating an identity of a client associated with the credentials;
 - (c) transmitting the API request along with the token to a server configured to initiate a scraping process on the web scraping system;
 - at the server configured to initiate the web scraping system:
 - (d) analyzing the token to determine whether the client is authorized to conduct the request; and
 - (e) when the client is authorized, causing the web scraping system to scrape the target website.
92. The device of claim 91, the operations further comprising:
- (f) passing the API request between a plurality of servers, each configured to perform a function of the web scraping system, the server configured to initiate the web scraping system being included in the plurality of servers;
 - at each of the respective servers:

- 109 -

(g) analyzing the token to determine whether the client is authorized to conduct the function performed by the respective server; and

(e) when the client is authorized to conduct the function, performing the function.

93. The device of claim 92, wherein the plurality of servers includes a server configured to service API requests formatted as a web proxy request.
94. The device according to any one of the claims 92 and 93, in particular according to claim 92, wherein the plurality of servers includes a server configured to service synchronous API requests, leaving a connection between the web scraping system and the client computing device open while the web scraping system scrapes the target website.
95. The device according to any one of the claims 92 to 94, in particular according to claim 92, wherein the plurality of servers includes a server configured to service asynchronous API requests, closing a connection between the web scraping system and the client computing device before the web scraping system scrapes the target website.
96. The device according to any one of the claims 92 to 95, in particular according to claim 92, wherein the generating (b) comprises generating the token to include a role of the client.
97. The device according to any one of the claims 92 and 96, in particular according to claim 92, wherein the generating (b) comprises generating the token to include a digital signature that cryptographically guarantees that the identity of the client has not been tampered with.
98. The device according to any one of the claims 92 and 97, in particular according to claim 91, wherein the API request is a first API request, and the token is a first token, the operations further comprising:

- 110 -

- (f) validating credentials received with a second API request, the second API request asking to retrieve content that the web scraping system has previously scraped from the target website;
 - (g) when the credentials are validated, generating a second token indicating an identity of a client associated with the credentials received with the second API request;
 - (h) determining whether the first and second tokens indicate that the first and second requests came from the client; and
 - (e) when the first and second tokens indicate that the first and second requests came from the client, returning the scraped content in response to the second API request.
99. The device of claim 98, the operations further comprising, when the first and second tokens do not indicate that the first and second requests came from the client, refusing to return the scraped content.
100. The device according to any one of the claims 92 and 99, in particular according to claim 91, wherein the entry point is a load balancer that selects the server from a plurality of parallel servers.
101. A method for aggregating web results scraped from a plurality of data centers, comprising:
- at a first data center of the plurality of data centers, comprising:
 - (a) generating a first web request for a target website;
 - (b) transmitting the first web request to the target website;
 - (c) in response to the first web request, receiving a first content from the target website;
 - (d) storing the first content in an archive;
 - (e) storing a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;

- 111 -

(f) analyzing a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;

at a second data center of the plurality of data centers:

(g) generating a second web request for the target website;

(h) transmitting the second web request to the target website;

(i) in response to the second web request, receiving a second content from the target website;

(j) storing the second content in the archive;

(k) storing a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;

(l) analyzing the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and

(m) in response to the analyzing (f) and (l), providing to the client a file with the first and second links.

102. The method of claim 101, wherein the providing (m) comprises:

storing the file in a database separate from the archive; and
providing the client read access to the database.

103. The method according to any one of the claims 101 and 102, in particular according to claim 101, wherein the file includes the metadata of the first and second entries.

104. The method o according to any one of the claims 101 to 103, in particular according to f claim 101, further comprising:

(n) parsing a response from the target website to the first web request to determine the first content; and

(o) parsing a response from the target website to the second web request to determine the second content.

- 112 -

105. The method according to any one of the claims 101 and 104, in particular according to claim 101, further comprising:
- (n) determining whether the steps (a)-(e) have occurred during a recent time period, wherein steps (j)-(k) occur only when the steps (a)-(e) have not occurred during the recent time period.
106. The method of claim 105, wherein the determining (n) comprises, at the second data center:
- (i) receiving, from the first data center, a data feed with additions to the first index file;
 - (ii) storing the additions in a prior scraped file; and
 - (ii) comparing the metadata describing the second content to the prior scraped file.
107. The method of claim 106, further comprising removing metadata from the prior scraped file at an expiration date selected based on the recent time period.
108. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:
- at a first data center of the plurality of data centers, comprising:
 - (a) generating a first web request for a target website;
 - (b) transmitting the first web request to the target website;
 - (c) in response to the first web request, receiving a first content from the target website;
 - (d) storing the first content in an archive;
 - (e) storing a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;

- 113 -

(f) analyzing a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;

at a second data center of the plurality of data centers:

(g) generating a second web request for the target website;

(h) transmitting the second web request to the target website;

(i) in response to the second web request, receiving a second content from the target website;

(j) storing the second content in the archive;

(k) storing a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;

(l) analyzing the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and

(m) in response to the analyzing (f) and (l), providing to the client a file with the first and second links.

109. The device of claim 108, wherein the providing (m) comprises:

storing the file in a database separate from the archive; and
providing the client read access to the database.

110. The device according to any one of the claims 108 and 109, in particular according to claim 108, wherein the file includes the metadata of the first and second entries.

111. The device according to any one of the claims 108 to 110, in particular according to claim 108, the operations further comprising:

(n) parsing a response from the target website to the first web request to determine the first content; and

(o) parsing a response from the target website to the second web request to determine the second content.

- 114 -

112. The device according to any one of the claims 108 to 111, in particular according to claim 108, the operations further comprising:
- (n) determining whether the steps (a)-(e) have occurred during a recent time period, wherein steps (j)-(k) occur only when the steps (a)-(e) have not occurred during the recent time period.
113. The device of claim 112, wherein determining (n) comprises, at the second data center:
- (i) receiving, from the first data center, a data feed with additions to the first index file;
 - (ii) storing the additions in a prior scraped file; and
 - (ii) comparing the metadata describing the second content to the prior scraped file.
114. The device of claim 113, the operations further comprising removing metadata from the prior scraped file at an expiration date selected based on the recent time period.
115. A system for aggregating web results scraped from a plurality of data centers, comprising:
- a processor;
 - at least one memory;
 - an archive database stored in the at least one memory;
 - a first index file stored in at least one memory at a first data center;
 - a first scraper at the first data center, the first scraper configured to (i) generate a first web request for a target website, (ii) transmit the first web request to the target website, and (iii) in response to the first web request, receive a first content from the target website;
 - a first writer at the first data center, the first writer configured to (i) store the first content in an archive and (ii) store a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;

- 115 -

a first dumper at the first data center, the first dumper configured to analyze a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;

a second index file stored in the at least one memory at a second data center;

a second scraper at the second data center, the second scraper configured to (i) generate a second web request for the target website, (ii) transmit the second web request to the target website, and (iii) in response to the second web request, receiving a second content from the target website;

a second writer at the second data center, the second writer configured to (i) store the second content in an archive and (ii) store a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;

a second dumper at the second data center, the second dumper configured to analyze the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and

a dump database configured to provide to the client a file with the first and second links.

116. The system of claim 115, wherein the file includes the metadata of the first and second entries.

117. The system according to any one of the claims 115 and 116, in particular according to claim 115, the operations further comprising:

a first parser at the first data center, the first parser configured to parse a response from the target website to the first web request to determine the first content; and

a second parser at the second data center, the second parser configured to parse a response from the target website to the second web request to determine the second content.

- 116 -

118. The system according to any one of the claims 115 to 117, in particular according to claim 115, wherein the second writer is configured to determine whether the first writer has recorded the target website in the archive during a recent time period, wherein the second writer is configured to record the target website in the archive only when the first writer has not during the recent time period.
119. The system according to any one of the claims 115 to 118, in particular according to claim 115, wherein further comprising, at the second data center:
- a prior scraped file stored in the at least one memory;
 - a feed receiver configured to receive, from the first data center, a data feed with additions to the first index file and to storing the additions in the prior scraped file, wherein the writer is configured to compare the metadata describing the second content to the prior scraped file.
120. The system according to any one of the claims 115 to 119, in particular according to claim 115, wherein metadata is removed from the prior scraped file at an expiration date selected based on the recent time period.
121. A method for scraping web results via a proxy, comprising:
- (a) receiving an API request from a client computing device, the API request specifying a target website to scrape;
 - (b) based on the target website, selecting a scraper specifying how to generate a sequence of HTTP requests to access content requested by the API request;
 - (c) based on the selected scraper and received API request, generating an earlier HTTP request in the sequence of HTTP requests;
 - (d) transmitting, to the target website and via a proxy selected from a group of proxies, the earlier HTTP request;
 - (e) receiving, from the target website and via the proxy selected from the group of proxies, a response to the earlier HTTP request, the response including a session data;

- 117 -

(f) based on the selected scraper and received API request, generating a later HTTP request in the sequence of HTTP requests such that the later HTTP request includes the session data;

(g) receiving, from the target website and via the proxy selected from the group of proxies, from the a response to the later HTTP request, the response including the scraped web results; and

(h) transmitting to the client computing device the scraped web results.

122. A method as in claim 121, further comprising:

according to the API request and the selected scraper, generating a path to use in the earlier HTTP request.

123. A method as in any of the claims 121 to 122, wherein the API request specifies a geographic location, further comprising:

selecting the proxy based on the group of proxies such that an IP address of the selected proxy is located in the geographic location.

124. A method as in any of the claims 121 to 123, wherein the API request specifies a geographic location, further comprising:

determining at least one header value to use in the earlier and later HTTP requests such that the at least one header identifies to the target website that the earlier and later HTTP requests originated from the geographic location.

125. A method as in any of the claims 121 to 124, further comprising:

determining at least one header value to use in the earlier and later HTTP requests to mimic what a browser at the proxy would generate.

126. A method as in any of the claims 121 to 125, further comprising:

determining whether the target website as requested in the API request has been recently scraped from the target website and the recent scraped results stored in a cache;

- 118 -

when the recent scraped results are determined not to be stored in the cache,
executing steps (b)-(h); and

when the recent scraped results are determined to be stored in the cache,
transmitting the recently scraped results as a response to the API request.

127. A method as in any of the claims 121 to 126, wherein the API request is an HTTP request from the client computing device, wherein the transmitting (h) comprises transmitting the scraped web results as an HTTP response to the HTTP request from the client computing device.
128. A method as in any of the claims 121 to 127, when the receiving (g) is complete transmitting, to the client computing device, a callback message.
129. A method as in any of the claims 121 to 128, wherein the API request is an HTTP proxy protocol request, and wherein the transmitting (h) comprises transmitting the scraped results as a response to the HTTP proxy protocol request.
130. A method as in any of the claims 121 to 129, further comprising:
retrieving, by a headless browser, additional content that would be needed to render a web page specified by the scraped results.
131. A method as in any of the claims 121 to 130, further comprising:
selecting the proxy from the group of proxies to prevent any proxy from the group of proxies from being identified as a bot.
132. A method as in any of the claims 121 to 131, the API request specifying a user-agent, further comprising:
determining at least one header value to use in the earlier and later HTTP requests to mimic the user agent.

- 119 -

133. A method as in any of the claims 121 to 132, further comprising:
receiving both synchronous and asynchronous API requests to scrape web data;
and
prioritizing processing of the synchronous API requests above processing of the asynchronous API requests.
134. A method as in any of the claims 121 to 133, wherein the response to the earlier HTTP request includes an HTML page with a form to submit information to a URI, wherein the later HTTP request includes a path with the URI.
135. A method as in any of the claims 121 to 134, wherein the API request includes a value, wherein the response to the earlier HTTP request includes an HTML page with a form with a parameter name, wherein the later HTTP request indicates that the parameter name corresponds to the value.
136. A method as in any of the claims 121 to 135, wherein the API request includes a value matched to one of several parameter names, further comprising normalizing the API request to a common format.
137. A method as in any of the claims 121 to 136, further comprising generating a job identifier based on a server servicing the API request.
138. A method as in any of the claims 121 to 137, further comprising storing status of a web scraping job specified in the API request in an in-memory table.
139. A method as in any of the claims 121 to 138, further comprising storing status of a web scraping job in a horizontally partitioned database.
140. A system for scraping web results via a proxy, comprising:

- 120 -

a means for receiving an API request from a client computing device, the API request specifying a target website to scrape;

a means for, based on the target website, selecting a scraper specifying how to generate a sequence of HTTP requests to access content requested by the API request;

a means for, based on the selected scraper and received API request, generating an earlier HTTP request in the sequence of HTTP requests;

a means for transmitting, to the target website and via a proxy selected from a group of proxies, the earlier HTTP request;

a means for receiving, from the target website and via the proxy selected from the group of proxies, a response to the earlier HTTP request, the response including a session data;

a means for, based on the selected scraper and received API request, generating a later HTTP request in the sequence of HTTP requests such that the later HTTP request includes the session data;

a means for receiving, from the target website and via the proxy selected from the group of proxies, from the a response to the later HTTP request, the response including the scraped web results; and

a means for transmitting to the client computing device the scraped web results.

AMENDED CLAIMS

received by the International Bureau on 27 December 2022 (27.12.2022)

1. A computer-implemented method for executing a plurality of web scraping requests over time, comprising:
 - (a) during a first time period, receiving a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;
 - (b) placing job descriptions for the plurality of web scraping requests on a queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;
 - (c) incrementally over a second time period longer than the first time period, removing the job descriptions from the queue;
when each respective job description is removed from the queue:
 - (d) generating a web request for the target website as specified in the job description; and
 - (e) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.
2. The method of claim 1, wherein the removing (c) occurs in increments selected to avoid the group of proxies from becoming blocked by the target website.
3. The method according to any one of the preceding claims, in particular according to claim 1, further comprising:
 - (f) in response to the web request, receiving content transmitted from the target website; and
 - (g) transmitting the received content to the client computing device.

4. The method according to any one of the preceding claims, in particular according to claim 1, wherein the plurality of web scraping requests each include a client identifier indicating a client from which the web scraping request originates, further comprising:
 - (f) selecting the queue from a plurality of queues based on the client identifier.
5. The method of claim 4, wherein the plurality of web scraping requests each further include an identifier of the target website, further comprising:
 - (g) selecting the queue from the plurality of queues based on the identifier of the target website.
6. The method according to any one of the claims 4 and 5, in particular according to claim 5, further comprising selecting the queue from the plurality of queues based on whether the job description specifies that web scraping be made from a headless browser.
7. The method according to any one of the claims 4 to 6, in particular according to claim 6, further comprising, when the job description requests that web scraping be made from the headless browser:
 - (h) in response to the web request, receiving content specifying a web page, the content transmitted from the target website;
 - (i) analyzing the content to determine web addresses for additional content needed to render the web page; and
 - (j) requesting the additional content from the web addresses.
8. The method according to any one of the preceding claims, in particular according to claim 1, wherein incrementally removing (c) comprises repeated removing a job description from the queue at a time window.
9. The method of claim 8, wherein incrementally removing (c) comprises repeated removing a plurality of job descriptions from the queue at the time window.

10. The method of claim 9, wherein an amount of job descriptions in the plurality of job descriptions are specified based on a client that sent the plurality of web scraping requests.
11. The method according to any one of the preceding claims, in particular according to claim 1, further comprising:
 - (f) determining that web scraping requests from the client computing device no longer need to be throttled;
 - (g) determining when the queue is empty; and
 - (h) when the queue is determined to be empty, deleting the queue.
12. The method according to any one of the preceding claims, in particular according to claim 1, further comprising:
 - (f) determining which of the plurality of web scraping requests are synchronous requests and which are asynchronous requests; and
 - (g) prioritizing web scraping requests determined in (f) to be synchronous requests over web scraping requests determined in (f) to be asynchronous requests to help ensure that results from synchronous requests are returned before a connection with the client computing device times out.
13. A system for executing a plurality of web scraping requests over time, comprising:
 - at least one processor;
 - a memory configured to store a queue;
 - a request intake manager configured to, during a first time period, receive a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;
 - a throttler configured to place job descriptions for the plurality of web scraping requests on the queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;

a job remover configured to, incrementally over a second time period longer than the first time period, remove the job descriptions from the queue;

a scraper configured to, when each respective job description is removed from the queue, generate a web request for the target website as specified in the job description, and transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.

14. The system of claim 13, wherein the job remover is configured to remove in increments selected to avoid the group of proxies from becoming blocked by the target website.
15. The system according to any one of the claims 13 and 14, in particular according to claim 13, wherein the scraper is configured to, in response to the web request, receive content transmitted from the target website, and the request intake manager is configured to notify the client that the received content is available for transmission to the client computing device;
16. The system according to any one of the claims 13 to 15, in particular according to claim 13, further comprising a memory configured to store a plurality of queues, wherein the plurality of web scraping requests each include a client identifier indicating a client from which the client originates, further comprising:
 - a queue selector that selects the queue from a plurality of queues based on the client identifier.
17. The system of claim 16, wherein the plurality of web scraping requests each further include an identifier of the target website, the queue selector is configured to select the queue from the plurality of queues based on the identifier of the target website.
18. The system according to any one of the claims 16 and 17, in particular according to claim 17, further comprising the queue selector configured to select the queue from the plurality

of queues based on whether the job description specifies that web scraping be made from a headless browser.

19. The system according to any one of the claims 13 to 18, in particular according to claim 13, further comprising:

a throttle configuration database that stores a regular time window associated with the queue, wherein the job remover repeatedly removes a job description from the queue at the time window.

20. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations comprising:

(a) during a first time period, receiving a plurality of web scraping requests from a client computing device, each of the plurality of web scraping requests specifying a target website to capture content from;

(b) placing job descriptions for the plurality of web scraping requests on a queue, the job descriptions each describing at least one web scraping job specified by a request from the plurality of requests;

(c) incrementally over a second time period longer than the first time period, removing the job descriptions from the queue;

when each respective job description is removed from the queue:

(d) generating a web request for the target website as specified in the job description; and

(e) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies.

21. A computer-implemented method for tracking user activity, comprising:

(a) receiving a web scraping request from a client computing device, the web scraping request specifying a target website to capture content from;

- (b) based on the web scraping request, generating a web request for the target website;
 - (c) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies;
 - (d) in response to the web request, receiving, via the proxy, content transmitted from the target website;
 - (e) counting an amount of data in the received content to determine a current traffic total for a client of the client computing device; and
 - (f) transmitting the received content to the client computing device,
- wherein the web request is a second web request, and the received content is a second content, and the method further comprises:
- (g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate a first web request and the second web request;
 - (f) generating the first web request for the target website according to the instructions;
 - (g) transmitting the first web request such that the web request reaches the target website via the proxy; and
 - (h) in response to the first web request, receiving, via the proxy, a first content including a data transmitted from the target website via the proxy,
- wherein the generating (b) comprises generating, based on the data, the second web request according to the instructions in the selected scraper, and
- wherein the counting (e) comprises excluding an amount of data in the first content to determine the current traffic total retrieved by the scraper for the client.
22. The method of claim 21, further comprising:
- (g) based on the current traffic total, generating an invoice for a client corresponding to the client computing device.

23. The method according to any one of the claims 21 and 22, in particular according to claim 21, wherein the current traffic total is for a time period, further comprising:
- (g) receiving an additional web scraping request from the client computing device;
 - (h) determining whether the current traffic total exceeds a maximum allowable for a client corresponding to the client computing device; and
 - (i) when the current traffic total is determined to exceed the maximum allowable in (h), refusing to service the additional web scraping request.
24. The method according to any one of the claims 21 to 23, in particular according to claim 21, wherein the current traffic total is for a time period, further comprising:
- (g) receiving an additional web scraping request from the client computing device;
 - (h) determining whether the current traffic total exceeds a maximum allowable for a client corresponding to the client computing device; and
 - (i) when the current traffic total is determined to exceed the maximum allowable in (h), terminating the additional web scraping request.
25. The method according to any one of the claims 21 to 24, in particular according to claim 21, further comprising:
- (g) determining whether the target website has refused to serve the web request from the proxy, wherein steps (b)-(f) are conducted when the target website is determined in (g) not to have refused to serve the web request from the proxy.
26. The method of claim 25, further comprising:
- (i) when the target website is determined in (g) to have refused to serve the web request from the proxy, retrying to send the web request to the target website via a different proxy.

27. The method according to any one of the claims 21 to 26, in particular according to claim 21, further comprising:
- (g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate the web request to extract data from the target website,
 - wherein the generating (b) comprises generating the web request according to the instructions in the selected scraper, and
 - wherein the counting (e) comprises counting the amount of data in the received content to determine a current traffic total retrieved by the scraper for the client.
28. The method according to any one of the claims 21 to 27, in particular according to claim 21, wherein the counting (e) comprises determining the amount of data in the received content as compressed for transmission.
29. The method of claim 27, wherein the counting (e) further comprises:
- (i) determining a type of data represented by the received content;
 - (ii) based on the type of data, determining a compression factor representing an amount of compression expected when the type of data is transmitted over a network; and
 - (ii) based on the compression factor, determining the amount of data in the received content as compressed for transmission.
30. The method according to any one of the claims 21 to 29, in particular according to claim 21, further comprising:
- (g) analyzing the content to determine web addresses for additional content needed to render a web page; and
 - (h) retrieving the additional content from the web addresses,
 - wherein the counting (e) comprises including an amount of data in the additional content in the current traffic total for a client of the client computing device.

31. The method according to any one of the claims 21 to 30, in particular according to claim 21, further comprising:
- (g) receiving a request from a client corresponding to the client computing device for an amount of data remaining;
 - (h) determining the amount of data remaining as a difference between the current traffic total and a maximum allowable for the client; and
 - (i) returning the amount of data remaining to the client.
32. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, comprising:
- (a) receiving a web scraping request from a client computing device, the web scraping request specifying a target website to capture content from;
 - (b) based on the web scraping request, generating a web request for the target website;
 - (c) transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies;
 - (d) in response to the web request, receiving, via the proxy, content transmitted from the target website;
 - (e) counting an amount of data in the received content to determine a current traffic total for a client of the client computing device; and
 - (f) transmitting the received content to the client computing device,
- wherein the web request is a second web request, and the received content is a second content, the operations further comprising:
- (g) selecting a scraper from a plurality of scrapers based on the target website such that the selected scraper includes instructions on how to generate a first web request and the second web request;
 - (h) generating the first web request for the target website according to the instructions;
 - (i) transmitting the first web request such that the web request reaches the target website via the proxy; and

- (j) in response to the first web request, receiving, via the proxy, a first content including a data transmitted from the target website via the proxy,
- wherein the generating (b) comprises generating, based on the data, the second web request according to the instructions in the selected scraper
- wherein the counting (e) comprises excluding an amount of data in the first content to determine the current traffic total retrieved by the scraper for the client.
33. The device of claim 32, the operations further comprising:
- (g) determining whether the target website has refused to serve the web request from the proxy, wherein steps (b)-(f) are conducted when the target website is determined in (g) not to have refused to serve the web request from the proxy.
34. The device of claim 33, the operations further comprising:
- (h) when the target website is determined in (g) to have refused to serve the web request from the proxy, retrying to send the web request to the target website via a different proxy.
35. The device according to any of the claims 32 to 34, in particular according to claim 32, wherein the counting (e) comprises determining the amount of data in the received content as compressed for transmission.
36. The device according to any of the claims 32 to 35, in particular according to claim 32, wherein the counting (e) further comprises:
- (i) determining a type of data represented by the received content;
- (ii) based on the type of data, determining a compression factor representing an amount of compression expected when the type of data is transmitted over a network; and
- (ii) based on the compression factor, determining the amount of data in the received content as compressed for transmission.

37. The device according to any of the claims 32 to 36, in particular according to claim 32, further comprising:
- (g) analyzing the content to determine web addresses for additional content needed to render a web page; and
 - (h) retrieving the additional content from the web addresses,
- wherein the counting (e) comprises including an amount of data in the additional content in the current traffic total for a client of the client computing device.
38. A computer-implemented method for determining which servers are available to process web scraping jobs, comprising:
- repeatedly checking health of each of a plurality of database servers;
 - based on the health checks, determine whether each of a plurality of database servers are to be enabled or disabled in a table, the plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;
 - when a web scraping request is received from a client computing device:
 - selecting one of the database servers identified as enabled in the table; and
 - sending a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job,
 - wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a number of messages queued within the respective database server's message broker.
39. The method of claim 38, wherein each of the repeated checking comprises, for each of the plurality of database servers, connecting to the portion of the job database for the respective database server.

40. The method according to any of the claims 38 and 39, in particular according to claim 38, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.
41. The method according to any of the claims 38 to 40, in particular according to claim 38, wherein each of the plurality of database servers is a shard managing storage in a horizontal partition of the jobs database.
42. The method according to any of the claims 38 to 41, in particular according to claim 38, wherein each of the plurality of database servers do not synchronize states to one another.
43. The method according to any of the claims 38 to 42, in particular according to claim 38, wherein the plurality of database servers are executed by a plurality of different computing devices.
44. The method according to any of the claims 38 to 43, in particular according to claim 38, further comprising:
determining whether a number of database servers that are disabled in the plurality of database servers exceeds a threshold; and
when the number of database servers that are disabled exceeds the threshold, alerting an administrator.
45. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, comprising:
repeatedly checking health each of a plurality of database servers;

based on the health checks, determining whether each of a plurality of database servers are to be enabled or disabled in a table, the plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;

when a web scraping request is received from a client computing device:

selecting one of the database servers identified as enabled in the table; and

sending a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job,

wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a number of messages queued within the respective database server's message broker.

46. The device of claim 45, wherein each of the repeatedly checking comprises, for each of the plurality of database servers, connecting to the portion of the job database for the respective database server.
47. The device according to any one of the claims 45 and 46, in particular according to claim 45, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and each of the repeatedly checking comprises, for each of the plurality of database servers, checking a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.
48. The device according to any one of the claims 45 to 47, in particular according to claim 45, wherein each of the plurality of database servers is a shard managing storage in a horizontal partition of the jobs database.

49. The device according to any one of the claims 45 and 48, in particular according to claim 45, wherein each of the plurality of database servers do not synchronize states to one another.
50. The device according to any one of the claims 45 and 49, in particular according to claim 45, wherein the plurality of database servers are executed by a plurality of different computing devices.
51. The device according to any one of the claims 45 and 50, in particular according to claim 45, further comprising:
- determining whether a number of database servers that are disabled in the plurality of database servers exceeds a threshold; and
 - when the number of database servers that are disabled exceeds the threshold, alerting an administrator.
52. A system for determining which servers are available to process web scraping jobs, comprising:
- a processor;
 - a job database that stores the status of web scraping jobs while the web scraping jobs are being executed;
 - a memory that stores the job database;
 - a plurality of database servers operating independently of one another, each database servers configured to manage data storage to at least a portion of the job database;
 - a database monitor configured to repeatedly check health of each of the plurality of database servers and, based on the results of the health checks, determine whether each of the plurality of database servers are to be enabled or disabled in a table;
 - a database server selector configured to, when a web scraping request is received from a client computing device, select one of the database servers identified as enabled in the table; and

a request intake manager configured to send a job description specified by the web scraping request to the selected database server for storage in the job database as a pending web scraping job,

wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the jobs database, and the database monitor is configured to, for each of the plurality of database servers, check a number of messages queued within the respective database server's message broker.

53. The system of claim 52, wherein the database monitor is configured to, for each of the plurality of database servers, check a connection between the request intake manager and the jobs database.
54. The system according to any one of the claims 52 and 53, in particular according to claim 52, wherein each of the plurality of database servers comprises a message broker that queues job descriptions to be stored in the job database, and the database monitor is configured to, for each of the plurality of database servers, check a connection between a server that receives web scraping requests from client computing devices and the respective database server's message broker.
55. A computer-implemented method for distributing web scraping requests over a plurality of data centers, comprising:
 - (a) receiving an API request from a client computing device, the API request relating to a web scraping job that was previously requested of a web scraping system;
 - (b) identifying a hostname addressing a host at a data center of a plurality of data centers executing the web scraping system, the data center servicing the web scraping job;
 - (c) directing the API request to the identified hostname; and
 - (d) at the identified data center, servicing the API request,wherein the API request is a second API request, and
wherein the method further comprises, at the identified data center:

receiving a first API request to request that the web scraping system scrape content from a target website via a proxy, wherein the first API request is received prior to the second API request;

generating a web request for the target website as specified in the first API request;

transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies; and

in response to the web request, receiving scraped content from the target website via the proxy; and

wherein the data center is at a first data center wherein steps (a)-(c) are executed at a second data center different from the first data center.

56. The method according to claim 55, wherein the second API request is to retrieve state of the web scraping job specified by the API request.
57. The method of claim 56, wherein the second API request is to retrieve the received scraped content.
58. The method according to any one of the claims 55 to 57, in particular according to claim 56, further comprising:
 - sending a message to the client computing device, the message specifying a sub-domain for the client computing device to retrieve the scraped content.
59. The method according to any one of the claims 55 to 58, in particular according to claim 55, wherein the hostname addresses a load balancer that balances requests among a plurality of servers that accept web scraping requests.
60. The method according to any one of the claims 55 to 59, in particular according to claim 55, wherein steps (a)-(c) are executed at a data center balancer that selects which data center of the plurality of data centers to service the web scraping job.

61. The method of claim 60, wherein the directing (c) comprises sending an instruction to the client computing device directing the client computing device to resend the API request to the sub-domain.
62. The method of claim 60, wherein the directing (c) comprises forwarding the API request to the sub-domain for the first data center processing the request, further comprising:
at the second data center, forwarding the scraped content from the first data center to the client computing device.
63. The method according to any one of the claims 55 and 62, in particular according to claim 55, wherein the plurality of data centers executing the web scraping system each operate to accept job requests independently of one another, wherein status and results from the API request are stored on one of a plurality of independent, horizontally partitioned database shards.
64. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:
- (a) receiving an API request from a client computing device, the API request relating to a web scraping job that was previously requested of a web scraping system;
 - (b) identifying a hostname addressing a host at a data center of a plurality of data centers executing the web scraping system, the data center servicing the web scraping job; and
 - (c) directing the API request to the identified hostname for servicing at the identified data center,
- wherein the API request is a second API request,
and the operations comprise, at the identified data center:

receiving a first API request to request that the web scraping system scrape content from a target website via a proxy, wherein the first API request is received prior to the second API request;

generating a web request for the target website as specified in the first API request;

transmitting the web request such that the web request reaches the target website via a proxy selected from a group of proxies; and

in response to the web request, receiving scraped content from the target website via the proxy,

wherein the data center is at a first data center wherein steps (a)-(c) are executed at a second data center different from the first data center.

65. The device according to claim 64, wherein the second API request is to retrieve the received scraped content.
66. The device according to any of the claims 64 to 65, in particular according to claim 64, the operations further comprising:
- sending a message to the client computing device, the message specifying the sub-domain for the client computing device to retrieve the scraped content.
67. The device according to any of the claims 64 to 66, in particular according to claim 64, wherein the plurality of data centers executing the web scraping system each operate to accept job requests independently of one another, wherein status and results from the API request are stored on one of a plurality of independent, horizontally partitioned database shards.
68. The device according to any of the claims 64 to 67, in particular according to claim 64, wherein the hostname addresses a load balancer that balances requests among a plurality of servers that accept web scraping requests.

69. A system for distributing web scraping requests over a plurality of data centers, comprising:
- at least one processor and memory;
 - a data center balancer configured to (i) receive a first API request to scrape content from a target website and a second API request to retrieve the content once scraped, (ii) direct the first API request to a first data center selected from a plurality of data centers, and (iii) direct the second API request to a second data center selected from the plurality of data centers;
 - a first request intake manager executed at the first data center, the first request intake manager configured to (i) in response to the first API request, initiate web scraping from the target website and to (ii) in response to the second API request, return content scraped from the target website to a client computing device; and
 - a second request intake manager executed at the second data center, the second request intake manager configured to in response to the second API request, (i) determine which of the plurality of data centers processed the first API request, and to (ii) direct the second API request to the determined data center.
70. The system of claim, 69 further comprising:
- a first load balancer executed at the first data center, the first load balancer configured to, in response to the first API request: (i) select the first request intake manager from a first plurality of first request intake managers executed at the first data center, and (ii) direct the first API request to the selected first request intake manager; and
 - a second load balancer executed at the second data center, the second load balancer configured to, in response to the second API request: (i) select the second request intake manager from a second plurality of first request intake managers executed at the second data center, and (ii) direct the second API request to the selected second request intake manager;

71. A computer-implemented method for securing a web scraping system, comprising:
- at an entry point to the web scraping system, performing the following:
 - (a) validating credentials received with an API request from a client computing device, the API request asking that the web scraping system scrape content from a target website;
 - (b) when the credentials are validated, generating a token indicating an identity of a client associated with the credentials;
 - (c) transmitting the API request along with the token to a server configured to initiate a scraping process on the web scraping system;
 - at the server configured to initiate the web scraping system:
 - (d) analyzing the token to determine whether the client is authorized to conduct the request; and
 - (e) when the client is authorized, causing the web scraping system to scrape the target website,
- wherein the entry point is a load balancer that selects the server configured to initiate the web scraping system from a plurality of parallel servers.
72. The method of claim 71, further comprising:
- (f) passing the API request between a plurality of servers, each configured to perform a function of the web scraping system, the server configured to initiate the web scraping system being included in the plurality of servers;
 - at each of the respective servers:
 - (g) analyzing the token to determine whether the client is authorized to conduct the function performed by the respective server; and
 - (e) when the client is authorized to conduct the function, performing the function.
73. The method of claim 72, wherein the plurality of servers includes a server configured to service API requests formatted as a web proxy request.

74. The method according to any one of the claims 72 and 73, in particular according to claim 72, wherein the plurality of servers includes a server configured to service synchronous API requests, leaving a connection between the web scraping system and the client computing device open while the web scraping system scrapes the target website.
75. The method according to any one of the claims 72 to 74, in particular according to claim 72, wherein the plurality of servers includes a server configured to service asynchronous API requests, closing a connection between the web scraping system and the client computing device before the web scraping system scrapes the target website.
76. The method according to any one of the claims 72 to 75, in particular according to claim 72, wherein the generating (b) comprises generating the token to include a role of the client.
77. The method according to any one of the claims 72 to 76, in particular according to claim 72, wherein the generating (b) comprises generating the token to include a digital signature that cryptographically guarantees that the identity of the client has not been tampered with.
78. The method according to any one of the claims 72 to 77, in particular according to claim 71, wherein the API request is a first API request, and the token is a first token, further comprising:
- (f) validating credentials received with a second API request, the second API request asking to retrieve content that the web scraping system has previously scraped from the target website;
 - (g) when the credentials are validated, generating a second token indicating an identity of a client associated with the credentials received with the second API request;
 - (h) determining whether the first and second tokens indicate that the first and second API requests came from the client; and

- (i) when the first and second tokens indicate that the first and second API requests came from the client, returning the scraped content in response to the second API request.
- 79. The method of claim 78, further comprising, when the first and second tokens do not indicate that the first and second requests came from the client, refusing to return the scraped content.
- 80. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:
 - at an entry point to a web scraping system, performing the following:
 - (a) validating credentials received with an API request from a client computing device, the API request asking that the web scraping system scrape a target website;
 - (b) when the credentials are validated, generating a token indicating an identity of a client associated with the credentials;
 - (c) transmitting the API request along with the token to a server configured to initiate a scraping process on the web scraping system;
 - at the server configured to initiate the web scraping system:
 - (d) analyzing the token to determine whether the client is authorized to conduct the request; and
 - (e) when the client is authorized, causing the web scraping system to scrape the target website,
 - wherein the entry point is a load balancer that selects the server configured to initiate the web scraping system from a plurality of parallel servers.
- 81. The device of claim 80, the operations further comprising:

- (f) passing the API request between a plurality of servers, each configured to perform a function of the web scraping system, the server configured to initiate the web scraping system being included in the plurality of servers;
 - at each of the respective servers:
 - (g) analyzing the token to determine whether the client is authorized to conduct the function performed by the respective server; and
 - (e) when the client is authorized to conduct the function, performing the function.
- 82. The device of claim 81, wherein the plurality of servers includes a server configured to service API requests formatted as a web proxy request.
- 83. The device according to any one of the claims 81 and 82, in particular according to claim 81, wherein the plurality of servers includes a server configured to service synchronous API requests, leaving a connection between the web scraping system and the client computing device open while the web scraping system scrapes the target website.
- 84. The device according to any one of the claims 81 to 83, in particular according to claim 81, wherein the plurality of servers includes a server configured to service asynchronous API requests, closing a connection between the web scraping system and the client computing device before the web scraping system scrapes the target website.
- 85. The device according to any one of the claims 81 to 84, in particular according to claim 81, wherein the generating (b) comprises generating the token to include a role of the client.
- 86. The device according to any one of the claims 81 and 85, in particular according to claim 81, wherein the generating (b) comprises generating the token to include a digital signature that cryptographically guarantees that the identity of the client has not been tampered with.

87. The device according to any one of the claims 81 and 86, in particular according to claim 81, wherein the API request is a first API request, and the token is a first token, the operations further comprising:
- (f) validating credentials received with a second API request, the second API request asking to retrieve content that the web scraping system has previously scraped from the target website;
 - (g) when the credentials are validated, generating a second token indicating an identity of a client associated with the credentials received with the second API request;
 - (h) determining whether the first and second tokens indicate that the first and second requests came from the client; and
 - (e) when the first and second tokens indicate that the first and second requests came from the client, returning the scraped content in response to the second API request.
88. The device of claim 87, the operations further comprising, when the first and second tokens do not indicate that the first and second requests came from the client, refusing to return the scraped content.
89. A method for aggregating web results scraped from a plurality of data centers, comprising:
- at a first data center of the plurality of data centers, comprising:
 - (a) generating a first web request for a target website;
 - (b) transmitting the first web request to the target website;
 - (c) in response to the first web request, receiving a first content from the target website;
 - (d) storing the first content in an archive;
 - (e) storing a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;
 - (f) analyzing a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;

at a second data center of the plurality of data centers:

- (g) generating a second web request for the target website;
- (h) transmitting the second web request to the target website;
- (i) in response to the second web request, receiving a second content from the target website;
- (j) storing the second content in the archive;
- (k) storing a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;
- (l) analyzing the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and
- (m) in response to the analyzing (f) and (l), providing to the client a file with the first and second links.

90. The method of claim 89, wherein the providing (m) comprises:
- storing the file in a database separate from the archive; and
 - providing the client read access to the database.
91. The method according to any one of the claims 89 and 90, in particular according to claim 89, wherein the file includes the metadata of the first and second entries.
92. The method according to any one of the claims 89 to 91, in particular according to claim 89, further comprising:
- (n) parsing a response from the target website to the first web request to determine the first content; and
 - (o) parsing a response from the target website to the second web request to determine the second content.
93. The method according to any one of the claims 89 and 92, in particular according to claim 89, further comprising:

- (n) determining whether the steps (a)-(e) have occurred during a recent time period, wherein steps (j)-(k) occur only when the steps (a)-(e) have not occurred during the recent time period.
94. The method of claim 93, wherein the determining (n) comprises, at the second data center:
- (i) receiving, from the first data center, a data feed with additions to the first index file;
 - (ii) storing the additions in a prior scraped file; and
 - (ii) comparing the metadata describing the second content to the prior scraped file.
95. The method of claim 94, further comprising removing metadata from the prior scraped file at an expiration date selected based on the recent time period.
96. A non-transitory computer-readable device having instructions stored thereon that, when executed by at least one computing device, cause the at least one computing device to perform operations, the operations comprising:
- at a first data center of the plurality of data centers, comprising:
 - (a) generating a first web request for a target website;
 - (b) transmitting the first web request to the target website;
 - (c) in response to the first web request, receiving a first content from the target website;
 - (d) storing the first content in an archive;
 - (e) storing a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;
 - (f) analyzing a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;
 - at a second data center of the plurality of data centers:

- (g) generating a second web request for the target website;
 - (h) transmitting the second web request to the target website;
 - (i) in response to the second web request, receiving a second content from the target website;
 - (j) storing the second content in the archive;
 - (k) storing a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;
 - (l) analyzing the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and
 - (m) in response to the analyzing (f) and (l), providing to the client a file with the first and second links.
97. The device of claim 96, wherein the providing (m) comprises:
- storing the file in a database separate from the archive; and
 - providing the client read access to the database.
98. The device according to any one of the claims 96 and 97, in particular according to claim 96, wherein the file includes the metadata of the first and second entries.
99. The device according to any one of the claims 96 to 98, in particular according to claim 96, the operations further comprising:
- (n) parsing a response from the target website to the first web request to determine the first content; and
 - (o) parsing a response from the target website to the second web request to determine the second content.
100. The device according to any one of the claims 96 to 99, in particular according to claim 96, the operations further comprising:

(n) determining whether the steps (a)-(e) have occurred during a recent time period, wherein steps (j)-(k) occur only when the steps (a)-(e) have not occurred during the recent time period.

101. The device of claim 100, wherein determining (n) comprises, at the second data center:

- (i) receiving, from the first data center, a data feed with additions to the first index file;
- (ii) storing the additions in a prior scraped file; and
- (ii) comparing the metadata describing the second content to the prior scraped file.

102. The device of claim 101, the operations further comprising removing metadata from the prior scraped file at an expiration date selected based on the recent time period.

103. A system for aggregating web results scraped from a plurality of data centers, comprising:

- a processor;
- at least one memory;
- an archive database stored in the at least one memory;
- a first index file stored in at least one memory at a first data center;
- a first scraper at the first data center, the first scraper configured to (i) generate a first web request for a target website, (ii) transmit the first web request to the target website, and (iii) in response to the first web request, receive a first content from the target website;
- a first writer at the first data center, the first writer configured to (i) store the first content in an archive and (ii) store a first entry in a first index file, the first entry including metadata describing the first content and a first link addressing the first content in the archive;

a first dumper at the first data center, the first dumper configured to analyze a search request from a client to determine that search criteria in the search request describes the metadata in the first entry;

a second index file stored in the at least one memory at a second data center;

a second scraper at the second data center, the second scraper configured to (i) generate a second web request for the target website, (ii) transmit the second web request to the target website, and (iii) in response to the second web request, receiving a second content from the target website;

a second writer at the second data center, the second writer configured to (i) store the second content in an archive and (ii) store a second entry in a second index file, the second entry including metadata describing the second content and a second link addressing the second content in the archive;

a second dumper at the second data center, the second dumper configured to analyze the search request from the client to determine that the search criteria in the search request describes the metadata in the second entry; and

a dump database configured to provide to the client a file with the first and second links.

104. The system of claim 103, wherein the file includes the metadata of the first and second entries.

105. The system according to any one of the claims 103 and 104, in particular according to claim 103, the operations further comprising:

a first parser at the first data center, the first parser configured to parse a response from the target website to the first web request to determine the first content; and

a second parser at the second data center, the second parser configured to parse a response from the target website to the second web request to determine the second content.

106. The system according to any one of the claims 103 to 105, in particular according to claim 103, wherein the second writer is configured to determine whether the first writer has recorded the target website in the archive during a recent time period, wherein the second writer is configured to record the target website in the archive only when the first writer has not during the recent time period.
107. The system according to any one of the claims 103 to 106, in particular according to claim 103, wherein further comprising, at the second data center:
- a prior scraped file stored in the at least one memory;
 - a feed receiver configured to receive, from the first data center, a data feed with additions to the first index file and to storing the additions in the prior scraped file, wherein the writer is configured to compare the metadata describing the second content to the prior scraped file.
108. The system according to any one of the claims 103 to 107, in particular according to claim 103, wherein metadata is removed from the prior scraped file at an expiration date selected based on the recent time period.
109. A method for scraping web results via a proxy, comprising:
- (a) receiving an API request from a client computing device, the API request specifying a target website to scrape;
 - (b) based on the target website, selecting a scraper specifying how to generate a sequence of HTTP requests to access content requested by the API request;
 - (c) based on the selected scraper and received API request, generating an earlier HTTP request in the sequence of HTTP requests;
 - (d) transmitting, to the target website and via a proxy selected from a group of proxies, the earlier HTTP request;
 - (e) receiving, from the target website and via the proxy selected from the group of proxies, a response to the earlier HTTP request, the response including a session data;

- (f) based on the selected scraper and received API request, generating a later HTTP request in the sequence of HTTP requests such that the later HTTP request includes the session data;
- (g) receiving, from the target website and via the proxy selected from the group of proxies, from the a response to the later HTTP request, the response including the scraped web results; and
- (h) transmitting to the client computing device the scraped web results.

110. A method as in claim 109, further comprising:
according to the API request and the selected scraper, generating a path to use in the earlier HTTP request.
111. A method as in any of the claims 109 to 110, wherein the API request specifies a geographic location, further comprising:
selecting the proxy based on the group of proxies such that an IP address of the selected proxy is located in the geographic location.
112. A method as in any of the claims 109 to 111, wherein the API request specifies a geographic location, further comprising:
determining at least one header value to use in the earlier and later HTTP requests such that the at least one header identifies to the target website that the earlier and later HTTP requests originated from the geographic location.
113. A method as in any of the claims 109 to 112, further comprising:
determining at least one header value to use in the earlier and later HTTP requests to mimic what a browser at the proxy would generate.
114. A method as in any of the claims 109 to 113, further comprising:
determining whether the target website as requested in the API request has been recently scraped from the target website and the recent scraped results stored in a cache;

when the recent scraped results are determined not to be stored in the cache, executing steps (b)-(h); and

when the recent scraped results are determined to be stored in the cache, transmitting the recently scraped results as a response to the API request.

115. A method as in any of the claims 109 to 114, wherein the API request is an HTTP request from the client computing device, wherein the transmitting (h) comprises transmitting the scraped web results as an HTTP response to the HTTP request from the client computing device.
116. A method as in any of the claims 109 to 115, when the receiving (g) is complete transmitting, to the client computing device, a callback message.
117. A method as in any of the claims 109 to 116, wherein the API request is an HTTP proxy protocol request, and wherein the transmitting (h) comprises transmitting the scraped results as a response to the HTTP proxy protocol request.
118. A method as in any of the claims 109 to 117, further comprising:
retrieving, by a headless browser, additional content that would be needed to render a web page specified by the scraped results.
119. A method as in any of the claims 109 to 118, further comprising:
selecting the proxy from the group of proxies to prevent any proxy from the group of proxies from being identified as a bot.
120. A method as in any of the claims 109 to 119, the API request specifying a user-agent, further comprising:
determining at least one header value to use in the earlier and later HTTP requests to mimic the user agent.

121. A method as in any of the claims 109 to 120, further comprising:
receiving both synchronous and asynchronous API requests to scrape web data;
and
prioritizing processing of the synchronous API requests above processing of the asynchronous API requests.
122. A method as in any of the claims 109 to 121, wherein the response to the earlier HTTP request includes an HTML page with a form to submit information to a URI, wherein the later HTTP request includes a path with the URI.
123. A method as in any of the claims 109 to 122, wherein the API request includes a value, wherein the response to the earlier HTTP request includes an HTML page with a form with a parameter name, wherein the later HTTP request indicates that the parameter name corresponds to the value.
124. A method as in any of the claims 109 to 123, wherein the API request includes a value matched to one of several parameter names, further comprising normalizing the API request to a common format.
125. A method as in any of the claims 109 to 124, further comprising generating a job identifier based on a server servicing the API request.
126. A method as in any of the claims 109 to 125, further comprising storing status of a web scraping job specified in the API request in an in-memory table.
127. A method as in any of the claims 109 to 126, further comprising storing status of a web scraping job in a horizontally partitioned database.
128. A system for scraping web results via a proxy, comprising:

a means for receiving an API request from a client computing device, the API request specifying a target website to scrape;

a means for, based on the target website, selecting a scraper specifying how to generate a sequence of HTTP requests to access content requested by the API request;

a means for, based on the selected scraper and received API request, generating an earlier HTTP request in the sequence of HTTP requests;

a means for transmitting, to the target website and via a proxy selected from a group of proxies, the earlier HTTP request;

a means for receiving, from the target website and via the proxy selected from the group of proxies, a response to the earlier HTTP request, the response including a session data;

a means for, based on the selected scraper and received API request, generating a later HTTP request in the sequence of HTTP requests such that the later HTTP request includes the session data;

a means for receiving, from the target website and via the proxy selected from the group of proxies, from the a response to the later HTTP request, the response including the scraped web results; and

a means for transmitting to the client computing device the scraped web results.

100




FIG. 1




FIG. 2

104




FIG. 3

104




FIG. 4




FIG. 5




FIG. 6




FIG. 7






FIG. 8

**FIG. 9**

1000**FIG. 10**




FIG. 11




FIG. 12A






FIG. 12B

**FIG. 13A****FIG. 13B**




FIG. 14




FIG. 15

1600




FIG. 16

1700




FIG. 17

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2022/067331

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F16/958 G06F16/951

ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2019/174613 A1 (BEIJING DIDI INFINITY TECHNOLOGY & DEV CO LTD [CN]) 19 September 2019 (2019-09-19)	1-20
A	the whole document	21-140
A	US 2014/122682 A1 (FALKENBERG MATTHIAS [DE] ET AL) 1 May 2014 (2014-05-01) the whole document	1-140
X	WO 2021/064480 A1 (METACLUSTER LT UAB [LT]) 8 April 2021 (2021-04-08)	21-40
A	the whole document	1-20, 41-140
A	US 10 965 770 B1 (VILCINSKAS EIVYDAS [LT] ET AL) 30 March 2021 (2021-03-30) the whole document	1-140
	----- -/-	



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

19 October 2022

Date of mailing of the international search report

27/10/2022

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2

NL - 2280 HV Rijswijk

Tel. (+31-70) 340-2040,

Fax: (+31-70) 340-3016

Authorized officer

Ntarlagiannis, V

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2022/067331

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2020/117757 A1 (YANAMANDRA SANGEETHA [IN] ET AL) 16 April 2020 (2020-04-16)	41-60
A	the whole document	1-40, 61-140

X	US 2012/324113 A1 (PRINCE MATTHEW BROWNING [US] ET AL) 20 December 2012 (2012-12-20)	61-100
A	the whole document	1-60, 101-140

X	X B Bai ET AL: "D3.4: Distributed Searching in Future Internet", Information and Communication Technologies (ICT) Programme Project No: FP7--‐ ICT--‐ 248036 COAST, 15 June 2012 (2012-06-15), pages 1-74, XP055157749, Retrieved from the Internet: URL: http://www.coast-fp7.eu/public/COAST_D3.4_YAHOO_BM_FF.pdf [retrieved on 2014-12-10]	101-120
A	the whole document	1-100, 121-140

X	US 2007/050338 A1 (STROHM ALAN C [US] ET AL) 1 March 2007 (2007-03-01)	121-140
A	the whole document	1-120

INTERNATIONAL SEARCH REPORT

International application No.
PCT/EP2022/067331

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims;; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☒ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-20

group I: claims 1-20, directed to an invention focusing on reducing the blocking time of interactive services of the client computing devices

2. claims: 21-40

group II: claims 21-40, directed to an invention focusing on providing traffic monitoring

3. claims: 41-60

group III: claims 41-60, directed to an invention focusing on providing resilient server selection

4. claims: 61-80

group IV: claims 61-80, directed to an invention focusing on an alternative solution for balancing web scraping requests

5. claims: 81-100

group V: claims 81-100, directed to an invention focusing on providing safe crawling

6. claims: 101-120

group VI: claims 101-120, directed to an invention focusing on providing optimised usage of storage resources

7. claims: 121-140

group VII: claims 121-140, directed to an invention focusing on providing an alternative solution for performing multi-resource crawling

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2022/067331

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2019174613 A1	19-09-2019	CN 110309389 A JP 2021515950 A US 2020410031 A1 WO 2019174613 A1	08-10-2019 24-06-2021 31-12-2020 19-09-2019
US 2014122682 A1	01-05-2014	GB 2507294 A US 2014122682 A1 US 2017118278 A1	30-04-2014 01-05-2014 27-04-2017
WO 2021064480 A1	08-04-2021	CA 3139804 A1 CN 114097207 A EP 3900240 A1 EP 4047905 A1 LT 3900240 T SG 11202112754Q A US 10601948 B1 US 10637956 B1 US 10798209 B1 WO 2021064480 A1	08-04-2021 25-02-2022 27-10-2021 24-08-2022 27-06-2022 30-12-2021 24-03-2020 28-04-2020 06-10-2020 08-04-2021
US 10965770 B1	30-03-2021	DE 202021103600 U1 EP 3994870 A1 US 10965770 B1 US 11140235 B1 US 2022086248 A1 US 2022247829 A1 WO 2022053880 A1	20-08-2021 11-05-2022 30-03-2021 05-10-2021 17-03-2022 04-08-2022 17-03-2022
US 2020117757 A1	16-04-2020	NONE	
US 2012324113 A1	20-12-2012	US 2012324113 A1 US 2016156588 A1 US 2018131668 A1 US 2020228490 A1	20-12-2012 02-06-2016 10-05-2018 16-07-2020
US 2007050338 A1	01-03-2007	BR PI0616595 A2 CA 2621031 A1 CN 101971172 A JP 5015935 B2 JP 5474038 B2 JP 2009510547 A JP 2012069163 A KR 20080043865 A US 2007050338 A1 US 2010125564 A1 US 2014046925 A1	23-08-2016 08-03-2007 09-02-2011 05-09-2012 16-04-2014 12-03-2009 05-04-2012 19-05-2008 01-03-2007 20-05-2010 13-02-2014