

Location-Independent Naming for Virtual Distributed Software Repositories*

Shirley Browne[†], Jack Dongarra, Stan Green, Keith Moore
Theresa Pepin, Tom Rowan, and Reed Wade
University of Tennessee
Eric Grosse
AT&T Bell Laboratories

Abstract

A location-independent naming system for network resources has been designed to facilitate organization and description of software components accessible through a virtual distributed repository. This naming system enables easy and efficient searching and retrieval, and it addresses many of the consistency, authenticity, and integrity issues involved with distributed software repositories by providing mechanisms for grouping resources and for authenticity and integrity checking. This paper details the design of the naming system, describes a prototype implementation of some of the capabilities, and describes how the system fits into the development of the National HPCC Software Exchange, a virtual software repository that has the goal of providing access to reusable software components for high-performance computing.

1 Introduction

Well-maintained software repositories are central to software reuse because they make high-quality software widely available and easily accessible. One such repository is Netlib¹, a collection of high-quality publicly available mathematical software [6, 4]. Netlib, in operation since 1985, currently processes over 300,000 requests a day. Netlib is serving as a prototype for development of the National HPCC Software Exchange (NHSE)², which has the goal of encompassing all High Performance Computing Consortium (HPCC) software repositories and of promoting reuse of software components developed by Grand Challenge and other scientific computing researchers [5]. Other network-accessible software

*The work described in this paper is sponsored by NASA under Grant No. NAG 5-2736, by the National Science Foundation under Grant No. ASC-9103853, and by AT&T Bell Laboratories.

[†]Author to whom correspondence should be directed. 107 Ayres Hall, Computer Science Department, University of Tennessee, Knoxville, TN 37996-1301, (615) 974-5886, browne@cs.utk.edu

¹Accessible from a World Wide Web browser at <http://www.netlib.org/>

²Accessible at <http://www.netlib.org/nse/>

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association of Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

SSR '95, Seattle, WA, USA

© 1995 ACM 0-89791-739-1/95/0004...\$3.50

repositories include ASSET³, CARDS⁴, DSRS⁵, ELSA⁶, the GAMS Virtual Software Repository⁷, and STARS⁸. ASSET, CARDS, DSRS, and ELSA are participating in an interoperability experiment that allows a user of any one of these repositories to access software exported from the other repositories.

The software reuse marketplace is expanding in at least two dimensions. One dimension is the expansion from intra-organizational reuse to inter-organizational reuse. For example, various federal agencies have established their own internal software reuse programs. Several efforts are now underway to promote reuse of software across agencies. Similarly, companies are becoming interested in accessing software produced by academic and government research groups. Another dimension of expansion is from reuse within a particular application domain to interdisciplinary reuse. Reuse of software from other disciplines is being fostered, for example, by efforts to solve interdisciplinary Grand Challenge problems. Solution of such problems will require collaboration by scientists from different disciplines, as well as sharing of software produced by application and computer scientists.

Another recent development that affects the software reuse marketplace is the growth of the World Wide Web (WWW), together with the ease with which individuals may make resources available on a WWW server. A contributor need only make the files composing an resource available on a file server and make available a descriptive HTML file containing pointers to the resource files.

Growth in the popularity of the Internet and the World Wide Web, as well as the wide availability of WWW client and server software, has accelerated the shift from centrally maintained software repositories to virtual, distributed repositories. For example, the GAMS Repository, once a central repository, is now a virtual repository that catalogs software maintained by other repositories [2]. Similarly, the NHSE will provide a uniform interface to a virtual HPCC software repository that will be built on top of a distributed set of discipline-oriented repositories [5], as shown in Figure 1.

The main advantage of distributing a repository is to

³ Accessible at <http://source.asset.com/>

⁴ Accessible at <http://dealer.cards.com/>

⁵ Accessible at <http://ssed1.ims.disa.mil/srp/dsrspage.html>

⁶ Accessible at

http://rbse.mountain.net/ELSA/elsa_lob.html

⁷ Accessible at <http://gams.nist.gov/>

⁸ Accessible at

<http://www.stars.ballston.paramax.com/index.html>

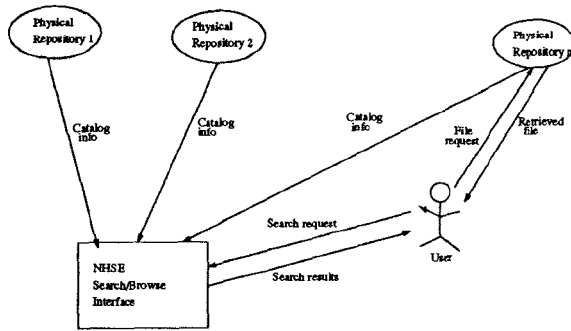


Figure 1: Virtual Repository Architecture

allow the software to be maintained by those in the best position to keep it up-to-date. Also, copies of popular software packages may be mirrored by a number of sites to increase availability (e.g., if one site is unreachable, the software may be retrieved from a different site) and to prevent bottlenecks.

Despite the benefits, distributed maintenance and mirroring of software poses the following challenges.

- Maintaining the quality of software and of indexing information and presenting a uniform searching and browsing interface become much more difficult.
- The WWW mechanism of specifying a file by its Uniform Resource Locator (URL) is inadequate for ensuring the consistency and currency of mirrored copies, as a URL for an independently mirrored copy of a software package may point to an out-of-date copy and give no indication that it is not up-to-date. Furthermore, mirror copies of a file cannot be located from a URL reference, since each copy has a different URL.
- Consistency between a set of files that are meant to be used together must be maintained. For example, the Netlib Software Repository provides dependency checking that allows the user to retrieve a top-level routine plus all routines in its dependency tree (i.e., those routines that are called directly or indirectly by the top-level routine). Another example is a graphical parallel programming environment that relies on an underlying parallel communications support package. The problem becomes more complex when different pieces might be retrieved from different physical repositories. Ideally, the user should be able to have a consistent set retrieved automatically without having to scan documentation to verify that compatible pieces have been retrieved.
- As the number of reuse libraries grows, users cannot be expected to access each of them separately using a different interface. Thus, scalable interoperability between separately managed repositories is needed.
- In the environment of accessing a few well-established repositories that the user knows and trusts, a user is assured of the integrity and authenticity of a retrieved file because these properties are provided by the administrative procedures of that repository. With a large number of less familiar repositories, however, it becomes necessary to establish interoperable trust mechanisms and to reduce the number of parties with whom the user must establish trust.

- The more decentralized and smaller the individual repositories become, the less practical it becomes for each individual repository to provide the full range of search and authentication services.

Most of the above problems can be alleviated by implementing a location-independent naming system that includes mechanisms for authenticity and integrity checking. We have designed a naming system that provides for two levels of naming. The binding between a lower-level name (called a LIFN) and file contents is unchangeable and verifiable. A lower-level name may be resolved to multiple, mirrored copies. In the case where it represents a set of files, the name may be resolved to a list of other names. A higher-level name (called a URN) is associated with a cataloging record that includes the lower-level name as well as other descriptive information. This record may be cryptographically signed by the publisher so that users may verify the authenticity of a retrieved resource. At any given time, a higher-level name is associated with exactly one lower-level name, but this binding may change over time. Higher-level names allow for long-lived human-readable references, while lower-level names permit reliable caching and mirroring as well as permitting precise references when needed. Location-independent names will be the basis of transparent mirroring. They will also provide a unique key to which third parties may attach value-added information such as additional cataloging information and quality assessments. This paper describes the design of our naming system. We also describe our implementation of a prototype name-to-location service and of a modified WWW client that does name resolution. A glossary of acronyms and terms used in this paper is included as an appendix.

2 Related Work

The use of a public-key encryption technique for authenticating the source of a software component and for ensuring that the component has not been altered subsequent to its publication is proposed in [9]. Cryptographic information, in the form of a digital signature created by signing the hashed digest of the contents of a component, is included within the component's unique identifier. The proposed method is intended to prevent not only changes by unauthorized parties, but also changes by the original author – i.e., the author is not permitted to modify a component without assigning a new unique identifier. The method assumes that each author has been assigned a globally unique Author ID, has chosen an asymmetric public/private key pair, and has publicized the public key to the community of potential re-users. A newly chosen symmetric encryption key is used to encrypt the component itself. Then the symmetric key, the hashed digest of the component, and the Author ID are concatenated and encrypted using the asymmetric private key, and the result is concatenated to the cleartext version of the Author ID to create the unique identifier for the component. The method does not address name-to-location resolution, other than to say that the encrypted component is made available along with the unique identifier and any other cleartext information. The proposed unique identifier is similar to our LIFN, and encryption of the hash digest and Author ID is similar to our method of having the author cryptographically sign a catalogue record that includes the author name and the file's MD5 signature. Our method allows a choice of encryption algorithms, however, and allows the digital signature used for authentication to

be generated independently and at a different time from the component's identifier.

Functional requirements for Uniform Resource Names (URNs) are proposed in [12] by the IETF Uniform Resource Identification (URI) Working Group. According to [12], the function of a URN is to provide a globally unique, persistent identifier used for recognition of and for access to characteristics of a resource or to the resource itself. URN assignment is delegated to naming authorities, the names of which are persistent and globally unique, and who may assign names directly or delegate their authority to sub-authorities. Global uniqueness of URNs is guaranteed by requiring each naming authority to guarantee uniqueness within its portion of the URN namespace. It is left up to each naming authority to determine the conditions under which it will issue a URN (for example, whether or not to issue a new URN when the contents of a file change). Some test implementations of URNs are underway by members of the URI Working Group at Georgia Tech and Bunyip Corporation⁹. The Georgia Tech testbed uses the whois++ protocol for URN to URC resolution. A URC, or Uniform Resource Characteristic, is a catalog record which includes locations, or URLs, at which the resource may be accessed. The URC server supports searching by other attributes, in addition to URN lookup, via the whois++ protocol. A modified version of Mosaic that does URN to URC resolution is available. A proxy server based on CERN httpd that does caching by URNs is also running at Georgia Tech.

As part of the Computer Science Technical Report (CSTR) project [8], which is developing an architecture for distributed digital document libraries, the Corporation for National Research Initiatives (CNRI) is implementing a name-to-location resolution service called the Handle Management System (HMS)¹⁰. CNRI's *handle* is a name for a *digital object* and is analogous to IETF's URN. The HMS includes a Handle Generator that a naming authority may run and use to create globally unique handles, Handle Servers that process update requests from naming authorities and query requests from clients to resolve handles, and a Handle Server Directory that maps a handle to the appropriate Handle Server. The distribution of handles to Handle Servers is based on a hashing algorithm. An electronic mail interface is used by handle administrators to add, delete, and modify handle entries in the Handle Server database. Clients use a UDP datagram interface to request location data associated with a handle. A modified version of Mosaic that does handle resolution is available from CNRI. The types of location information stored by Handle Servers include URL, repository name, email address, and X.500 Distinguished Name. Use of a repository name by a client requires another round of name-to-location resolution. CNRI's *properties record* that describes the properties of a digital object is analogous to IETF's URC. The properties record is not stored by the HMS, but rather by an Information and Reference (IR) Server that is to be maintained by each repository. Each naming authority may also maintain an IR server containing a properties record for each digital object within its authority.

3 Publishing and Name Assignment

Internet-accessible resources are currently referenced using Uniform Resource Locators (URLs). Because URLs are lo-

⁹ More information is available at <http://www.gatech.edu/iir/>

¹⁰ More information is available at <http://www.cnri.reston.va.us/>

cations rather than names, their use as references presents at least two problems. One problem is that files get moved, changing their URLs. Then pointers that contain the old URLs become stale. One can leave a forwarding address at the old URL, but forwarding addresses are an awkward and inelegant solution. Another problem with using URLs as references is that mirrored copies of files cannot be located from a URL reference, since each copy has a different URL.

It has been widely recognized that a solution to the above problems is to assign location-independent names to files and to provide a name-to-location service that, given a name, returns a list of locations for that name. A resource provider who moves some files need only delete the old name-to-location bindings and register the new bindings with the name-to-location service. Likewise, a site that mirrors a copy of a file need only register its location with the name-to-location service. Then a user attempting to retrieve the file corresponding to a location-independent name may query the name-to-location service for a list of alternative locations to be tried.

Our work is similar to the IETF's Uniform Resource Identifier Working Group's work on Uniform Resource Names (URNs) [12] and to CNRI's work on unique document identifiers for digital libraries [8]. However, neither of these groups has addressed the reliability and consistency issues addressed by our two-level naming system. Our system includes a lower-level name called Location Independent File Name (LIFN) and a higher-level name called a Uniform Resource Name (URN).

An important question is whether the byte contents of the file referred to by a location-independent name should be fixed or be allowed to change. If the byte contents are allowed to change, then a further question arises as to what should be the consistency requirements for alternative locations for the same name. Valid arguments for both cases can be made for different situations. For example, for software resources it is desirable to have an unambiguous reference to the fixed byte contents for the purpose of attaching a review or reporting experimental or performance results. Fixed contents also make it possible to compute a file digest that may be cryptographically signed by the author of the resource, allowing verification of the integrity of a retrieved file. On the other hand, it is desirable to have a reference to a software package that need not be changed every time a bug fix or minor revision takes place, especially if the cataloging information (e.g., title, author, abstract) does not change. The cataloging information for a software package might contain a reference to a Web page describing and/or documenting the package. The author of the Web page would like to be able to update the page without having to change all the references to it. A non-software example where it would be desirable to allow contents to change is a name that refers to a file containing the "current weather map".

Because both types of name are needed, we have implemented both. The type of name that refers to fixed byte contents is called a Location Independent File Name, or LIFN. Once a LIFN has been assigned to a particular sequence of bytes, that binding may not be changed. The type of name for which the contents to which it refers may change is called a Uniform Resource Name, or URN.

We divide the file access system into two levels. The upper level is where publishing, cataloging, and searching activities take place. These upper-level activities are concerned with the semantic, or intellectual, contents of files. The lower level is where distribution, mirroring, and caching

activities occur. These lower-level activities are not concerned with the semantic contents of files, only with ensuring that files may be accessed efficiently and that the byte contents of files are not corrupted.

The above arguments about the need for two types of name pertain to the upper level. At the lower level, there is a need for LIFNs, but not for URNs. Mirror sites use LIFNs and their associated file digests to ensure that their copies of files have not been corrupted. A cache site needs to be able to tell a user or client program whether it holds a copy of a requested file, and for this purpose it can answer whether or not it holds a copy of a particular LIFN.

The above considerations led us to implement LIFNs at the lower level of the file access system and URNs at the upper level, but to make LIFNs visible at the upper level as well. A publisher will be responsible for assigning both a URN and a LIFN to any resource for which cataloging information is provided. For other files, only LIFNs need be provided. At any given time, a URN that refers to a file or a set of files is associated with exactly one LIFN. A URN may be associated with a set of different LIFNs over the URN's lifetime, but we require that the set be in the form of a linear sequence, with the sequence order given by increasing time.

The LIFN and URN name spaces are subdivided among several *publishers*, also called *naming authorities*, who are responsible for ensuring the uniqueness of names assigned within their portions of the name spaces. A name is formed by concatenating the registered naming authority identifier with a unique string assigned by the naming authority. The LIFN and URN are formatted as

```
LIFN:<publisher id>:string
URN:<publisher id>:string
```

The *publisher id* portion of the name is used to locate appropriate URN and LIFN servers for that publisher. Given a URN, a URN server returns a Uniform Resource Citation (URC) for that URN that includes its currently associated LIFN, as well as other cataloging information. Given a LIFN, a LIFN server returns a list of locations for that LIFN. More information about accessing URCs and files from their URNs and LIFNs may be found in Section 4.

The publisher provides cataloging information for each URN it assigns. The catalog record includes information such as title, author, abstract, etc. A recommended set of attributes for software assets is given by the Reuse Library Interoperability Group (RIG) Basic Interoperability Data Model [1]. In addition, the catalogue record for a URN includes its currently associated LIFN, as well as an MD5 or similar fingerprint for that LIFN. This fingerprint is a 128-bit quantity resulting from applying the MD5 function to the contents of the file. The function is designed to make it computationally infeasible to find a different sequence of bytes that produces the same fingerprint [10]. To enable authentication, the entire description may be cryptographically signed, as discussed in Section 5. Portions of the catalog record may be exported to resource discovery servers, such as a Harvest Broker [3], which provide search services based on resource descriptions. The URN exported to the search service provides a unique long-lived key, so that descriptions may be unambiguously associated with a resource, and so that a resource turns up at most once in a list of search hits.

For a name to be useful, there must be some means of resolving a name to a location from which the resource can be retrieved or accessed. Thus, the publisher, as well as

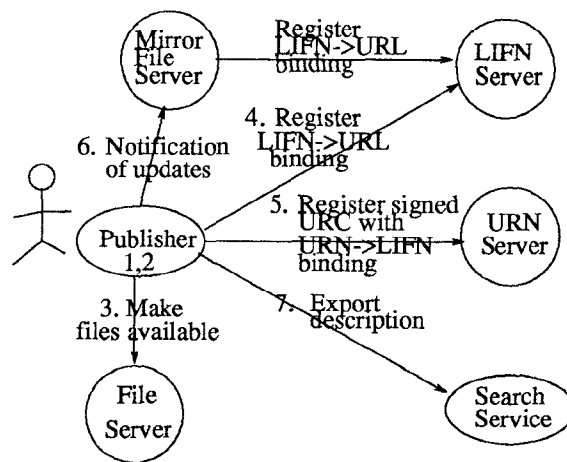


Figure 2: Publishing steps

any other parties that mirror the resource, must register such locations with the appropriate name-to-location lookup services. Such name-to-location services are discussed in Section 4.

Thus, publishing a resource involves the following steps, shown in Figure 2:

1. creating the resource's catalog record in the form of a URC,
2. signing the catalog record with the publisher's private key,
3. making the resource files available on one or more file servers,
4. registering the file locations with the LIFN server,
5. registering the URC with the URN server,
6. informing mirror sites of the new or updated file,
7. exporting relevant portions of the URC to search services.

Steps 1 and 5 have been discussed above. Steps 2 is discussed in Section 5, and Steps 3, 4, and 5 are discussed in Section 4.

4 Name Resolution and File Mirroring

Resources available from the virtual repository will be named by URNs and/or LIFNs, rather than by URLs. Thus, WWW clients will need a means of resolving a URN or LIFN to one or more locations, expressed in the form of a URL, to be able to access the resource. Access to files is provided by conventional file servers, using protocols such as HTTP, Gopher, and FTP.

For a non-file resource, such as a database service, a list of locations is associated directly with the URN for that resource. For a file resource, such as a file containing a piece of software, the relationship between the URN and the locations is indirect, via a LIFN – the URN is associated with a LIFN, and the LIFN is associated with a list of URLs.

The LIFN-to-location mapping service is provided by a network of LIFN servers, collectively called the LIFN database. These servers process queries for locations of

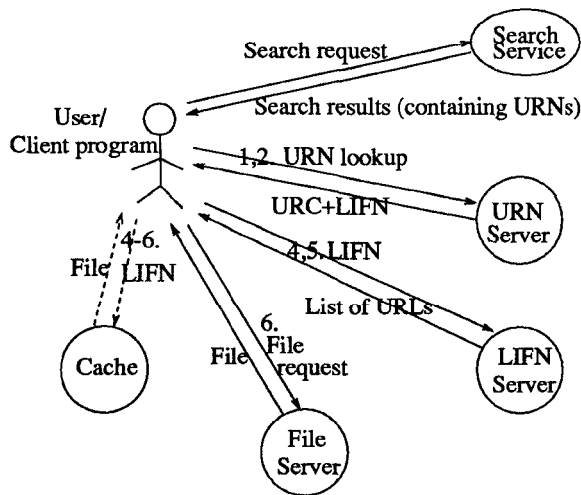


Figure 3: File access steps

LIFNs. They also accept updates from file servers containing new locations for LIFNs, as well as requests to delete old LIFN-to-location mappings. A naming authority may run its own LIFN servers, or it may find another organization willing to provide the service on its behalf.

The URN service is similar to the LIFN service, except that it maps a name either to a list of locations or to a URC that includes a LIFN. For fault tolerance and availability, the URN service is also provided by a network of servers.

Mappings from naming authority identifiers to URN and LIFN servers are stored in the the Domain Name System (DNS) name space, so that a client program can determine which URN (LIFN) server to query for a particular URN (LIFN). Our current client uses an ordinary DNS lookup for IP address records. The publisher identifier is prepended to the string `.LIFN.NETLIB.ORG` (for a LIFN) or `.URN.NETLIB.ORG` (for a URN). The resulting string is treated as if it were the name of an Internet host, and DNS is queried to find the IP addresses of that host. For example, to find a LIFN server for the naming authority `foo`, the client would look up the IP addresses for `foo.LIFN.NETLIB.ORG`. Several IP addresses may be listed for any one naming authority. Our client attempts to query each IP address until it finds one that can satisfy the LIFN or URN lookup request.

Thus, the steps involved in resolving a URN so as to access a copy of the file it names are as follows, as shown in Figure 3:

1. Use DNS to locate an appropriate URN server.
2. Query the URN server to retrieve the URC which contains the currently associated LIFN.
3. Authenticate the URC if desired.
4. Use DNS to locate an appropriate LIFN server.
5. Query the LIFN server to retrieve a list of locations.
6. Choose a location from which to retrieve the file.

In practice, Steps 4 through 6 will often be replaced by using the LIFN to access a local cache server. Because the binding between a LIFN and the byte contents it points to is fixed, the cached copy is sure to be correct.

A file server can mirror a file by acquiring a copy of it and posting an update to a LIFN server for the file's naming authority. If a file server moves or deletes a file, then it would post that information as well. It is not necessary to keep all LIFN servers for a particular naming authority perfectly synchronized. Such synchronization would entail too much overhead. Instead, location updates are posted to a any LIFN server and propagated to other peer servers using a batch update protocol.

Updates to the URN server are posted by the publisher and by others authorized by the publisher to update the catalog record for a given URN. In order to ensure a consistent linear history of updates to the catalog record for a URN (e.g., the sequence of LIFNs associated with that URN), replicated URN servers use a master-slave update protocol.

One of the most important aspects of our use of LIFNs is that it assures the user of retrieving the most up-to-date copy of a file referenced by a URN, without the overhead of a replica control protocol between file servers mirroring that file, which in general will not all be under the control of the URN's naming authority. This assurance is modulo the time required for the master-slave update protocol for the replicated URN servers, but if the user insists on contacting the master URN server, he is ensured of getting the most up-to-date copy.

5 Authenticity, Integrity, and Consistency of Resources

Authentication of a resource verifies that the resource was published by its purported publisher. Verifying the integrity of a file ensures that the file has not been modified. Provisions for authenticity and integrity checking are necessary for a software repository because there have been instances of software packages stored on a public repository that were modified by intruders to introduce security holes which were then spread to other systems¹¹. Our authentication and integrity mechanisms are similar to those described in [11] and [9].

Recall from Section 3 that a publisher cryptographically signs the catalog record for a resource. In the case of a file resource, this record includes the file's LIFN and MD5 fingerprint. Any client in possession of the publisher's public key can verify the authenticity of the resource description. Publishers are expected to widely advertise their public keys to make it difficult for an attacker to substitute rogue keys. In addition, publishers may have their keys certified by trusted third parties to further establish their authenticity, as in [11].

Assuming that the association between a LIFN and a file signature (e.g., the MD5 fingerprint) is known to be correct (either because the signature is part of the LIFN or because of the description authentication described in the preceding paragraph), a client may perform an integrity check on a retrieved file by computing the signature for the file and comparing it with the one known to be associated with the file's purported LIFN. Recall from Section 4 that a LIFN server returns a list of locations for a given LIFN but does not guarantee the correctness of those locations. A location may be incorrect if it no longer exists or if the contents of that location are wrong. In the former case, no file will be returned from that location. The latter condition may be detected by the client performing an integrity check.

To ensure consistency within a group of related files, we allow a URN to refer to a set of files. There are at least two

¹¹For an example, see the CERT advisory at ftp://ftp.cert.org/pub/cert_advisories/CA-94:07.wuarchive.ftpd.trojan.horse

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.