



[54] **DISTRIBUTED INTERFACE ARCHITECTURE FOR PROGRAMMABLE INDUSTRIAL CONTROL SYSTEMS**

[75] Inventors: **Kenneth C. Crater**, North Grafton; **Craig E. Goldman**, Natick, both of Mass.

[73] Assignee: **Control Technology Corporation**, Hopkinton, Mass.

[21] Appl. No.: 655,469

[22] Filed: **May 30, 1996**

[51] Int. Cl.⁶ **G06F 13/14**

[52] U.S. Cl. **364/138**; 364/131; 340/825.07; 395/200.49; 395/200.5; 395/200.58; 345/346

[58] **Field of Search** 364/138, 131, 364/146, 147, 136, 141, 183; 395/701, 200.31, 821, 356, 200.49, 200.5, 200.58; 340/825.07; 345/346

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,319,338	3/1982	Grudowski et al.	364/131
4,937,777	6/1990	Flood et al.	395/287
4,953,074	8/1990	Kametani et al.	364/132
5,012,402	4/1991	Akiyama	364/192
5,072,412	12/1991	Henderson, Jr. et al.	395/159
5,122,948	6/1992	Zapolin	364/131
5,157,595	10/1992	Lovrenich	364/136
5,225,974	7/1993	Mathews et al.	364/140
5,297,257	3/1994	Struger et al.	364/188
5,307,463	4/1994	Hyatt et al.	395/821

5,321,829	6/1994	Zifferer	395/183.22
5,398,336	3/1995	Tantry et al.	395/614
5,406,473	4/1995	Yoshikura et al.	364/140
5,420,977	5/1995	Sztipanovits et al.	395/356
5,598,536	1/1997	Slaughter, III et al.	395/200.16
5,613,115	3/1997	Gihl et al.	395/701
5,623,652	4/1997	Vora et al.	395/610
5,625,781	4/1997	Cline et al.	395/335

Primary Examiner—Hezrone E. Williams

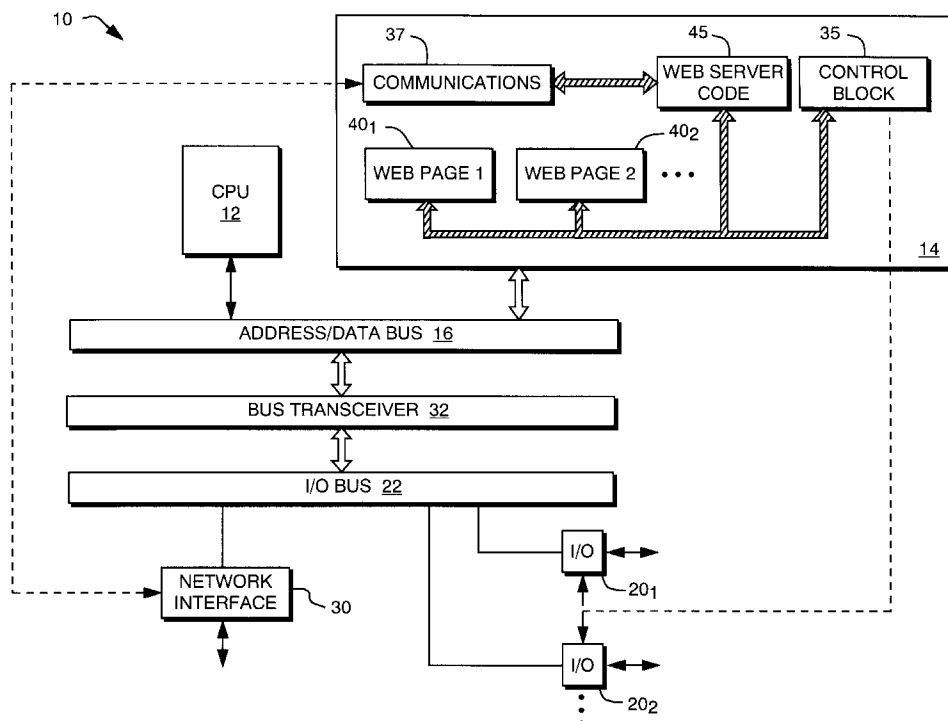
Assistant Examiner—Hien Vo

Attorney, Agent, or Firm—Cesari & McKenna, LLP

[57] **ABSTRACT**

An integrated control system comprises one or more controllers each equipped to perform a control function and to gather data (ordinarily from sensors) relevant to the control function. Each controller contains computer storage means, such as computer memory, for storing the relevant data and instructions, associated with the data, for causing a remote computer to generate a visual display incorporating the data in a predetermined format; and a communication module for establishing contact and facilitating data interchange with the remote computer. The remote computer, in turn, also includes a communication module compatible with the controller-borne module, and which enables the remote computer to download the data and associated instructions from one or more controllers. The remote computer also includes a facility for processing the instructions to create a user interface encoded by the instructions, and which incorporates the data. In this way, controller data is coupled to instructions for displaying that data, and this totality of information is continuously accessible, on a freely selective basis, to the remote computer.

22 Claims, 2 Drawing Sheets



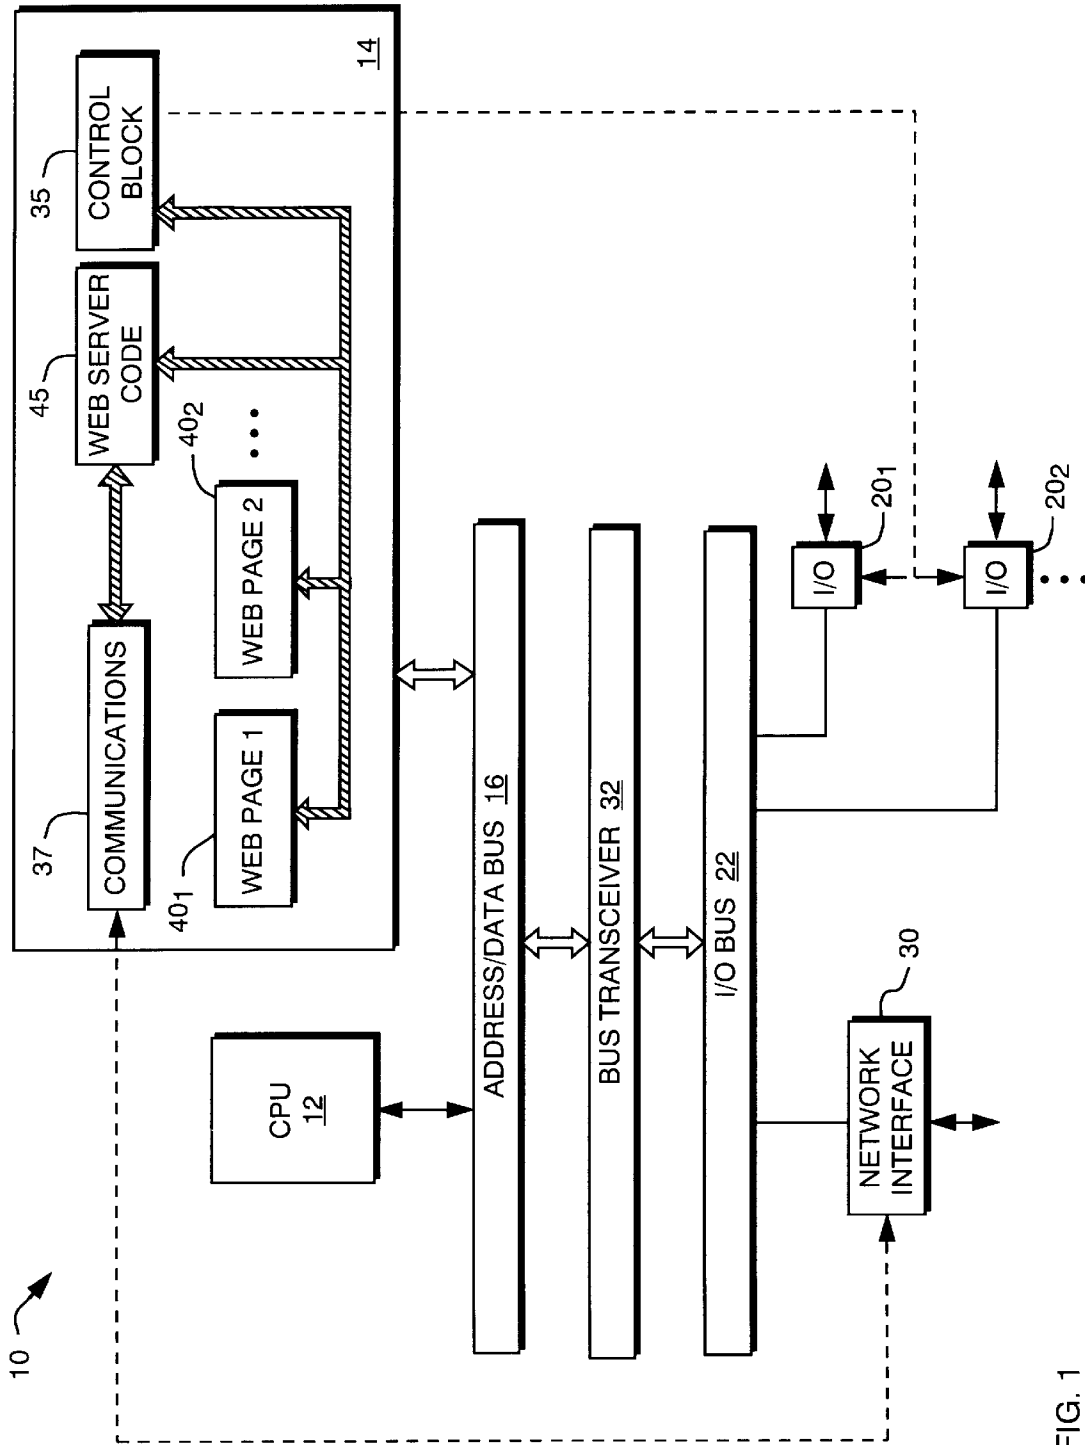


FIG. 1

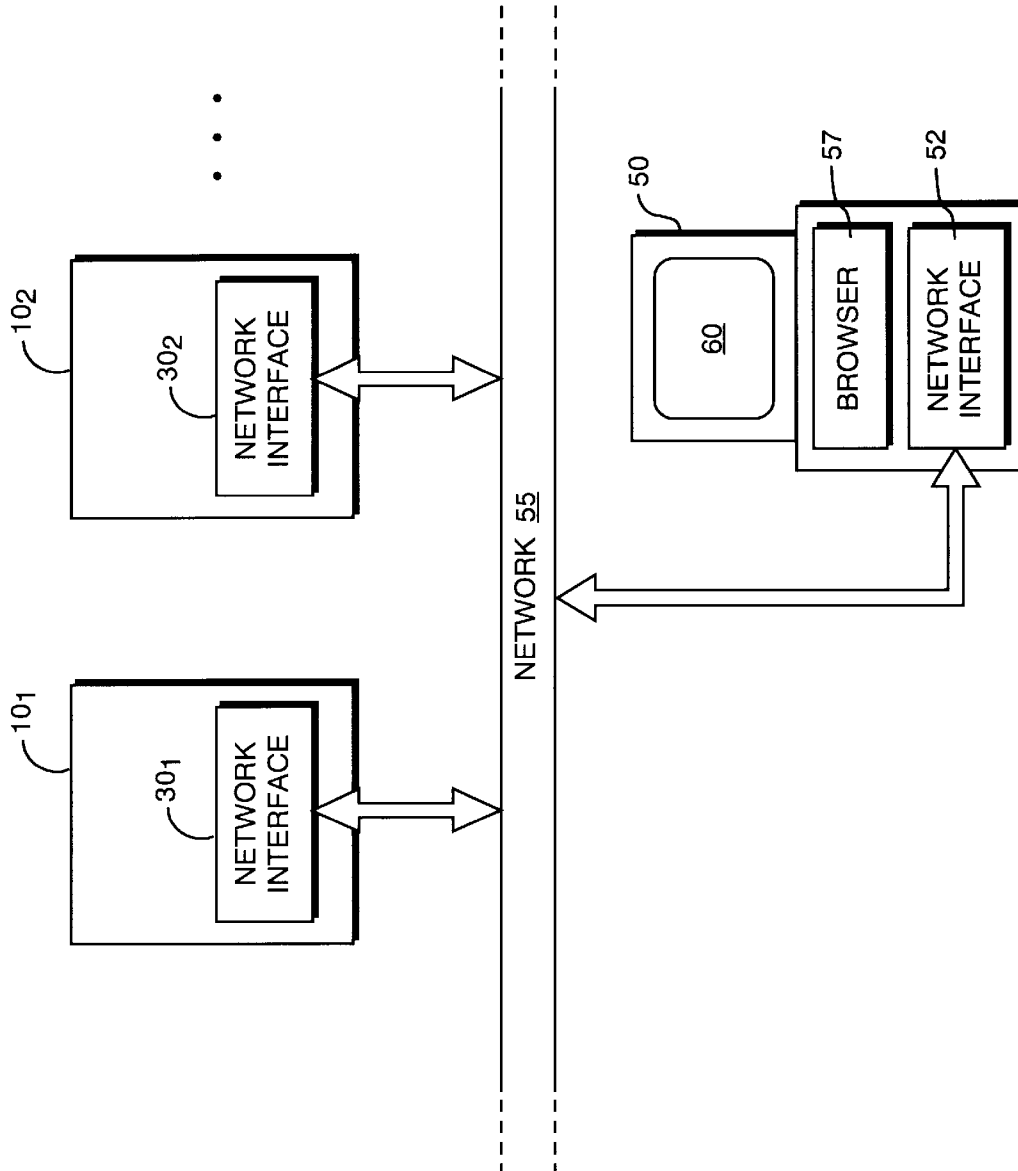


FIG. 2

**DISTRIBUTED INTERFACE
ARCHITECTURE FOR PROGRAMMABLE
INDUSTRIAL CONTROL SYSTEMS**

FIELD OF THE INVENTION

The present invention relates to industrial automation, and in particular to communication with and among programmable controllers for operating and monitoring industrial processes and equipment.

BACKGROUND OF THE INVENTION

Sophisticated industrial processes, such as oil refining, automobile assembly or power generation, require the cooperative execution of numerous interdependent tasks by many different pieces of equipment. The enormous complexity of ensuring proper task sequencing and management, which requires not only procedural logic but constant monitoring of equipment states to organize and distribute operations and detect malfunction, has resulted in the widespread adoption of programmable controllers. These controllers operate elaborate industrial equipment in accordance with a stored control program. When executed, the program causes the controller to examine the state of the controlled machinery by evaluating signals from one or more sensing devices (e.g., temperature or pressure sensors), and to operate the machinery (e.g., by energizing or deenergizing operative components) based on a procedural framework, the sensor signals and, if necessary, more complex processing. The “inputs” to a particular controller can extend beyond the sensed state of the equipment the controller directly operates to include, for example, its environment, the state of related machinery or the state of its controllers.

Control requirements become even more complex when different aspects of the same overall process are assigned to remotely situated equipment. Such configurations often require reliable, high-bandwidth serial communication links to provide the necessary interconnection and handle data transfer among controllers and the sensors relevant to their operation.

Ordinarily, process operation is monitored, at least intermittently, by supervisory personnel by means of one or more central management stations. Each station samples the status of controllers (and their associated sensors) selected by the operator and presents the data in some meaningful format. The management station may or may not be located on the same site as the monitored equipment; frequently, one central station has access to multiple sites (whether or not these perform related processes). Accordingly, communication linkage can be vital even in traditional industrial environments where process equipment is physically proximate, since at least some supervisory personnel may not be.

To facilitate the necessary communication, the controller processors and related computers (such as monitoring stations) are arranged as a computer network. A network, basically, is a collection of interconnected computers that use some consistent protocol to communicate with one another. Typically, the network is organized such that any computer may communicate with any other network computer. The communication protocol provides a mechanism by which messages can be decomposed and routed to a destination computer identified by some form of address. The protocol may place a “header” of routing information on each component of a message that specifies source and destination addresses, and identifies the component to facilitate later reconstruction of the entire message by the destination computer. This approach to data transfer permits the

network to rapidly and efficiently handle large communication volumes without reducing transfer speed in order to accommodate long individual messages, or requiring every network computer to process every network message. The degree of routing depends on the size of the network. Each computer of a local network typically examines the header of every message to detect matches to that computer’s identifier; multiple-network systems use routing information to first direct message components to the proper network.

Controllers have been interconnected by means of computer networks for some time; see, e.g., U.S. Pat. No. 5,307,463. In typical systems, a monitoring computer, which may be remotely located from any or all of the controllers to which it has access, periodically queries the controllers to obtain data descriptive of the controlled process or machine, or the controller itself. This data is then available for analysis by the monitoring computer. Heretofore, however, the type of information obtainable, on demand, from a controller has been limited, while the interface used to present the information on the monitoring computer is typically crude. The latter condition results from the multiplicity of data types offered by the controller. Were each type of data to be rendered in a format suited to that data, it would be necessary to equip the monitoring computer with multiple interfaces, and configure the operating application to ascertain the type of data before selecting and launching the proper interface. Not only does this scheme impose a substantial real-time support burden on the monitoring computer, but would also require a constantly expanding repertoire of graphical capabilities keyed to new forms of data reported by the controllers. This can be especially cumbersome in highly—even internationally—distributed environments, since the controllers and their functionalities may be programmed by personnel having no contact with those responsible for central monitoring.

DESCRIPTION OF THE INVENTION

BRIEF SUMMARY OF THE INVENTION

The present invention utilizes the capabilities of the Internet and, more particularly, the interactive capabilities made available by resources such as the World Wide Web to shift the burden of providing user interfaces for changing forms of data from monitoring computers to the controllers that actually gather and report the data. By combining data with functionality for displaying that data at the individual controller sites, the need to equip monitoring computers with specialized graphic capabilities is eliminated, along with the need for intensive, ongoing cooperation between engineers responsible for programming controllers and those who configure the computers that perform monitoring. Moreover, because Internet users are typically billed for connectivity at a single rate, the long-distance charges that would accrue through use of telephone lines for data communication are eliminated.

In accordance with the invention, an integrated control system comprises one or more controllers each equipped to perform a control function and to gather data (ordinarily from sensors) relevant to the control function. “Relevant” data includes, at a minimum, any information upon which control decisions are made or states shifted, but can also include information obtained from sensors not directly connected to the controller (e.g., involving other controlled machines) but which is nonetheless meaningful to supervisory personnel. For example, a chemical synthesis process may be carried out at a temperature controlled to stay within an operating range, but the optimal temperature may depend

on the output of a previous process feeding into the synthesis; in this case, the temperature of the synthesis process as well as the output of the previous process are relevant control parameters with respect to the synthesis process.

Each controller contains computer storage means, such as computer memory (volatile and/or non-volatile, such as random-access memory (“RAM”), programmable read-only memory (“ROM”) or Flash ROM), or a mass storage device such as a hard disk or CD-ROM, for storing the relevant data and instructions, associated with the data, for causing a remote computer to present the data (e.g., by generating a visual display incorporating the data) in a predetermined format; and a communication module for establishing contact and facilitating data interchange with the remote computer. The remote computer, in turn, also includes a communication module compatible with the controller-borne module, and which enables the remote computer to download the data and associated instructions from one or more controllers. The remote computer also includes a facility for processing the instructions to create a user interface—that is, a visual display or other presentation having a predetermined format—encoded by the instructions, and which incorporates the data. In this way, controller data is coupled to instructions for presenting that data, and this totality of information is continuously accessible, on a freely selective basis, to the remote computer.

It should be stressed that the system may include more than one remote monitoring computers, and any of these may not, in fact, be “remote”(in terms of distance) at all; instead, they may reside at the same site as the controlled process or equipment. Ordinarily, the monitoring computer or computers operate as peers with respect to the controllers on a flat network topology.

The invention facilitates a complete window into the operation of one or more controllers and, therefore, the industrial equipment they operate. Using the invention, remotely located personnel can monitor the efficiency or overall behavior of the equipment, perform diagnostic checks, or even effect certain maintenance operations. For widely dispersed control and supervisory operations, supervisory computers interact with the controllers over the Internet, with the controllers continuously connected to the Internet as “nodes.” In local operations, the flexibility conferred by Internet formalisms can be retained on a restricted, internal network.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing discussion will be understood more readily from the following detailed description of the invention, when taken in conjunction with the accompanying drawings, in which:

FIG. 1 schematically depicts a controller in accordance with the present invention; and

FIG. 2 schematically depicts a system incorporating multiple controllers as shown in FIG. 1 and a browser-equipped computer capable of accessing data associated therewith.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The Internet, which can provide the communication medium of the present invention, is a worldwide “network of networks” that links millions of computers through tens of thousands of separate (but intercommunicating) networks. Via the Internet, users can access tremendous amounts of stored information and establish communication linkages to other Internet-based computers.

Much of the Internet is based on the client-server model of information exchange. This computer architecture, developed specifically to accommodate the “distributed computing” environment that characterizes the Internet and its component networks, contemplates a server (sometimes called the host)—typically a powerful computer or cluster of computers that behaves as a single computer—that services the requests of a large number of smaller computers, or clients, which connect to it. The client computers usually communicate with a single server at any one time, although they can communicate with one another via the server or can use the server to reach other servers. A server is typically a large mainframe or minicomputer cluster, while the clients may be simple personal computers. Servers providing Internet access to multiple subscriber clients are referred to as “gateways”; more generally, a gateway is a computer system that connects two computer networks.

In order to ensure proper routing of messages between the server and the intended client, the messages are first broken up into data packets, each of which receives a destination address according to a consistent protocol, and which are reassembled upon receipt by the target computer. A commonly accepted set of protocols for this purpose are the Internet Protocol, or IP, which dictates routing information; and the transmission control protocol, or TCP, according to which messages are actually broken up into IP packets for transmission for subsequent collection and reassembly. TCP/IP connections are quite commonly employed to move data across telephone lines.

The Internet supports a large variety of information-transfer protocols. One of these, the World Wide Web (hereafter, simply, the “web”), has recently skyrocketed in importance and popularity; indeed, to many, the Internet is synonymous with the web. Web-accessible information is identified by a uniform resource locator or “URL,” which specifies the location of the file in terms of a specific computer and a location on that computer. Any Internet “node”—that is, a computer with an IP address (e.g., a server permanently and continuously connected to the Internet, or a client that has connected to a server and received a temporary IP address)—can access the file by invoking the proper communication protocol and specifying the URL. Typically, a URL has the format `http://<host>/<path>`, where “http” refers to the HyperText Transfer Protocol, “host” is the server’s Internet identifier, and the “path” specifies the location of the file within the server. Each “web site” can make available one or more web “pages” or documents, which are formatted, tree-structured repositories of information, such as text, images, sounds and animations.

An important feature of the web is the ability to connect one document to many other documents using “hypertext” links. A link appears unobtrusively as an underlined portion of text in a document; when the viewer of this document moves his cursor over the underlined text and clicks, the link—which is otherwise invisible to the user—is executed and the linked document retrieved. That document need not be located on the same server as the original document.

Hypertext and searching functionality on the web is typically implemented on the client machine, using a computer program called a “web browser.” With the client connected as an Internet node, the browser utilizes URLs—provided either by the user or a link—to locate, fetch and display the specified documents. “Display” in this sense can range from simple pictorial and textual rendering to real-time playing of audio and/or video segments or alarms, mechanical indications, printing, or storage of data for subsequent display. The browser passes the URL to a

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.