

X1[®] Enterprise Edition Scalability

February 2006

X1 Technologies, Inc.



Abstract

The purpose of this paper is to illustrate how to increase the power and performance of the X1 Platform to accommodate large enterprise deployments without significant extra resources or incurring significant costs. Three key search components - the X1 Enterprise Client, the X1 Enterprise Server and the X1 Browser Client; plus the X1 Deployment Manager administrative tool, work together to enable virtually unlimited scalability at low or no cost and simplify the process for IT administrators.

Table of Contents

Abstract	2
Introduction	4
X1 Enterprise Edition Components	4
X1[®] Enterprise Client	4
X1[®] Deployment Manager	5
X1[®] Enterprise Server	5
X1[®] Browser Client.....	6
Scaling an X1 Installation.....	6
Scaling the X1 Enterprise Client	6
Scaling the X1 Enterprise Server	6
Scaling the Default X1 Browser Client	8
Conclusion	8
About X1 Technologies	9

X1[®] Enterprise Edition Scalability

Introduction

The proliferation of systems and data has led to a more complex enterprise environment that makes information search ever more complex. The organization's capacity to collect data has scaled exponentially, but the employees' ability to find that data has not. Employees can spend up to 25% of their workweek conducting searches within multiple systems, sometimes without effect. Remote users require equivalent search capabilities but face security and bandwidth limitations.

Typical desktop search tools search local hard drives and email, but not packaged or custom applications or relational databases. Enterprise versions of these tools typically require significant additional investments in order to scale to meet performance demands.

The X1 Enterprise Edition helps organizations recover this lost productivity by providing unified, actionable search for internal and external end users in complex, heterogeneous environments, with an architecture that was designed from the ground up to be scalable.

With the X1 solution, end users can conduct searches through a local Enterprise Client, a Browser Client or Mobile Client. X1 enables unified information access, wherever users are, for whatever they need, across all their mission-critical business information (including the desktop, server hosted applications and storage). This dramatically improves productivity, reduces wait and browse times and eliminates the cost of work duplication.

X1 embraces the IT department's requirement for tools that are easy to setup, configure and deploy in small- and large-scale environments. X1 tools empower IT administrators to add to and manage large clusters of servers and server components without committing to a single hardware platform, operating system or proprietary platform. Deployment tools allow IT administrators to pre-select, lock down, or create department-specific configurations of the X1 Client. Finally, the "zero-footprint" Browser Client enables administrators to deploy X1 rapidly to a very large number of users.

X1 Enterprise Edition Components

The X1 Enterprise Edition at its core consists of the award-winning X1 Enterprise Client, the high-performance X1 Enterprise Server and the optional X1 Browser Client. Each of these components plays a part in enabling extreme scalability of the search solution.

X1[®] Enterprise Client

The award-winning X1 Enterprise Client enables knowledge workers to search and instantly find information across the environment, preview it in its native format, and take action on it immediately. X1 provides unified, actionable search across an individual's desktop or notebook computer and critical enterprise locations. X1 not only solves the problem of information overload for overwhelmed end users, but also enables IT to centrally manage desktop search and enables application teams to encourage efficient use of their systems.

Files, email messages, attachments and contacts are all indexed locally or, at the administrator's option, on one or more X1 Enterprise Servers. The Enterprise Client includes support for creating silent installers that are compatible with most major desktop software management systems used by large organizations.

The X1 Enterprise Client integrates with one or more X1 Enterprise Servers using its “Add Server” feature. When connected to an X1 Enterprise Server, the X1 Enterprise Client can search across local and server-indexed information through existing searches in the client, new server-defined searches and unified searches which look across both local and server-indexed information in a single search.

X1 can act as a client to multiple heterogeneous search systems at the same time. This approach allows our customers to leverage their existing investments in search technology. For example, the X1 search framework can be configured to allow X1 Enterprise Client users to simultaneously perform searches across X1 Enterprise Servers, IBM OmniFind systems and Oracle Secure Enterprise Search installations. The unique search filtering, sorting, preview and result highlighting capabilities of the X1 platform complement back-end technologies that were designed for more traditional browser-based, page-at-a-time results interaction. Through a standardized set of protocols and APIs, X1 is capable of unifying search results from any back-end search platform – Verity, Autonomy, Fast, Alta Vista, Infoseek, FileNet, Open Text and others.

X1[®] Deployment Manager

The X1 Deployment Manager empowers the IT administrator to completely configure X1 Enterprise Clients with features, settings and customizations appropriate to their particular enterprise environment, and deploy the Client across the network. The Deployment Manager works with existing software deployment processes and is compatible with Microsoft SMS, Windows Script Host[®] and other popular enterprise deployment tools. Enterprise Clients can be deployed “silently” to minimize impact on end users.

X1 Enterprise Client features can be turned on and off or grayed out and left visible in the interface, configured with default settings, or added to toolbars and menus. Customizations, including support for new X1 Content Connectors for server-hosted information, can be added through the X1 Enterprise Client’s plug-in facility. An administrator can customize the Client for multiple different groups and users, saving each exact configuration for later reference or reuse. The X1 Deployment Manager produces an MSI-compliant installer that can be pushed out to individual desktops using the organization’s software installation method of choice, including logoff/logon scripts and software configuration management tools.

Once deployed, the X1 user configurations can be managed using Microsoft’s Group Policy Objects and extended to support other system management frameworks.

The Client Extension Framework allows IT departments and third parties to build further enhancements and their own functionality into the X1 Enterprise Client.

Deployment Manager was designed to address the needs of organizations with thousands of desktops, multiple geographic locations and large mobile groups. It allows tightly controlled, centralized deployments using staged rollouts and specific customizations, both functional and operational, for the needs of different types of users.

X1[®] Enterprise Server

In organizations where large quantities of content are stored on network data sources such as file shares, relational databases or server-based applications, allowing end users to index this content on each desktop wastes space, consumes excessive bandwidth, opens security holes and creates administrative headaches.

As the heart of the X1 Platform, the X1 Enterprise Server automatically and transparently indexes the content housed in multiple information repositories, federates search results and delivers those results as quickly as the end user types. It is a central indexing engine for network content sources, which enables users to search for network content from either Enterprise Clients or Browser Clients in an “index-once-query-many” fashion. The X1 Enterprise Server can connect to new data sources by adding pre-built X1

Content Connectors, such as for Microsoft Exchange, or by using X1's public APIs and SDKs to build new X1 Content Connectors to other systems.

The X1 Enterprise Server also includes the X1 Cluster Manager, which features real-time management and monitoring capabilities across all X1 Enterprise Server components – Index Engines, Connector Engines and Federation Engines. The system administrator can access and preview indexed information, define searches¹ against this information, publish searches and define the security on published searches. In addition, the X1 Cluster Manager acts as configuration and management tool for the entire X1 Cluster, allowing an administrator to obtain a single, unified view of their X1 cluster and perform administrative tasks on it in a rapid and cohesive manner.

X1[®] Browser Client

For many employees the most relevant information resides on the network rather than local desktops, so the ability to find server-based information can be critical. For example, an IT department may prohibit local storage of documents, and teams may utilize a shared workspace. Moreover, some employees need remote access capabilities for their data whether they are on the road, work remotely or have special access needs.

For those users, X1 offers the X1 Browser Client, a Web-based application that enables end users to access their critical server-based information. The Browser Client delivers the X1 style "find-as-you-type" interface and the dynamic preview capabilities of the Enterprise Client. The X1 Browser Client only requires Internet Explorer browser to operate and does not require installation of any custom ActiveX Controls.

The X1 Browser Client runs on Internet Explorer and Firefox and supports mobile and remote office access. The Browser Client maps to end users' existing permissions and can be extended or customized to fit within existing portal or intranet systems to preserve a company's standard look and feel.

Scaling an X1 Installation

There are several independent but related scaling challenges for the X1 Enterprise Client, the X1 Enterprise Server and the default X1 Browser Client. Each challenge involves several different factors and dynamics, including how these components operate with each other.

Scaling the X1 Enterprise Client

The key scaling parameter for the X1 Enterprise Client is determining which custom configuration of the X1 Enterprise Client you need to provide to each user in your environment. By running the X1 Deployment Manager, an administrator can create customized X1 installers with the appropriate features and Content Connectors for specific departments and users.

For example, an administrator can create an installer for the Finance department that includes connectors for Oracle Financials and Microsoft Exchange, then create a separate installer for the Operations department that only includes the Microsoft Exchange connector. Administrators should use the X1 Deployment Manager to maintain and modify X1 Enterprise Client installations. The X1 Deployment Manager produces an MSI installer which can then be pushed out to end users using standard software configuration management packages such as Microsoft SMS and Novell ZenWorks.

Scaling the X1 Enterprise Server

Scaling the X1 Enterprise Server involves consideration of several different parameters, including:

¹ An X1 "Saved Search" is the exact logical equivalent to a RDBMS "View". It pre-selects a result set and allows users access to these named "Saved Searches" as buckets of pre-defined information. The X1 framework allows administrators to manage user and group access to saved searches using standard Windows Active Directory or LDAP user/group management.

- How much data will be indexed in total
- Which backend systems the X1 Enterprise Server must connect to
- How many X1 Enterprise Clients, X1 Browser Clients and custom integrations will simultaneously query the server
- What the network topology is within the deployment environment
- What requirements exist for fault tolerance and data redundancy

The X1 Server itself has no architectural or design limitations to scaling since it uses a near-linear, shared-nothing architecture that allows the scaling of an X1 cluster (an instance of the Enterprise Server) to an extremely large index and connection size. Customers have successfully indexed many terabytes of information in a single X1 cluster. In addition to the scalability of the X1 Enterprise Server, the X1 architecture allows for federation to external indexes of content from other X1 clusters and other search indexes such as Google, Autonomy or Fast. This ability not only protects existing knowledge management investments but also allows for a single point of access to all corporate information assets.

As described above, the X1 Enterprise Server consists of multiple components that work together to index enterprise content, federate results and deliver results. These components can be configured, deployed and administered using the centralized X1 Cluster Manager.

- The *X1 Connector Engine* runs the X1 Content Connectors, which authenticate and connect the Server to key corporate data repositories, such as files, e-mail, databases, and content management systems. One or more Connector Engines can execute in parallel. Connector Engines may execute on dedicated servers on in combination with all other X1 Enterprise Server components on the same server. Multiple Content Connectors can be simultaneously loaded, unloaded, scheduled and operated within each Connector Engine. Different back-end systems can take advantage of three different types of X1 Content Connectors:
 - Traditional scanners, or crawlers, treat data sources as virtual file systems and “walk” the hierarchy tree, indexing content as it is found.
 - LiveDB connectors use polling to query systems for new and changed information at predetermined and scheduled intervals.
 - Real-time connectors are listeners, which allow third-party systems to insert information into the index in an event-driven fashion.

This highly configurable and flexible environment allows performance and scalability tuning to meet even the most challenging configurations.

- The *X1 Index Engine* indexes all the documents sent to it by the Connector Engine while simultaneously responding to queries from the Federation Engine. Each Index Engine can create and use multiple Index Instances. This allows separate indexes to be located on separate disks, which maximizes the disk-to-RAM I/O bandwidth. An X1 cluster does not have a design limitation on the number of Index Engines and Index Instances that can be supported.

In general, for scalability purposes, seek a balance between the speed with which the Connector Engines read information from the source systems and the speed with which the Index Engines process that information. X1 uses Java Spaces technology to manage the flow of information between the physical and logical nodes of the X1 cluster. The X1 Cluster Manager is a unified monitoring console that shows, in real time, the throughput of all Engines, and allows the throttling of components to optimize scalability. The X1 Engines also implement self-management logic to regulate and optimize the flow of information through the Enterprise Server – for example, if the Index Engines are slower than the Connector Engines, and cannot “take” items from the Java Space as fast as items are “put” into it, then the Connector Engines will automatically be paused so that the system can catch up.

Additionally, all Engines are engineered to be self-restarting and the Java Space provides persistent storage of transient information. This means that the X1 cluster can be shut down and when restarted, any data that has been “put” for indexing will not be lost.

- The *X1 Federation Engine* receives queries from X1 Clients (or other custom-built applications) and compiles the responses from all the Index Engines in a cluster. Using X1 Software Development Kits, developers can use the Federation Engine to embed the X1 Enterprise Server into their applications.

The scalability aspects of the underlying design account for the fact that when results are federated from a large number of remote search subsystems, there cannot be a guarantee of how fast any one of them will return results. Federation is even possible, although not recommended, over slow WAN connections. To optimize query speed, the federation engine uses an asynchronous set of protocols, so that the results from the fastest Index Engine or third-party search subsystem are returned immediately without waiting for the slowest system to respond. This allows the X1 Enterprise Client to provide fast-as-you-type usability across huge sets of enterprise data without the latency typically associated with traditional search federation platforms.

- The X1 System Management Engine silently manages system administration, remote administration, performance monitoring, logging and other tasks for all the other Enterprise Server components.

To improve performance and scale the capacity, it is as simple as drag-and-drop using the X1 Cluster Manager. Administrators can remotely configure and install each of the components in the X1 Enterprise Server and monitor the performance of the entire cluster. Improving performance in any area is as simple as adding an additional node and configuring the appropriate X1 Engine.

Scaling the Default X1 Browser Client

The default X1 Browser Client is installed and maintained as part of an X1 Enterprise Server installation. The single most important criteria for sizing an X1 Browser Client installation is to estimate the number of Browser Client users as well as their usage level. It is a Microsoft IIS/ASP.NET application and scaling it should generally follow the best practices for operating these technologies.

Generally, customers with existing Intranet and Portal implementations would use the X1 query protocols and APIs to integrate the X1 Platform capabilities within their existing deployments, rather than operate the default X1 Browser Client. X1 has successfully integrated with third-party web applications and portals using Microsoft ASP, .NET, JSP, Java Servlets and Perl environments. X1 does not introduce additional scalability and hardware requirements in such implementations.

Conclusion

X1 delivers the only truly scalable solution for unified, actionable search on the market today through the X1 Enterprise Edition product set. The X1 Enterprise Edition can scale to the multi-terabyte data loads commonly encountered in large organizations while still providing administrators with a clear, easy and cohesive interface for managing the installation and providing quality services to end users.

About X1 Technologies

X1 Technologies, Inc., a recognized leader in desktop search solutions, was founded in 2003 to help business users easily access and act upon information that resides anywhere on a desktop or the corporate enterprise. Companies of all sizes use X1 Technologies' solutions for their enterprise desktop search initiatives. X1 Technologies was an early leader in the desktop search space with the launch of their flagship product, X1 Desktop Search, which garnered industry awards including the "Best of 2004" and "Technical Excellence" from PC Magazine. X1 Technologies has partnered with leading companies including IBM, Yahoo! and EarthLink to power their desktop search initiatives. Headquartered in Pasadena, California, X1 Technologies is an operating company of Idealab, and is backed by U.S. Venture Partners. For more information, visit www.x1.com or call 626-229-3050.



The X1 Platform

Employees spend as much as 25% of their work week searching, and often failing to find, critical data across the enterprise.

- The X1 Platform puts the full power of unified, actionable search into the hands of your users, enhances your information security and integrates with your critical business systems.
- X1 combines the most comprehensive search and preview abilities on the market, as-fast-as-you-type results, easy configuration and seamless integration. Organizations of all sizes, across all industries, have demonstrated significant productivity gains and reduced risk through X1's unified, actionable search.

X1 Works Like You Do, Not the Other Way Around

Knowledge workers have too much information and not enough time to find it. The X1 desktop search model recognizes that people work in different ways, with various organization skills. They generally seek the shortest way between Point A and Point B, irrespective of the "approved" path.

The Complete X1 Platform Consists Of:

- The award-winning X1 Client - bringing desktop and enterprise search together
- The high-performance X1 Enterprise Server - the engine of the X1 enterprise solution
- X1 Content Connectors - integrating search into your critical enterprise applications
- X1 Client and Server SDKs - to tailor the X1 experience to your end users