Scaling the Heights of E-Business Demands

PROGRESS SOFTWARE

Contents

Introduction

> *"Users hunger for a comprehensive solution that provides all of the functionality required to build, deploy, and maintain scalable e-business applications."*
>
> – **IDC Bulletin, "Enterprise Portals: From Mania Toward a Mature Ecosystem"**

The Internet is changing how we do business. It has evolved from being a place to gather information, to a crucial method for both e-businesses and brick and mortar companies to deliver applications, products, and services. Today's commercial Web sites have grown from simple information sites to sophisticated e-business sites that provide customers with applications that are crucial to their organization's growth and operations.

These new sites often have a worldwide user base and produce millions of dollars of revenue. To support this model, e-business solutions must have enterprise or better capacity for robustness, scalability, accuracy and reliability.

The adoption of the Web site as a means of doing business has irrevocably changed the minimum requirements an application must satisfy to be considered "usable." Many applications perform adequately in an environment where the number of users, number of transactions, and growth of the database are predictable. However, in an e-business environment, applications may need to support thousands of users, millions of transactions/day, and many GBs of data. All this must be supported at satisfactory performance levels. An application that is subjected to exponentially increasing demands quickly reveals resource shortcomings and performance bottlenecks.

The challenge for the application developer, therefore, is to provide full-featured, top-quality applications that can scale to accommodate an ever-increasing number of users and transactions. All without sacrificing performance.

Progress Software is a leading supplier of software technology for e-business solutions. From development to deployment to maintenance, the Progress Version 9 and WebSpeed Version 3 product families allow our users to design and develop high performing, scalable e-business applications.

What is Internet Scalability?

Simply put, scalability is the ability of an application to deliver a consistent set of functionality in a timely manner, independent of the number of users, transactions, or connections. The essence of scalability is the capability of an application to start small and grow larger and larger, without discontinuities, when business requirements change.

Scalability is an essential element of any e-business application. When businesses offer application access over the Internet, or over extranets, those applications must be ready for unpredictable fluctuations in user access. They must respond consistently, regardless of traffic. E-business applications must be able to handle rapid, unplanned changes in capacity.

Scalability is critical. Never before have there been applications with such high requirements for simultaneous user access, number of transactions, and amounts of data to be analyzed and stored. Managing scalability ensures that an e-business application will be able to handle the most demanding user and transaction loads and won't fail, as usage increases, because of performance issues.

Human factors research has shown that the average Web viewer will only wait 8 seconds for a Web page to load, before finding an alternative. As the Internet has become easier to use, more people with fewer technical skills are using it to obtain information and transact business. For these people, e-business applications must be scalable and retain high performance, regardless of the time of day or number of people accessing the application.

> *"The decreasingly technical profile of the user of such sites will translate into decreasing tolerance for wait times, which means that slow systems and downtime are not options. All systems must respond quickly with up-to-the-minute information and must be available 24 hours a day."*
>
> – **IDC: "Information and Data Management Systems Market Forecast and Analysis, 2000 – 2004"**

Designing for Scalability

When designing a scalable solution, developers need to answer the following question:

> *How much effort and time is required to successfully expand the usage of this application to a large number of users with an ever increasing number of transactions?*

To provide a truly scalable solution, both the infrastructure and the application must be able to handle increases in traffic. Capacity planning should be used to determine the contents of the infrastructure (i.e. how many application servers, database servers, etc.). Most e-business solution providers plan for and maintain excess resource capacity.

> *Design the entire architecture for horizontal scalability. By partitioning Web site functionality into components that live on different systems, enterprises achieve higher levels of availability and "pluggable" scalability. The design goal is to be able to add more capacity by deploying another server and load balancing a portion of the connections to it.*
>
> – **Barb Gomolski, "Top 10 Recommendations on Building Scalable Hi-Performance Web Sites," E-Business Matters InfoWorld Column, 1/16/01**

Developers should also design their applications to handle large increases in users, transactions, and data by using tools and products that provide scalable elements.

Distributed Computing

Distributed computing is an architectural approach to designing an application that runs across a network. It maximizes an application's processing capability and increases its capacity to scale by allowing the business logic to be distributed throughout the enterprise.

Distributed enterprise applications offer the best scalability and performance, faster time-to-market (by re-using application components), and the ability to fully embrace the power of the Web.

Progress' Universal Application Architecture (UAA) enables application developers to choose virtually any client, any database and any platform for true plug-and-play development. With the explosion of interest in e-business, the Progress Version 9 and WebSpeed Version 3 product families enable our customers to expand their business by extending applications to the Web.

Extending applications to the Web often encompasses using the power of the Internet as a platform for distributed computing. You can deploy and apply resources throughout its extended network. The UAA helps application developers implement two important aspects of distributed computing--application partitioning and the n-tier model.

Application Partitioning

Application Partitioning is the process of segmenting the business logic from the user interface and data, and running it across multiple servers on an intranet, extranet, or the Internet.

All applications deal with three types of components: user interface (the client), business logic (server), and data (RDBMS). The process of splitting an application's code among clients, application servers, and database servers is commonly referred to as application partitioning.

By partitioning applications into components and moving logic to application servers and away from the client and user interface, developers can begin to reap the benefits of re-use by sharing components. Another benefit of the distributed architecture is the ability to locate the business logic in one place (on 1 or more servers), rather than in each client application.

N-Tier Model

The n-tier model supports the separation of user interface, business logic and data across three or more machines, as Figure 1 illustrates.

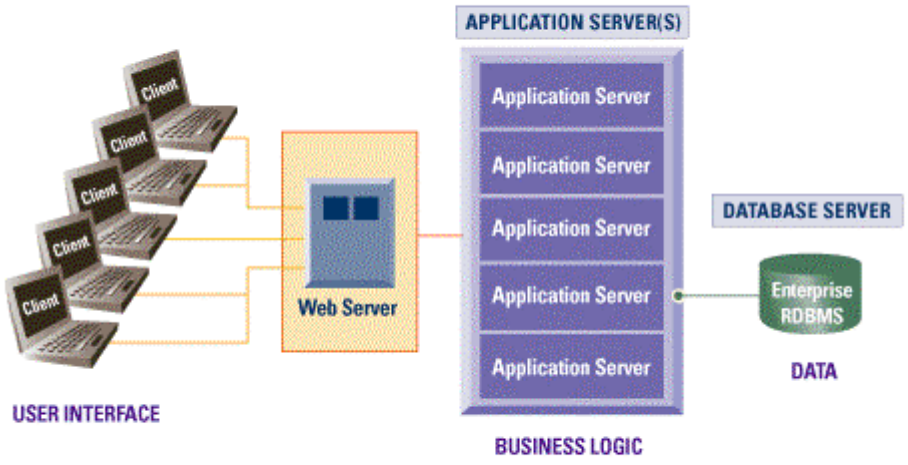

Figure 1: N-tier architecture

Figure 1 depicts an n-tier computing model where one or more application servers are inserted between the client (user interface) and the database. The application servers handle the processing of complex business logic, thereby freeing up resources on the client and minimizing network traffic of data from the server to the client.

By designing an n-tier application, developers can easily maintain or improve scalability and avoid performance bottlenecks that can occur at a single client or server machine.

The Progress Universal Application Architecture

Progress Software Corporation provides application developers with components they need to make their applications scalable and optimal for a distributed computing environment. The Universal Application Architecture (UAA) provides developers with the tools and strategic framework required to create high performing, scalable distributed applications, and/or to transform their existing business applications into open, Future Proof™ applications. These applications can communicate easily with other applications, databases, or services available anywhere on the network, and utilize multiple user interfaces.

In the UAA model, the business logic of an application resides primarily on the server, accessed by users with thin clients or Web browsers. Application code that is most suitable for client side execution, such as user interface logic, is distributed as needed to the client but managed on the server. This improves application performance and scalability by reducing network traffic while retaining a responsive user interface. It also makes administration and maintenance easier because all upgrades and changes are made to the application residing on the server, rather than to major application components residing on many clients scattered throughout the Internet.

A UAA application is network-, platform-, and UI- independent. It allows developers to build GUI, HTML, Web-based GUI, Java, ActiveX, or character-based interfaces in response to customer requirements.

All Progress products support the UAA model, and provide developers with tools that will allow them to design high-performance, scalable applications.

Scalable Components

When designing an application for e-business, scalability must be explicitly planned for and a central focus of each component of the solution. Whereas many times scalability is used in reference to the database only, all aspects of an e-business application (client, servers, databases) must be able to respond to a large number of users, who may connect at any time, and meet the performance demands of those customers. Scalability is a key focus for each component of a Progress Software application, including Progress clients, application servers, databases and messaging services.

Clients

Traditional fat-client applications quickly begin to suffer performance problems as the application becomes more complex and more data is required to be transmitted over the Web. In addition, the footprint and processing power of the client machine (usually a PC) is strained by the constant need for resources. Splitting out the business logic, from the UI, into components allows the logic to execute on multiple platforms without modification.

Progress provides 2 thin-client architectures, the WebClient and a Web browser-based architecture.

Progress WebClient. The Progress WebClient is a thin client that provides a full graphical user interface over the Web. WebClient provides a high-performance, rich UI, and robust, scalable architecture. Progress WebClient was developed with the expectation that an application's business logic is located on a Progress AppSever. Therefore, it requires that the application be developed with proper separation of UI and business logic. Progress WebClient runs as a free-standing application, on an end-user's PC, executing the Progress 4GL user-interface code and communicating to the Progress AppServer using HTTP, HTTPS tunneling, or TCP/IP.

Web Browser. A Web browser can be used as a client when an application is developed using Progress WebSpeed. Progress WebSpeed is a powerful solution for quickly building and deploying business applications across corporate intranets, extranets, and the Internet. WebSpeed provides a complete and highly-productive Integrated Development Environment. A developer can develop and deploy applications that use XML, HTML, DHTML, JavaScript, and/or Java. To deploy a WebSpeed application, a Web browser serves as the client and the business logic is located on the WebSpeed Transaction Server.

Application Servers

The application server is the middle tier of an n-tier architecture. It sits, where the business logic is executed, between the client and the database. Through load-balancing and stateless connections, application servers can greatly increase an application's performance and scalability.

Load balancing is used when multiple application servers are employed in a single application. This provides the application with the ability to balance out the transaction load across the application servers, ensuring that users are not left waiting for available processes to execute their business logic.

Stateless servers add powerful scalability to applications. A client application connected to an application server running in stateless mode forms a direct link for the duration of a single process. An application broker directs the request to any application server process that is not busy or puts the request into a message queue if all processes are in use. With stateless application servers, a few server processes can handle many client applications.

Progress Software provides two types of application servers, the Progress AppServer and the Progress WebSpeed Transaction Server. A third type of server, the Progress NameServer, provides load balancing and location transparency for the application servers.

Progress AppServer. The Progress AppServer allows the developer to distribute Progress (4GL) applications across multiple machines in a network. Through application partitioning, the AppServer enables the separation of business logic from user interface in network- or processing-intensive applications. The Progress AppServer includes stateless server technology, message queuing, an integrated name server, and support for multiple client application platforms. These features optimize networking performance, increase application scalability, and provide flexibility to support a diverse client population. The Progress AppServer can access several different clients, including the Progress WebClient, the GUI client and the OpenClient, which can be ActiveX, Java, or C++ based. Clients can communicate with the Progress AppServer via TCP/IP, HTTP or, in the case of the Progress Secure AppServer, HTTPS. Figure 2 depicts a typical n-tier configuration using the Progress WebClient and the Progress AppServer.

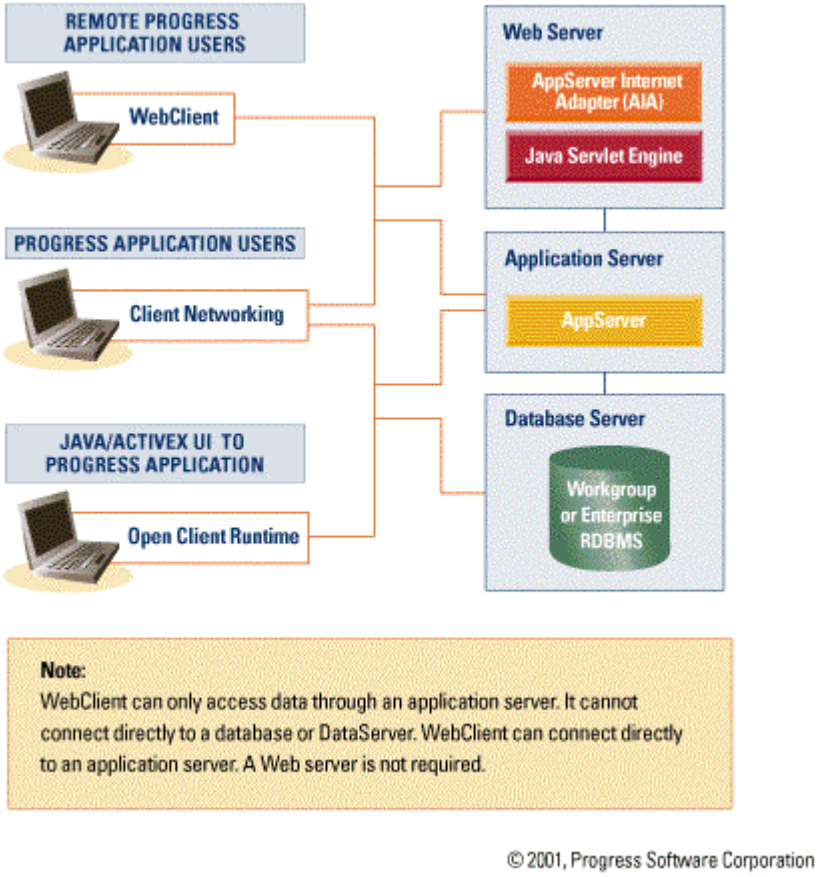

Figure 2: N-tier architecture with the Progress WebClient and the Progress AppServer

WebSpeed Transaction Server. The Progress WebSpeed Transaction Server provides an optimized transaction-processing environment for achieving high transaction volumes and rapid responses for Web applications. With high throughput, and dynamic load balancing, applications can have unsurpassed scalability to handle thousands of simultaneous users.

Dynamic load balancing ensures high availability of transaction processing resources in a distributed n-tier environment. The Progress WebSpeed Transaction Server also allows developers to leverage common business logic for Web applications. Flexible state management offers full support for extended database queries and updates using stateless, state-aware, or state-persistent Web objects. Figure 3 depicts a typical n-tier configuration using Progress WebSpeed Transaction Server and Web browsers.

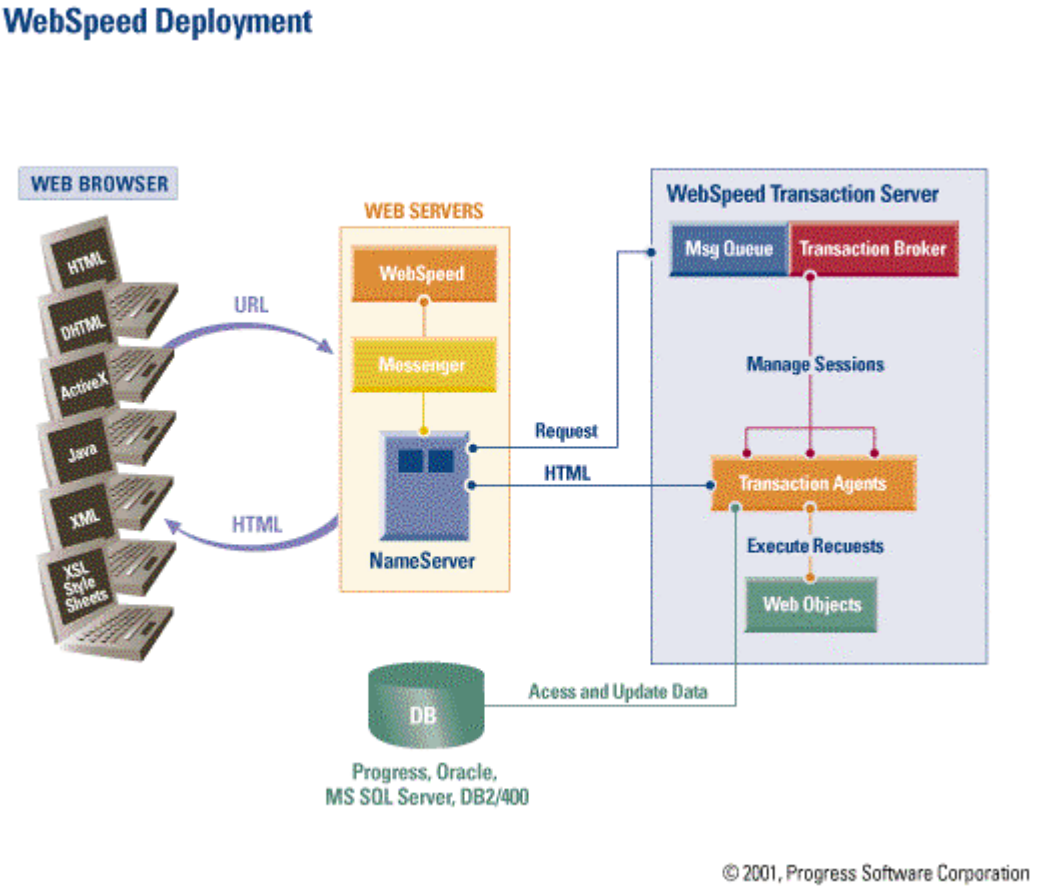

Figure 3: WebSpeed Configuration

Progress NameServer. The Progress NameServer supports the Progress AppServer and Progress WebSpeed Transaction Server by providing location transparency and load balancing.

With location transparency client applications don't have to know about all of the application servers that are on the middle tier. The client simply makes a connection to a NameServer, requesting to execute some business logic on an application server. The Progress NameServer is responsible for keeping track of where the business logic exists and then re-connecting the client with that 'service.' As additional application servers are added to the middle tier or moved to other locations, the name server records these events. The client application doesn't have to be modified. This allows developers to increase the resources of an application without requiring any changes to, or notification of, the clients.

The Progress NameServer also provides the ability to do load balancing across multiple AppServers or WebSpeed Transaction Servers. Again, this is all transparent to the client application.

Database

To meet the needs of today's e-business applications, an RDBMS server must provide the highest possible levels of performance in a variety of environments. In an e-business or enterprise-level environment, the RDBMS must scale gracefully as users are added. The RDBMS must take advantage of high performance hardware to provide the fastest response time and transactions per second. The requirement can be summarized as providing a fast user response time regardless of how many users are connected.

The Progress RDBMS leads the database market with low cost of ownership and the ability to scale to meet the demands of enterprise applications, e-commerce, and application service providers. The Progress RDBMS can scale from small workgroup to large enterprise and Web-based applications accessed by thousands of users. There are two fundamental areas where a database must be scalable: the number of concurrent users supported, and the storage capacity.

To achieve scalable e-business applications, Progress Software recommends the use of the Progress Enterprise RDBMS.

The Progress Enterprise Database Server is recommended for high-volume, distributed, enterprise and e-business applications. It includes a multi-user relational database server that provides database services for both production operation and development activities. The Enterprise Database Server should be purchased for an application that includes any of the following characteristics:

- A moderate (30) to high (5,000) number of concurrent users
- A moderate (2 Gigabytes) to large (250+ Gigabytes) database size, using multiple disk drives
- Real-time distributed updates performed across more than one database or across more than one computer system
- Executes in a symmetric multi-processor (SMP) environment that includes two or more CPUs

The Progress Enterprise RDBMS is highly scalable in a number of important ways, including storage and user-count scalability.

Storage Scalability. Almost all systems acquire more data over time. This creates a need for more and more storage space. A small database can be easily accommodated on a single disk drive, and can grow to use multiple disk drives. Flexible storage allocation policies allow control over layout and effective use of available storage media. By using Progress storage areas, administrators can group tables and indexes to take advantage of premium hardware or to allow easier logical data management.

User-Count Scalability. The Progress RDBMS and the Progress AppServer have been architected to handle the most demanding e-business requirements. Progress recently worked with NxTrend, makers of a leading Distribution Application, and IBM to provide a real-world benchmark of the Progress RDBMS. The benchmark simulated real users of the Trend Application. The results of the benchmark showed a linear scalability from 5 to 5,000 concurrent users and the capability to handle more than 10,000 concurrent users with superior reliability and multi-terabyte storage capacity.

> *We continually try to extend our business to larger and larger organizations. The ability to show that our system will scale to 5,000 users will give us a competitive advantage and help us expand our market.*
>
> – **Doc Chu, Research Engineer, NxTrend Technology, Inc.**

In additional benchmarking using the Progress ATM Benchmark application, the Progress database reached a performance record for single-server transactions by processing over **one billion transactions in a 24-hour period**. The database achieved a peak transaction rate of over 12,175 transactions per second and a sustained transaction rate of 11,522 transactions per second. The Progress ATM Benchmark simulates the third tier (database component) processing associated with a large ATM network with an update-intensive transaction processing workload. The transactions are financial strength providing full audit information and recoverability in the event of a system failure. During testing, 99.9% of all transactions were completed in 0.1 seconds or less, and all transactions were complete in 0.6 seconds or less. (For details on the Progress ATM Benchmark testing, see the report at www.progress.com/benchmark/v91.htm.)

Messaging Services

> *Messaging is critically important to the execution and management of business processes within, and external to, an enterprise. Its relevance to ebusiness is extremely high and goes hand in hand with integration of business processes, legacy assets, and disparate platforms, applications, and data sources.*
>
> – **IDC Bulletin, "The eBusiness Platform: When Application Servers and Application Integration Meet"**

Applications on the Internet must be able to integrate and interoperate with other applications. The process of integrating with other business systems is a daunting task without a communication system that is specifically designed for Internet-based integration. E-business applications require a reliable, fast and scalable communication between a broad variety of newly developed and legacy applications and data sources. Communication in this loosely coupled environment is made possible by a specific genre of middleware called Message-Oriented Middleware (MOM).

A messaging server must provide industrial-strength, scalable messaging capable of accommodating heterogeneous platforms with high levels of performance and integrity.

SonicMQ. Progress SonicMQ is the leading solution for E-Business Messaging (EBM), a rapidly growing category of messaging middleware designed to meet the unique requirements of information integration and exchange over the Internet. Offering a highly efficient, robust, cost effective messaging solution to address enterprise application integration, SonicMQ can scale up to support full-blown highly-distributed, highly-scalable Internet messaging requirements. Using the Progress SonicMQ Adapter, Progress users can add secure business-to-business communication and facilitate enterprise application integration to fully integrate SonicMQ with their Progress-based e-business applications.

SonicMQ provides the highest level of scalability available in a messaging server today. The Dynamic Routing Architecture (DRA), introduced in the latest version of SonicMQ (V3.0), provides the ability to dynamically add tens of thousands of business partners to an e-business application. Running SonicMQ allows users to participate in a massively scalable messaging infrastructure delivering millions of messages per day.

Benchmark tests have quantified this level of performance regardless of the messaging model an application uses (Publish and Subscribe or Point-to-Point). The results illustrated in Figure 4 show how SonicMQ performs compared with IBM MQSeries® messaging service. Performance gains increased for SonicMQ as the number of messages grew, verifying that SonicMQ not only scales to handle large numbers of messages, but that its performance scales to meet the increased demands as well.

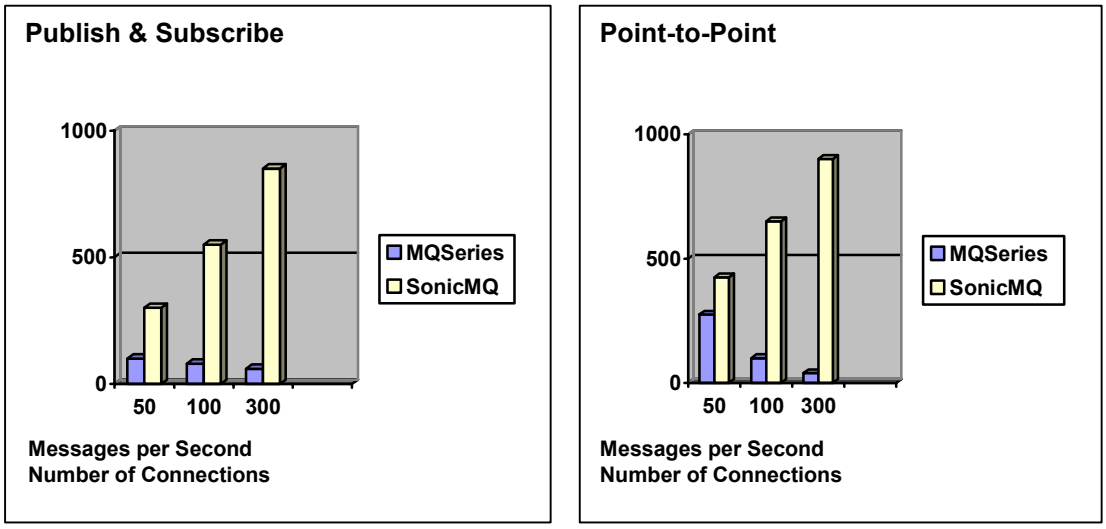

Figure 4: SonicMQ Performance and Scalability

Conclusion

> *The old guard – exemplified by traditional, monolithic applications residing on either back-end systems or on fat clients or servers in a client/server configuration – is giving way to a new way of building and deploying solutions, driven in large part by the Internet and resulting e-business initiatives*
>
> – **IDC Bulletin, "The eBusiness Platform: When Application Servers and Application Integration Meet"**

The Internet **is** changing how we do business. To be successful, applications must meet the challenges and requirements demanded by Internet and e-business users. 8 seconds to deliver the goods is a difficult requirement for a Web application. To achieve this, and to satisfy the other e-business requirements, scalability, robustness and application modularity (separation of user interface from business logic) must be the essential elements in every e-business application.

Progress Software delivers the embedded technology that enables developers to build scalable, robust, modular e-business Future Proof™ applications

Corporate Headquarters
Progress Software Corporation, 14 Oak Park, Bedford, MA 01730 USA Tel: 781 280 4000 Fax: 781 280 4095

Europe/Middle East/Africa Headquarters
Progress Software Europe B.V. Schorpioenstraat 67 3067 GG Rotterdam, The Netherlands Tel: 31 10 286 5700 Fax: 31 10 286 5777

Latin American Headquarters
Progress Software Corporation, 2255 Glades Road, One Boca Place, Suite 300 E, Boca Raton, FL 33431 USA Tel: 561 998 2244 Fax: 561 998 1573

Asia/Pacific Headquarters
Progress Software Pty. Ltd., 1911 Malvern Road, Malvern East, 3145, Australia Tel: 61 39 885 0544 Fax: 61 39 885 9473

PROGRESS SOFTWARE

www.progress.com