
RESONATE | WHITE PAPER

*Maximum Performance with
Fine-Grained Control: A Requirement for
E-Business Applications and Services*

Whitepaper – August 1999



Table of Contents

Introduction.....	3
Fine-Grained Control Imperative for E-Business	5
Distributed Software vs. Hardware-Based Products.....	7
The Distributed Software Advantage: Maximum Performance with Fine-Grained Control	9
The Bottom Line: Demand the Performance and Control Required for E-Business Success	10

Introduction

Performance and Scalability Requirements

A critical issue for all e-business applications is the ability to rapidly scale infrastructure to accommodate large and unpredictable growths in traffic. When every resource in an infrastructure can maximize its computing power, the e-business application has the ability to accommodate the greatest loads while reducing overall system requirements. Traffic management solutions play a critical role in determining how efficiently resources are utilized because they manage and direct all traffic going into a site. It is imperative that traffic management solutions be optimized for greatest performance and scalability.

Potentially in conflict with this objective is the further requirement for advanced traffic management features which are critical for state management, security and high reliability.

Therefore, traffic management solutions must be able to achieve required performance levels without tradeoffs in functionality or flexibility.

Maximize Performance with No Compromises

Traffic management solutions designed for e-business applications provide advanced features for achieving the fine-grained, comprehensive control over traffic scheduling and routing. These features enable site architects to meet the application and flexibility requirements specified by their site design and marketing objectives.

However, the architect of an e-business site may be forced to make painful tradeoffs between high performance and a rich feature set if they choose a traffic management solution that was not designed to deliver both high performance and advanced features simultaneously. This tradeoff may not be realized until after the initial implementation, when faster performance is achieved only by sacrificing critical features.

This performance vs. functionality limitation occurs when the architecture of the traffic management solution cannot deliver the processing power that the sophisticated, CPU intensive, features need. Hardware-based traffic management appliances typically suffer from this limitation.

Since the hardware-based traffic management device must do all the work itself, as traffic levels increase, there are fewer remaining CPU cycles available to perform the advanced traffic management features. Furthermore, since most hardware devices today have only a single CPU, the power available for processing the features is quite limited.

Fine-Grained Control with Unbounded Scalability and Performance

Resonate's proven, patented software solution is designed for e-business applications and services that require superior performance *while retaining all advanced traffic management functions*. Resonate's distributed software architecture enables the performance and throughput of a site to actually *increase* as servers are added to the site. This enables site architects to realize all the benefits of the advanced traffic management features at performance levels unmatched by any alternative products.

New advanced traffic management features offer powerful capabilities that extend well beyond basic load balancing server availability and performance functions for better, more flexible utilization of existing resources. For example, "delayed-binding" features such as SSL Session ID, shopping cart, and Cookie-based persistence, allow more control and enable architects to "bind" users to a specific server when state information about a particular user is stored on that server. Also, intelligent scheduling features, such as URL and cookie parsing, enable architects to separate traffic based upon the content being requested. E-businesses count on these types of functions to empower more scalable and more flexible site architectures.

Recently, several vendors have announced traffic management products that feature "direct path return" as a method for increasing performance. Direct path return is a technology implementation that allows outgoing data to be sent directly to a user from a content server, without having to route through a traffic management system. While direct path return improves the performance of the system, recent products announced do not allow the use of delayed-binding features to be used in conjunction with direct path return. Resonate Central Dispatch is the only traffic management solution which allows both direct path return and delayed binding features to be used simultaneously.

While Central Dispatch was built from the start with direct path return included, these recent direct path return announcements are an indication of the market recognition and demand for high performance given unpredictable, increasing traffic. However, it is unthinkable to expect e-business operations to choose between performance and features/functionality. Only Resonate's distributed architecture allows both direct path return and feature rich scheduling and routing – the essential combination for managing sophisticated E-business offerings.

Never Settle For Less

Benchmarking of Resonate Central Dispatch resulted in impressive throughput in excess of 750 Megabits per second from a single scheduler, *with all delayed binding features turned on*. Based upon these results, total single site throughput for a Central Dispatch site utilizing its Multiple Scheduler capability can exceed an incredible 12 Gigabits per second. Resonate's Multiple Scheduler capability is unique in the industry, allowing a single site to scale rapidly while maintaining the same network topology and a single management view of the site. Resonate's throughput performance exceeds other traffic management products by 150% – 200%, making it the premier choice for maximum throughput performance in the industry.

Fine-Grained Control Imperative for E-Business

What is Delayed-Binding?

A traffic management solution is deployed to distribute incoming TCP connections to a group of servers that appear as a single, virtual IP address (VIP). Typically, these products allow different criteria to be used when determining which of the servers will be chosen to satisfy a particular client request. Once that server selection decision has been made, all subsequent traffic associated with that request is directed to that server. This process is called **binding** a client to a server.

In the simplest implementation of a traffic management product, the binding decision is made as soon as the scheduler receives the initial SYN packet from the client. After making that decision, the scheduler simply forwards that SYN to the selected server, which will satisfy that request. With this approach, the scheduler is able to use only a limited amount of information about that request in order to select a server to bind to (the information contained in the SYN packet).

The most useful data contained in the SYN packet is the Source (Client) and Destination IP address, and the Source and Destination port number. Because the binding of the request to a server happens upon the receipt of this SYN packet, an implementation of this nature would be described as an “immediate-binding” technique.

A number of vendors have introduced features which take into account information at the application layer (e.g. HTTP, SSL session etc.) when making the decision about which server the client should bind to. This type of decision can only be made after the initial three-way TCP connection handshake, because the necessary information is not immediately available. Because the binding decision happens **after** the completion of the TCP handshake (rather than upon receipt of the initial SYN packet from a client system), this type of binding is generally called **delayed-binding**. The class of features that require this type of delayed binding are typically known as delayed-binding features.

In the following diagram, the sequence in which immediate-binding and delayed-binding traffic management systems select a server is shown.

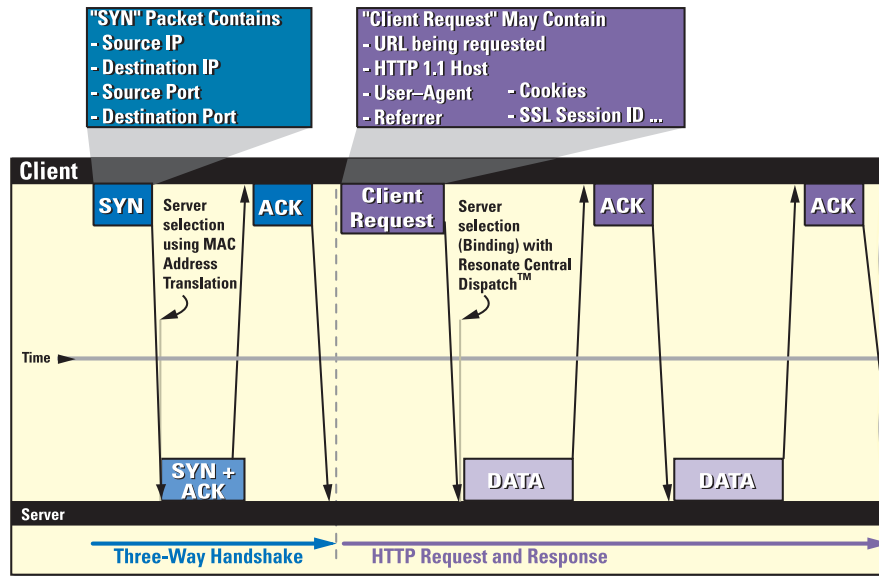


Figure 1: TCP Connection in Detail

With immediate-binding, the selection is performed immediately after the SYN packet is received from the client. Because of this, the only information, which can be used in making that server selection is the information contained in the SYN packet, which is the Source and Destination IP in addition to the Source and Destination Port.

With delayed-binding implementations, the three way handshake is performed between the client and the load balancing solution, at which time the client sends the actual HTTP request. Delayed-binding solutions typically wait for the receipt of at least one packet, which in the example of HTTP may contain the URL, Cookies, HTTP 1.1 Host and other information. Once this information is available to the traffic management solution, it can utilize any of this information to make a more granular or precise server selection.

The Value of Delayed-Binding

Delayed-binding features are valuable because they allow more specific, individualized and intelligent server selection for any particular request. With delayed-binding, it is possible to gain additional insight into the request being made. In the case of HTTP, it allows information such as the requested URL and Cookies to be used as the criteria in the server selection process.

For many sites, ensuring persistence is not only a requirement for SSL/HTTPS, but also for the HTTP protocol. This can be accomplished for both protocols using different delayed-binding techniques.

In the case of HTTP, the traffic management solution can look at a Cookie, (for example a User ID Cookie) which the client sends with each request, and remember to continue to send users with the same cookie value to the same server for each request.

With SSL, the SSL Session-ID can also be intercepted using delayed-binding techniques to ensure that all requests sharing a Session-ID are forwarded to the same server.

Distributed Software vs. Hardware-Based Products

In order to achieve maximum performance with delayed binding features turned on, a distributed architecture must be employed. The chart below compares Central Dispatch with hardware-based product traffic management techniques (Network Address Translation and MAC Address Translation), assessing the advanced functionality capabilities and the system scalability given the workload that each product performs and how the workload is distributed. With Central Dispatch, workload is distributed across all scheduler and content servers within the site, as compared to the equivalent workload that a single hardware device must perform.

Comparing Traffic Management Techniques			
	Distributed Software	NAT-Based Solutions	MAT-Based Solutions
Features Offered	High	High	Low
Workload Capacity	High Workload distributed across servers leveraging multiple servers for processing power. No effect on traffic as features are deployed.	Low Single point device handles all workload. Limited to the processing power in that device. Traffic and features create contention for power.	Medium Single point device handles all workload. Limited to the processing power in that device. Lack of features means lower workload.

Figure 2: Functionality and workload

Network Address Translation (NAT) Technique

NAT (Network Address Translation) based techniques perform all of the work necessary to accomplish the delayed-binding functionality¹ on the single device. This work includes tasks for managing packets traveling in both the inbound and outbound directions. Requiring a single device to handle the entire workload and perform all the activity necessary causes *a decrease in throughput and performance* as the site scales.

The Limitations of NAT

- Not scalable
- Workload bottleneck through single device

Functionality is still available for advanced traffic management features, but effort

¹ The tasks associated with translating IP addresses for correct client communication.

can not be offloaded or shared. The e-business using NAT retains features at the cost of performance degradation.

MAC Address Translation (MAT) Technique

One of the primary reasons many vendors introduced MAC Address Translation (MAT) techniques was because of the significant amount of work that must be performed by traffic management devices. MAT modes offer higher performance by allowing servers to respond directly back to clients, with outgoing packets bypassing the traffic management device entirely. In other words, a direct path return. In this mode, the traffic management device has a significantly lighter workload. *However, since MAT requires immediate-binding of the incoming request, all of the delayed-binding features are unavailable.* The e-business using MAT retains high performance at the cost of features.

The Limitations of MAT

- No delayed binding features
- All servers on the same subnet
- No awareness of server health
- No awareness of server responsiveness/speed

No Delayed Binding Features

Delayed-binding features are not possible with a MAT based solution. This is because with MAT, server selection must be performed at the time the initial SYN packet from the client arrives. Since delayed-binding features are enabled by information that arrives after the SYN, no delayed-binding features are possible. Delayed-binding features that are unavailable with MAT include scheduling based on URL contents, Cookies, HTTP Headers, or an SSL Session ID.

In spite of the fact that a particular vendor might support the use of any of these items for choosing a server in their normal operating mode, they cannot offer these features while the direct path return feature is used.

All Servers on the Same Subnet

MAT techniques require that the content servers be present on the same subnet as the traffic management device. If the servers are not on the same subnet, the device will not be able to modify the destination MAC address. Since the destination MAC address for a device on a different subnet will be the router's MAC address, the request will never get to the destination server.

No Awareness of Server Health

In addition, since the hardware-based traffic management solution is unable to see the response from the server, it needs to rely on active probing of the servers at some interval to ensure that those devices and services (e.g. web server on port 80) are healthy.

No Awareness of Server Responsiveness/Speed

Hardware-based traffic management solutions rely on seeing the response to initial connection setup (the SYN ACK packet) from the servers in order to estimate server responsiveness/speed. This is accomplished by measuring the time between the sending of the SYN packet (by the device to the content server) and the receipt of the SYN ACK packet from the content server.

With MAT-based techniques, the traffic management device is no longer establishing the connection to the server. It is therefore not able to gather any information on server responsiveness. It can only measure the rate at which the inbound packets arrive from the clients. Since these rates can be highly variable based upon client location, this information is not typically useful. This leaves the traffic management products using MAT with only the number of open connections on each server to use as a basis for choosing a server. Although some hardware products may use the responsiveness of their health checks as a metric, these values are inaccurate and cannot be measured frequently enough to be a comprehensive, timely metric.

The Distributed Software Advantage: Maximum Performance with Fine-Grained Control

With a distributed software solution like Central Dispatch, software is installed on the content servers as well as the scheduling system. This allows Central Dispatch to delegate responsibility for certain functions to the content servers, freeing the scheduler to continue handling new requests and leveraging the power of the content server's CPU. Thus, Central Dispatch's distributed software architecture is *designed* to offload processing work to the content server systems. This architecture enables the total processing power allocated to the management of traffic *to increase* as more servers are added, thus increasing total site throughput and performance as the site scales.

The e-business using distributed software like Resonate's Central Dispatch gets the best of both worlds: high performance combined with advanced traffic management features.

The Advantages of Distributed Software

The performance of a single Central Dispatch site has been benchmarked to determine its maximum throughput and hits per second. In these tests, the aggregate throughput of a site utilizing a single Central Dispatch scheduler exceeded 750 Megabits per second. This benchmark was performed fully utilizing delayed binding features which allow the examination of the URL being requested and directing certain URL patterns to specific servers. In addition to this impressive single

The Resonate Advantage

- Unmatched performance
- Leading edge features
- Comprehensive server monitoring
- Deployment flexibility
- Proven track record of leading customers

scheduler performance, Central Dispatch's capability to add additional scheduler nodes to a single site for increased performance capacity (while maintaining a consolidated management view), traffic throughput can exceed *12 Gigabits per second and 56,000 hits per second from a single Central Dispatch site.*

This performance is unmatched in the industry with absolutely no feature loss. Resonate's throughput performance exceeds other traffic management products by 150% – 200%.

In this benchmark series, a single Central Dispatch scheduler was configured on an Intel Pentium II, 450 Mhz Windows NT server. Throughput results reflect the total amount of traffic through the Central Dispatch-supported site, and hits per second represent the number of hits or information requests processed by the scheduler in one second. Central Dispatch supported platforms also include Solaris, AIX, and Linux. Running Central Dispatch on higher performing CPUs will deliver even higher performance results.

These results emphasize the fact that Resonate Central Dispatch can achieve incredible performance through direct path return without compromising delayed binding features. Central Dispatch also allows complete flexibility in the placement of servers. This means that servers can be on the same subnet or behind a router; nodes need only be reachable from the scheduler through IP.

The Bottom Line: Demand the Performance and Control Required for E-Business Success

It is obvious that successful, fast-growing e-businesses require advanced features to ensure uptime and availability. At the same time, throughput requirements for these same e-businesses are growing astronomically. Hardware solutions are responding to this growing traffic demand by offering products that “check the box” on direct path return, but force their customers to sacrifice all delayed binding features.

Resonate alone enables customers to achieve unmatched performance without compromising leading-edge e-business features.

Resonate is a registered trademark, the Resonate logo, Keeping E-Business Open for Business, Resonate Central Dispatch, Resonate Global Dispatch, and Resonate Commander are trademarks of Resonate, Inc. All other trademarks are the property of their respective owners. Copyright © 1998, 1999 Resonate, Inc. 09/99 WP902.