

---

# RESONATE | TECHNICAL WHITE PAPER

## *TCP Connection Hop*

---

April 2001

A core technology invention of Resonate which provides a powerful, flexible architecture for intelligent load balancing and traffic scheduling is called the “TCP Connection Hop.” The TCP Connection Hop capability is implemented through a powerful, efficient, and high performance software component called Resonate Exchange Protocol™ (RXP). RXP, which is installed on every machine in a Central Dispatch cluster, serves multiple functions depending upon the designation of the machine on which it is running.



## *Table of Contents*

Resonate Active Service Level Management Solution .....	3
RXP Functions .....	3
Functional Overview .....	4
Sample Data Flow Described .....	5
Key Advantages of the TCP Connection Hop .....	6
Summary .....	7

## *Resonate Active Service Level Management Solution*

Resonate delivers the first active service level management solutions for e-business applications. By combining real-time monitoring with policy-based control, Resonate Commander takes immediate corrective action to ensure maximized service levels of an e-business application. Resonate Commander ensures service levels of e-business applications, providing high availability and performance while leveraging existing infrastructure investments.

Resonate's Unique patented TCP Connection Hop is a core technology invention of Resonate which provides a powerful, flexible architecture for intelligent load balancing and traffic scheduling. It is an integral part of Resonate Commander as well as Resonate Central Dispatch and Resonate Global Dispatch.

The TCP Connection Hop capability is implemented through a powerful, efficient, and high performance software component called Resonate Exchange Protocol™ (RXP). RXP, which is installed on every machine in a Central Dispatch cluster, serves multiple functions depending upon the designation of the machine on which it is running.

### *RXP Functions*

1. **Primary Scheduler.** If the machine is designated as a primary scheduler, RXP functions as the master scheduler for all incoming requests to a cluster of servers.
2. **Server Node.** On all other (content) server nodes the RXP component serves as the recipient of traffic directed to it by the scheduler, ensuring that connections are efficiently and reliably made with the targeted server, and also that content sent to clients can go directly to clients without passing back through the scheduler.
3. **Backup Scheduler.** On a server designated as the backup scheduler, RXP also waits to take over the scheduling functions should the primary server fail. Note: A server can be designated to serve both scheduling and content serving functions. Because the same RXP software component is installed onto each machine, designations can change on the fly without disrupting service to clients.

Resonate traffic management service agents, which are not discussed in detail in this paper, are also installed onto each machine in order to communicate to RXP the health (CPU load, availability, network latency) of each server in the cluster.

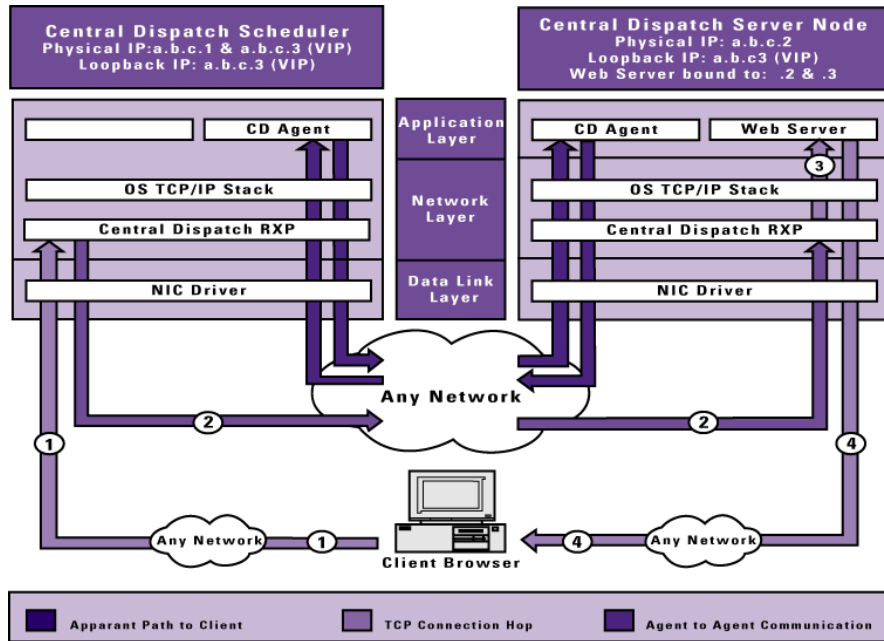
The TCP Connection Hop is an invention unique to Resonate and Resonate has been granted a U.S. Patent for this invention.

## *Functional Overview*

As the primary scheduler, RXP performs the following actions for each TCP request:

1. Receives the connection request (SYN packet) from requesting client.
2. Establishes a TCP connection in order to determine appropriate load balancing and scheduling decision.
3. Checks a table of customer specified scheduling rules to determine which server and port to send the request to. A resource-based scheduling is enabled, which can parse the URL to determine content being requested
4. If more than one server is available to serve the content, picks the least loaded server based upon load balancing metrics
5. “Hops” or transfers the TCP connection, using a Resonate protocol, to the RXP component on the targeted server.
6. Manages acknowledgements (ACK packets) from clients and persistent session information for ongoing requests.

## Sample Data Flow Described



1. The client browser makes an HTTP request to the virtual IP (VIP) address a.b.c.3. An IP packet is created by the client machine and sent through the network to the machine physically bound to the a.b.c.3 address, which is the active scheduler, in this case a.b.c.1.
2. The IP packet makes its way up the network stack until it reaches RXP. RXP recognizes that this packet is destined to a VIP, so it decides it must schedule and load balance that packet, as opposed to just letting it pass through. Based on the VIP, port and Layer 7 application information (e.g. URL, cookies, header value), RXP is aware of what nodes it can schedule that packet to. Then, based on real-time information about those nodes from the agents, and the scheduling policy, RXP decides to schedule the packet to the node at a.b.c.2. RXP then encapsulates the entire IP packet, as it was sent from the client, in an RXP packet, and sends it to a.b.c.2.

RXP on the server node gets this packet, sees that it is an RXP packet and strips off the RXP header (revealing the original IP packet), and then passes that packet up the network stack where it reached the web server. Since the Web server is bound to the VIP (a.b.c.3) via the Loopback Adapter, the Web server takes the packet and processes it as if it came directly from the client.

3. The Web server replies directly to the client, without having to pass back through the scheduler. The packet contains the destination address of the VIP, as expected by the client receiving the packet.

## *Key Advantages of the TCP Connection Hop*

The TCP Connection Hop technology of Resonate has many advantages, as listed below. It is a non-intrusive, efficient method for scheduling and load balancing all TCP-based traffic. The existence of any software product is completely transparent to the client at all times, and the triangular data flow is optimized for the asymmetric nature of Internet traffic (heavy server-to-client traffic, light client-to-server traffic).

1. **No Modifications to Native Kernel.** The original packet from the client is not modified, re-addressed, spoofed etc. It is moved from one machine to another. RXP is a Kernel Module, sitting between the NIC driver and the server's native TCP/IP stack; the native kernel or network software is not "hacked into" or modified in any way. RXP is a loaded kernel module, which can be loaded and unloaded easily.
2. **No Performance Hit From Using Native TCP/IP Stack.** Because RXP is below the TCP Stack, the scheduler machine's native TCP/IP stack never processes connections to the VIP, unless RXP determines that the scheduler machine is the least loaded and should directly serve that client's request.
3. **Efficient, High Performance Design.** The whole connection hop occurs at the network layer, is very efficient, and has very little overhead associated with it. This architecture is proven and scales to handle traffic to some of the largest Internet Web sites. Since RXP is acting like a traffic cop, a connection only passes through it for a few milliseconds. All of the scheduling rules and agent node statistics are in memory and RXP has direct access to them so it can schedule a packet on the fly.
4. **Ability to Cross Subnets.** Because RXP sits on both the scheduler and (content) server nodes it is able to perform the TCP Connection Hop across subnets. This eliminates the need for network redesign that many hardware solutions require and eases system and network management (changes in network topology) as sites grow, services are added, or servers/networks are reconfigured.
5. **Triangular Data Flow.** Packets go directly from content servers to the client without having to pass through the scheduler. In hardware solutions, packets must pass back through the hardware box and network address translation performed on each packet.
6. **Resource-Based Scheduling.** Because RXP performs delayed binding of the TCP connection, Layer 7 application information passed from the client can be parsed to provide more control over which server(s) the connection should be passed to. This provides a powerful mechanism for optimizing the allocation of server resources and easing content replication issues by segregating the content on large server farms.

7. **Flexible TCP Service Scheduling.** Multiple TCP services (and VIPs) can be scheduled to any set of available servers, in an overlapping or segregated design. For example, HTTP traffic can be scheduled to servers A, B, and C, while FTP traffic can be scheduled to servers C and D (with C having both HTTP and FTP traffic). Alternatively, HTTP traffic (servers A and B) can be segregated from FTP traffic (servers C and D).
8. **Real-Time Server Designation.** At any time, any server can be designated as either the primary or backup schedulers. Content servers can be enabled or disabled (e.g., for maintenance) on the fly without disrupting service.

## Summary

Resonate's patented TCP Connection Hop provides an efficient, reliable, high performance and flexible mechanism for intelligent scheduling and load balancing of TCP-based traffic. Operating at the network level, the scheduling function determines on-the-fly (before entering the native TCP/IP stack of a server) where each request should go. The receiving function of RXP ensures that the connection is handled reliably and that the client receives the requested content. Through the TCP Connection Hop performance is optimized by protecting the source behind a single VIP address.

These functions and capabilities enable Resonate to provide the most efficient load balancing and scheduling flexibility. This enables Resonate customers to support the greatest range of expected network topologies, TCP services, content and server configurations, both today and in the future.

Resonate and the Resonate logo are registered trademarks, Resonate Commander, Resonate Central Dispatch, and Resonate Global Dispatch are trademarks of Resonate, Inc. All other trademarks are the property of their respective owners. Copyright © 2002 Resonate, Inc. 06/01 TWP010-mod 8/02.

