



## White Paper

# Post-Deployment Rapid Resolution of Performance Problems Using Cisco Application Analysis Solution

**Optimizing the performance of business applications yields greater efficiency, productivity, and profitability. However, the increasing sophistication of today's applications and networks often makes it difficult to diagnose and repair problems when they arise. For businesses to maintain a competitive edge and ensure integrity with customers and vendors, IT professionals require tools to predict, avoid and mitigate network and application disruptions.**

## SUMMARY

Cisco® Application Analysis Solution (AAS) allows you to visualize, diagnose, and perform high-level predictive analysis of an application's performance. Useful for both pre-deployment analysis and for troubleshooting existing applications, this tool helps IT groups to analyze and optimize the interactions between the network, servers, and applications through offline modeling. It bases analyses on actual packet traces collected from the production network.

Cisco AAS is central to the Cisco Network Application Performance Analysis (NAPA) solution, which is a comprehensive set of tools and services that provides information about application and network performance. More information on the Cisco NAPA Solution can be found at <http://www.cisco.com/go/napas>.

This document discusses the challenges and workflow details for post-deployment troubleshooting of existing applications. In a post-deployment scenario, Cisco AAS assists in troubleshooting the performance of existing applications by analyzing actual application transactions from a production environment and determining if performance degradations are caused by a poorly written application or a limitation of the network infrastructure.

## POST-DEPLOYMENT CHALLENGE—TROUBLESHOOTING AN EXISTING APPLICATION

Employees in a remote location of a large company access the application server and Oracle server located in the company's headquarters. The connection is fast—they use a high-bandwidth, low-latency MAN. The measured latency between the application server and Oracle server is less than 0.9 ms. However users complain that the access is slow and that it takes more than 10 seconds to query the database.

To analyze the user's experience, trace captures are taken simultaneously at the client and application server using Cisco NAM or Cisco AAS capture agents. Cisco AAS intelligently collects and merges captures files containing data from critical points in the network. Both captures are filtered by host name to isolate the problematic application transaction.

## POST-DEPLOYMENT SOLUTION WORKFLOW

Cisco AAS rapidly illustrates whether performance problems are caused by network or application issues. With Cisco AAS, you can parse an application trace, analyze it at the application-message and network-packet levels, and graphically detail their interdependency. You can identify the root cause of a delay by breaking down multi-tier applications into component flows and automatically determining dependencies among application messages.



Troubleshooting application performance requires the following workflow:

1. Profile application response time
  - Capture and import application data from the network
  - Visualize the application
2. Identify performance issues and determine root cause
  - Examine summary of delays
  - Diagnose bottlenecks
  - Perform quantitative analysis
  - Analyze data transaction and import application data from network
3. Document analysis
  - Create and publish a report
4. Determine steps for improving application response time

### **Profile Application Response Time**

Captured files can be opened and analyzed using the Cisco AAS Application Characterization Environment, or “ACE.” ACE is an enabling technology for visualizing and analyzing application performance problems. Cisco AAS reads packet traces that can be captured from a Cisco NAM or using capture agents that come with Cisco AAS. The capture agent can be installed on Windows, UNIX, and Linux operating systems.

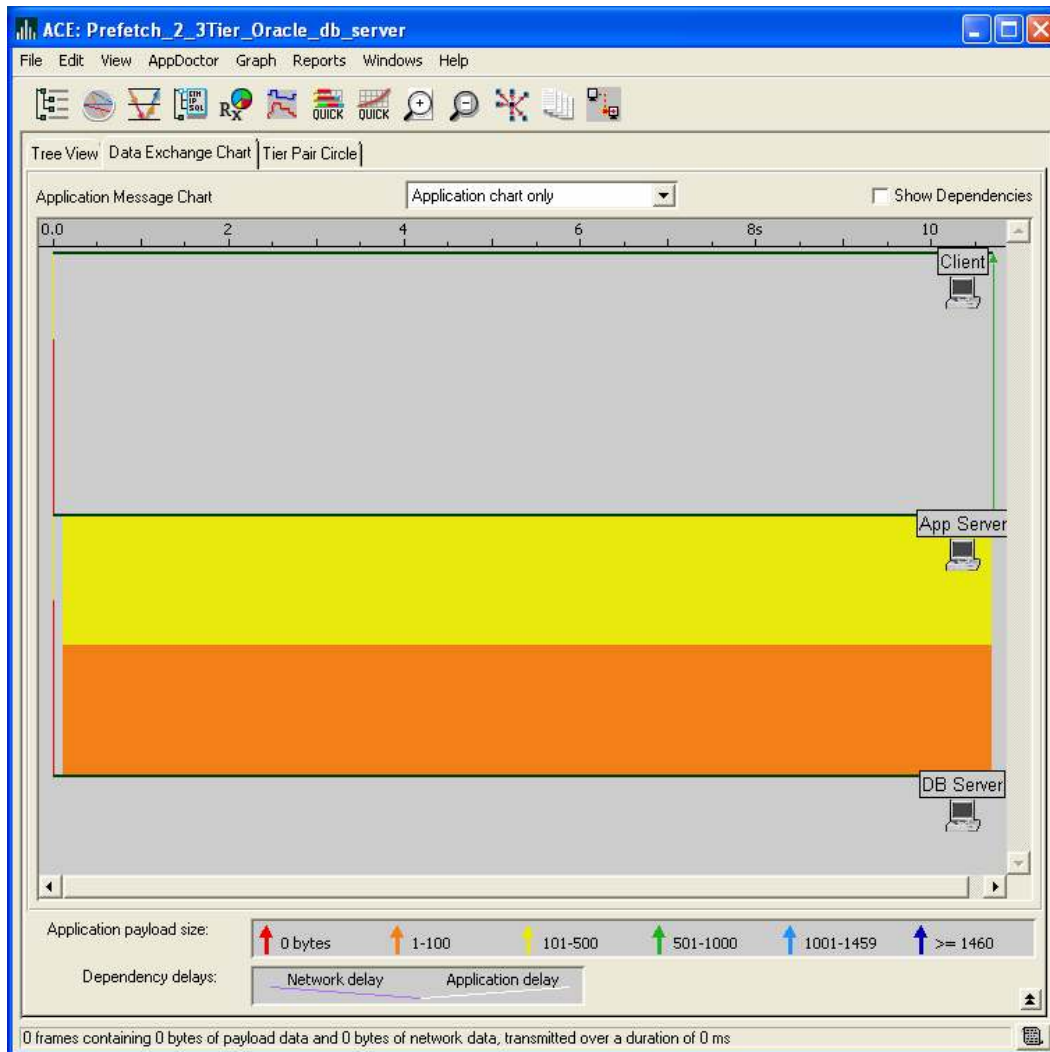
### **Capture and Import Application Data From the Network**

To evaluate the user’s experience with the application, trace captures are taken simultaneously at the client and servers. A Cisco NAM blade installed in a Cisco Catalyst® 6500 Series Switch as well as Cisco AAS agents are used for server and client application-transaction packet captures. Both types of captures are filtered by host name to isolate the application of interest; in this case the application is Oracle.

### **Visualize the Application**

The first step is to view the Data Exchange Chart, which shows the application and network packets between the tiers on a timeline. The tiers in this case are the user shown as “Client,” application server, and database server (see Figure 1).

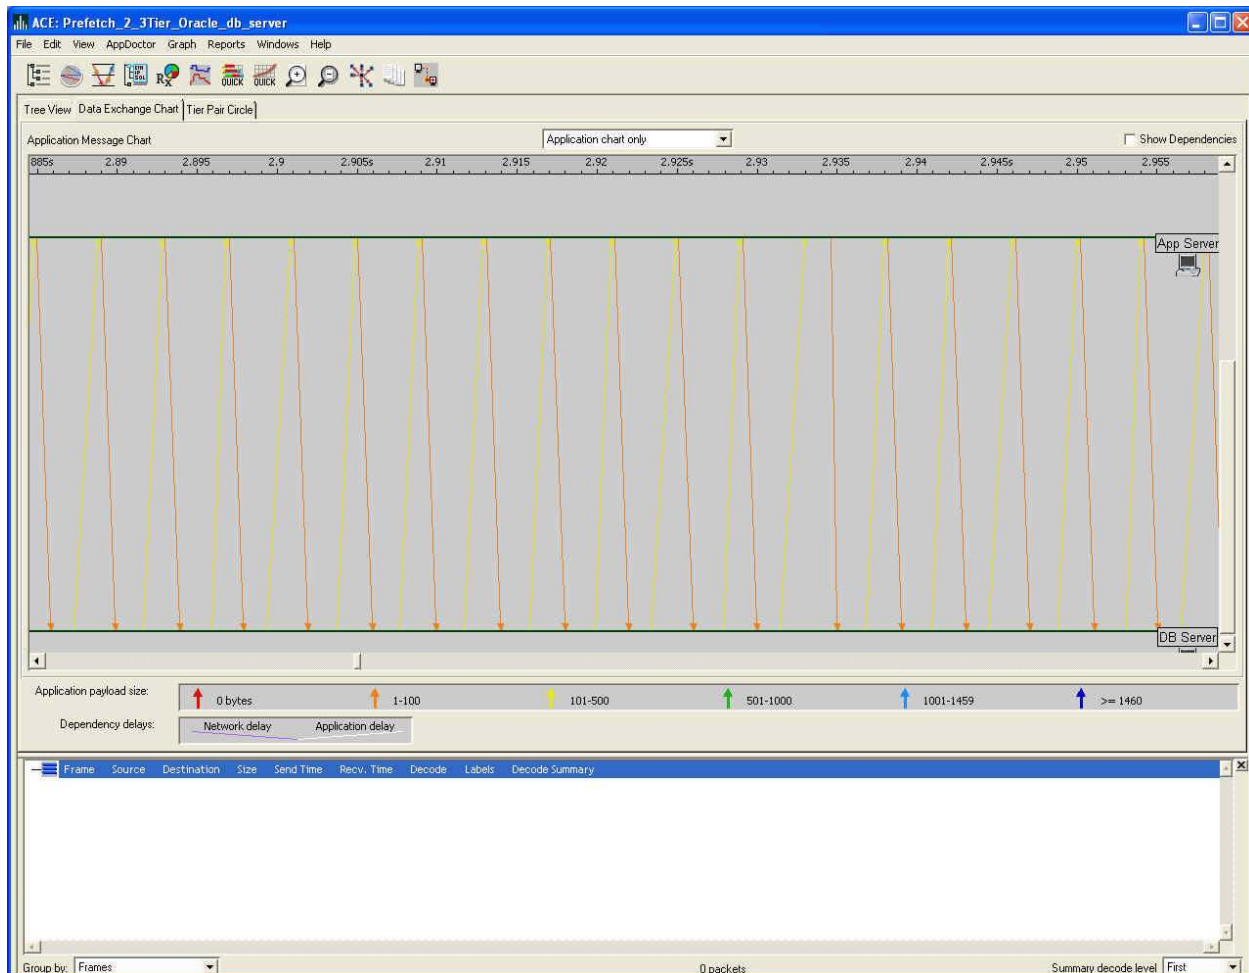
**Figure 1.** Data Exchange Chart Showing Tiers and Response Time



The Data Exchange Chart visually characterizes the application trace:

- The timeline on the top graphs shows the total response time of the application task. The response time for this application is almost 10.5 seconds.
- A horizontal line represents each application tier: client, application server, and database server.
- Message request and response times are shown between each tier.
- The messages are colored by their application payload size. The legend is shown at the bottom of the window.
- Each individual application message is represented by an arrow that starts at the source tier and ends at the destination tier (Figure 2). The arrow color represents the number of application bytes in the message. Each colored bar represents a group of messages; the amount of color represents the percentage of messages that fall into each application-payload size category.

**Figure 2.** Data Exchange Chart Shows Many Messages Between Application Server and Database Server



You can observe from the Data Exchange Chart the communication pattern between the client, application server, and database server. The client makes an initial request to the application server, initiating a series of messages between the application server and database server. Finally, a response is returned to the client. The majority of the total application response time results from communication between the application server and database server.

The chart also shows that about one-half of the messages between the application server and database server are orange, indicating the payload size is between 1 and 100 bytes per message. You can immediately conclude that the application and database are sending many small messages and the total response time is more than 10 seconds.

### Identify Performance Issues and Determine Root Cause

Visualizing the application helped determine that it is transmitting many small packets between the application server and database server and the total response time is significant. Now you can further examine how the small packet size affects the response time.

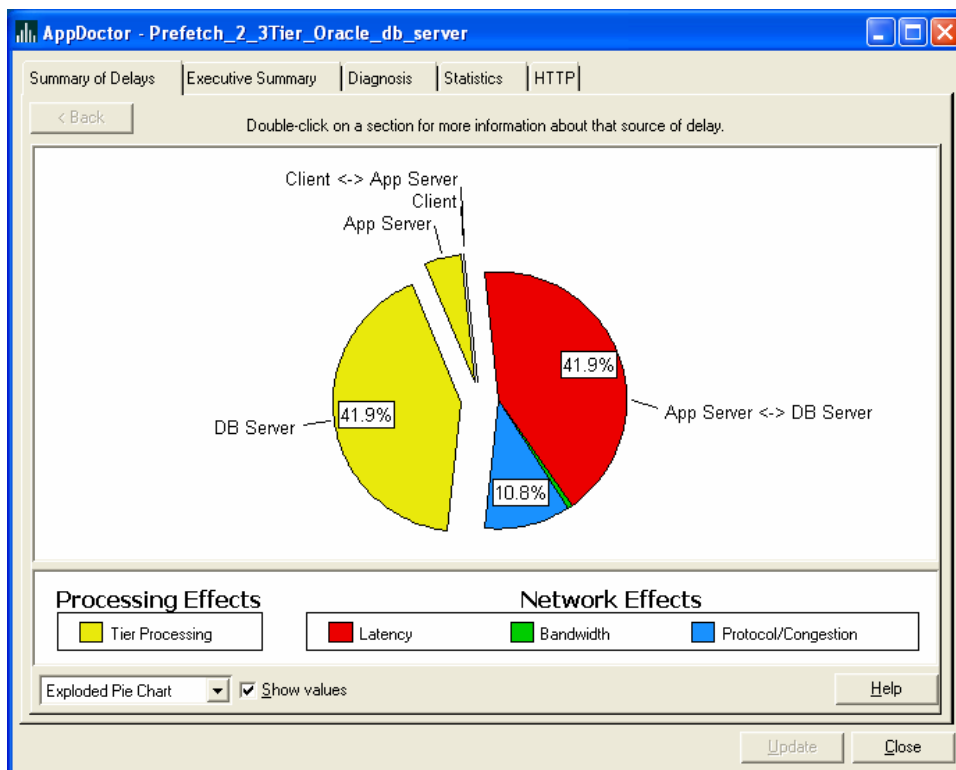
## Examine Sources of Delays

Sources of delay are summarized in convenient diagrams. Thresholds for key application statistics are used to generate informative reports that characterize problems. Cisco AAS performs these functions using its AppDoctor feature. The Summary of Delays function in AppDoctor shows you how individual delays contribute to the total application response time. AppDoctor's diagnosis then identifies the bottlenecks. Finally, more quantitative analysis can be performed looking at AppDoctor's statistics.

The Summary of Delays chart (Figure 3) breaks down the impact that the tiers and network have on the total application response time including:

- Tier processing delay—The total time it takes to process the application at each tier. This includes CPU processing time and user think time.
- Latency—Delay due to latency in the network. Latency is the time for one bit to be transmitted across the network.
- Bandwidth delay—Delay caused by the limited bandwidth of the network.
- Protocol and congestion delay—Delay as a metric of network restriction to packet flow. It may be caused by:
  - Packet queuing in the network and
  - Flow-control mechanisms imposed by network protocols. TCP, for example, has several built-in flow-control mechanisms including TCP window resizing.

**Figure 3.** Summary of Delays, Including Application Server-to-Database Server Latency



Notice in Figure 3 that a contributing factor to the application delay is network latency. It accounts for almost 42 percent of the total response time. *This is somewhat surprising because the measured latency between the application server and database server is less than 1 ms.*

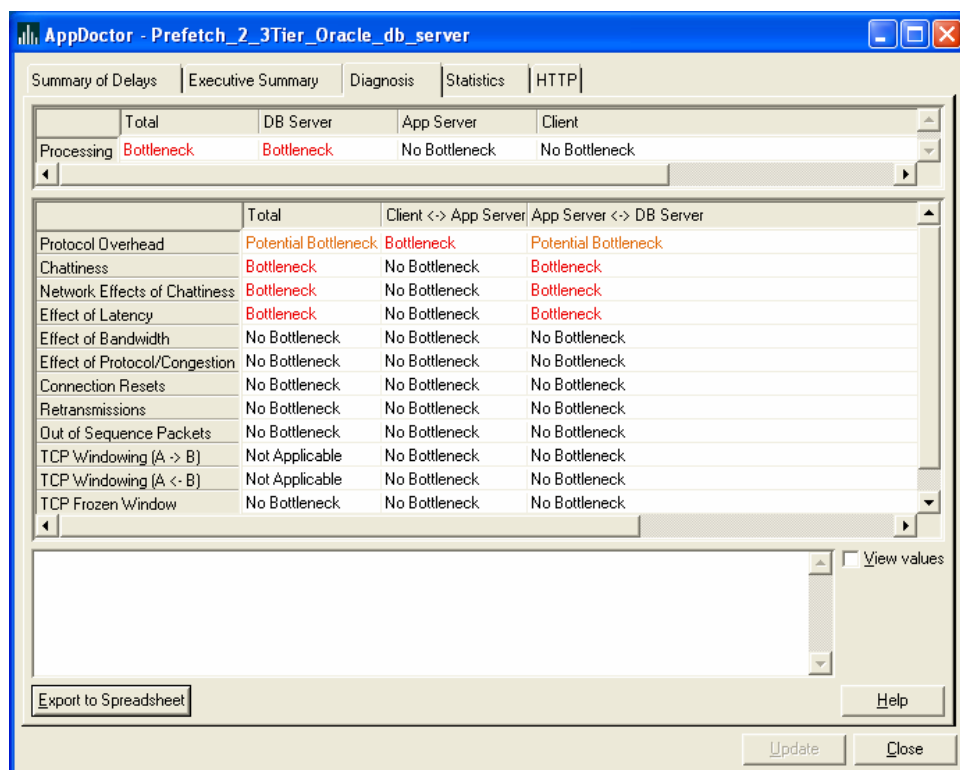
Also indicated in this chart is a significant amount of processing on the database server. The server “think” time accounts for nearly 42 percent of the total response time.

AppDoctor can be used to further analyze the processing delay on the database server.

## Diagnose Bottlenecks

The Diagnosis function in AppDoctor (Figure 4) provides a more detailed view of the potential bottlenecks affecting this transaction as well as recommendations on how to improve the application. The Diagnosis function tests the current transaction against issues that often cause performance problems for network-based applications, grouped by category. Values that exceed specific, user-configurable thresholds are marked as bottlenecks or potential bottlenecks.

**Figure 4.** Identified Bottlenecks



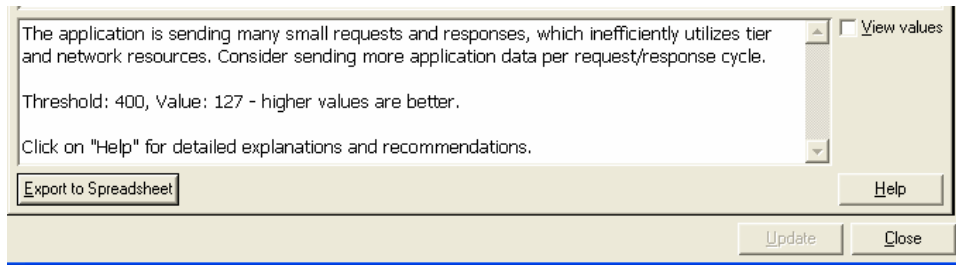
The following categories are identified as bottlenecks:

- Chattiness
- Network Effects of Chattiness
- Effect of Latency

When you select the Chattiness category, AppDoctor makes a recommendation on how to remove chattiness as a bottleneck. AppDoctor recommends sending more application data per application turn, which would decrease the impact of latency on the overall application response time (Figure 5).

If you select Effect of Latency, AppDoctor recommends moving the tiers closer together to reduce the propagation delay caused by geographical constraints.

**Figure 5.** Recommendation on How to Remove Chattiness

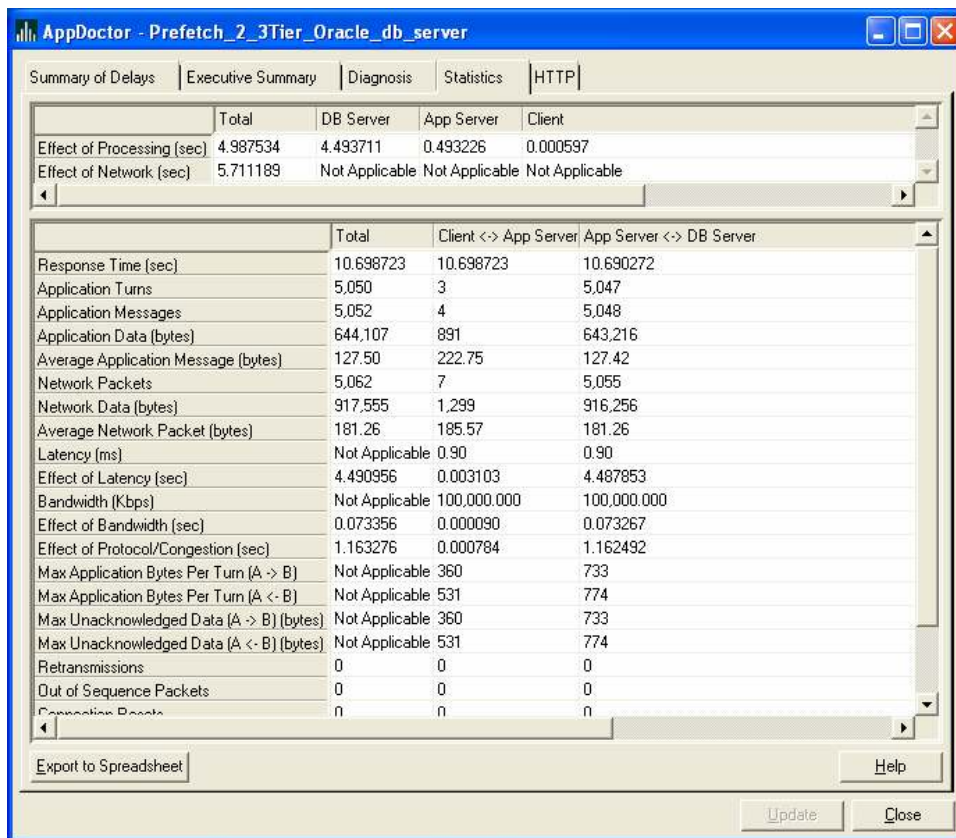


### Perform Quantitative Analysis

AppDoctor provides further details of the application transaction through summary statistics (Figure 6) including:

- Amount of time each tier spent processing the transaction
- Number of messages sent between application tiers
- Amount of data sent between application tiers
- Average network packet size
- Amount of data loss including packet retransmissions

**Figure 6.** Statistics Used to Identify Application Transaction Problem



Notice the Response Time is 10.698 seconds, which corresponds to the value that you saw previously in the Data Exchange Chart.

Several statistics are relevant to this study. This report examines two in particular: the number of application turns and the maximum data exchanged in each turn.

Each change in direction is called an *application turn* because the application changes the direction of data flow. Applications with many application turns are generally considered chatty and are sensitive to network delay. The sensitivity occurs because each message must be received at a tier before the corresponding response can be sent; as a result, each message is affected by network latency. Notice that the application used 5047 turns to exchange 643,216 bytes of data.

This confirms the previous analysis in which chattiness and the network latency are the primary causes of the poor application response time.


Also note that the *Effect of Latency* is 4.49 seconds. Here you see that the effect of latency alone, aggravated by the 5047 applications turns, accounts for about 4.49 seconds or 42 percent of the total transaction response time. Because latency is largely a product of geographical distance and network hops, adding bandwidth will have a minimal effect on the response time. To minimize this component, you can reduce the latency on the circuit or reduce the number of application turns.

### Analyze Data Transaction and Import Application Data from Network

In addition to improving the effect of latency on the application response time, it is necessary to reduce the amount of time the database server is processing the client's request by analyzing the application messages sent between the application and database servers using the Data Exchange Chart.

Examining the captured data will allow you to see how packets are exchanged. First examine the beginning of the transaction more carefully by zooming in on traffic between 0.24 and 0.26 seconds (Figure 7).

**Figure 7.** Zoom in on Captured Data

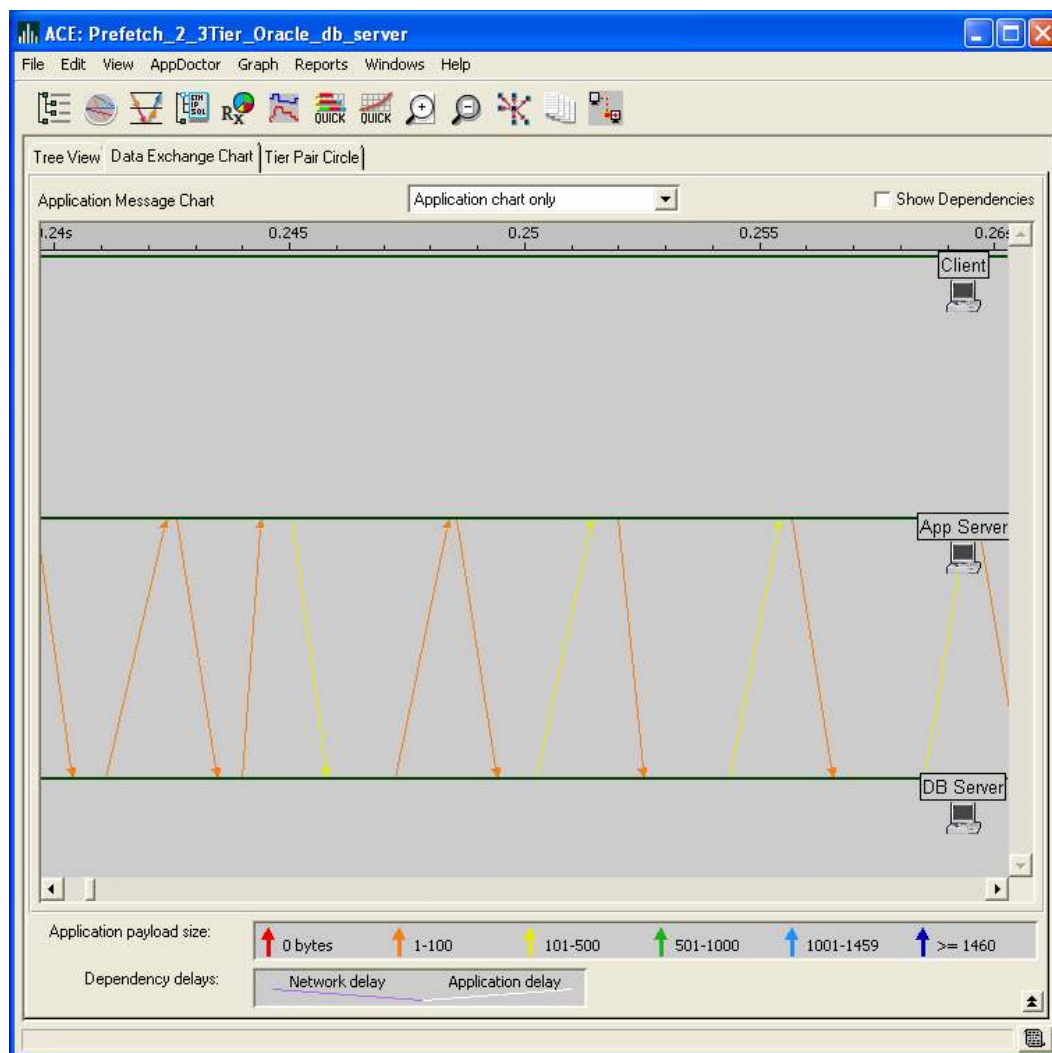


As you zoom into a group of messages in the Data Exchange Chart, the individual application messages are visualized.

Each arrow indicates the message direction between the Application Server and the Database Server (Figure 8).



**Figure 8.** Data Exchange Chart Showing Detailed Transaction



The Data Exchange Chart exhibits a simple request-response pattern. The application server sends a request and the database server receives the request, processes the request, and finally sends a response to the application server. The entire process takes approximately 9 seconds to complete.

You can view the processing delay for each message by selecting “Show Dependencies.” The processing and communication delay for each message is not significant, on average 0.1 ms. However, aggregated over the number of application turns, it becomes a significant bottleneck.

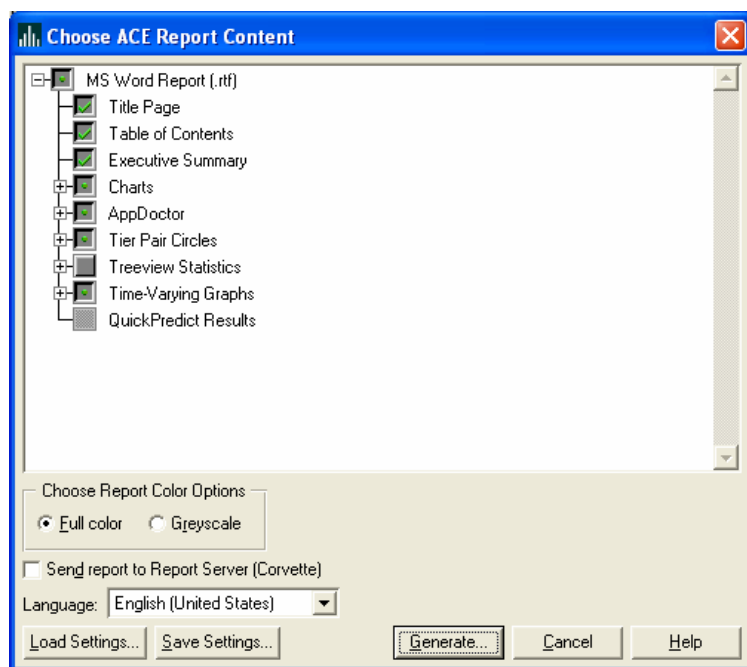
You can conclude from the analysis of the network latency and application chattiness that the application is not optimized for the current deployment. This data can be used to justify the expense of recoding the application.

## Document Analysis and Create and Publish a Report

The above analysis can be automatically captured and documented in either Rich Text Format (.rtf) or HTML format. This feature allows you to create a report that you can edit and view directly in Microsoft Word or publish to a report server.

You can specify the content to include in your report (Figure 9) and choose these report options: online (full color), print-friendly (grayscale), and English or another language.

**Figure 9.** Creating a Report



## Determine Steps for Improving Application Response Time

The post-application deployment scenario showed how Cisco AAS can be used to visualize and diagnose application performance problems. Before starting the analysis, users reported application response time was more than 10 seconds. An analysis of the communication pattern between each component of the application concluded the slow response was primarily due to application chattiness, aggravated by network latency. Although network latency between client and application server was known to be less than 1 ms, the aggregated delay was significant due to the number application turns. This information is critical to determining how to improve application performance and the end-user experience.

Cisco Application Analysis Solution is a vital tool for diagnosing and solving application performance problems. It can be used to predict application performance before deployment and to troubleshoot problems of production applications. For more information about Cisco AAS, contact your Cisco representative or visit <http://www.cisco.com/en/US/products/ps6362/index.html>.



**Corporate Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

**European Headquarters**  
Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
www.europe.cisco.com  
Tel: 31 0 20 357 1000  
Fax: 31 0 20 357 1100

**Americas Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Pacific Headquarters**  
Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
www.cisco.com  
Tel: +65 6317 7777  
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic  
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy  
Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal  
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden  
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

Printed in USA

C11-352337-00 06/06