

Oracle Web Conferencing Real-Time Collaboration System Architecture Overview

An Oracle White Paper
October 2003

Oracle Web Conferencing Real-Time Collaboration System Architecture Overview

ABSTRACT

Oracle Web Conferencing is the new product in the Oracle Collaboration Suite Release 10g. The Oracle Web Conferencing feature set is offered by the Oracle Real-Time Collaboration (RTC) system software (component of Oracle Collaboration Suite). It provides online conferences (interactive conferences, Web seminars, and large broadcasts) and the ability to record and play back any conference. In a conference, users can share all aspects of Microsoft Windows applications, cobrowse the Internet/intranet, write on a shared whiteboard, present documents, chat, conduct polls, and stream voice to others. It offers a secure environment for real-time collaboration and can be integrated via Web services into existing business applications and portals at the enterprise. The Real-Time Collaboration system can scale to any size corporation by adding more hardware.

Oracle Web Conferencing (Real-Time Collaboration system) software consists of consoles, a server, an application module, system services, and a repository. This software architecture makes it easy to deploy and manage, integrate with existing LDAP systems, work with existing firewall configurations, and uses the least amount of network bandwidth. This white paper gives an overview of the Real-Time Collaboration system software architecture.

INTRODUCTION

Oracle Web Conferencing system facilitates Real-Time collaboration between users connected by the IP networks. Real-Time Collaboration will become a widely used tool (similar to e-Mail) and platform for organizations as they realize the benefits of conducting communication over enterprise intranets and the public Internet. Every employee is connected to every other employee via the corporate intranet and the public Internet. The Oracle Real-Time Collaboration system is designed for the network computing infrastructure of today's enterprise. It does not require any reconfiguration for the enterprise firewalls. It works with any standards-based LDAP system for user management.

The product feature overview is given in the first section, the S/W architecture and the components description in the second section. The run-time execution of the various components of a Web conference and playback session is discussed in the third section. The last section covers the scalability architecture and sample deployments.

1. WEB CONFERENCING FEATURES

Oracle Web Conferencing provides support for cobrowsing, desktop sharing, document presentation, whiteboard usage, polling, chat, and voice streaming in a single Real-Time Collaboration system. Users can have public or restricted conferences. Each of the conferences can have one or more presenters. Any conference can be recorded and the recorded archive published for playbacks. Please refer to the Oracle Web Conferencing Data Sheet for more details on the above features.

2. SYSTEM SOFTWARE ARCHITECTURE

The Real-Time Collaboration system software architecture (Figure 1) consists of consoles, server, application module, system services, and the repository. Oracle Database, Oracle Application Server software components and parts of the Oracle Collaboration Suite components are intentionally left out to highlight the components of the Real-Time Collaboration system. The Web Conferencing Console and Playback Console run on any computer with the Microsoft Windows operating system. The Real-Time Collaboration Server, RTC Application Module, and RTC System Services are all software on the back-end servers, which run on mid-tiers with Oracle Application Server. The RTC Repository is the database schema in the Oracle 9i Database.

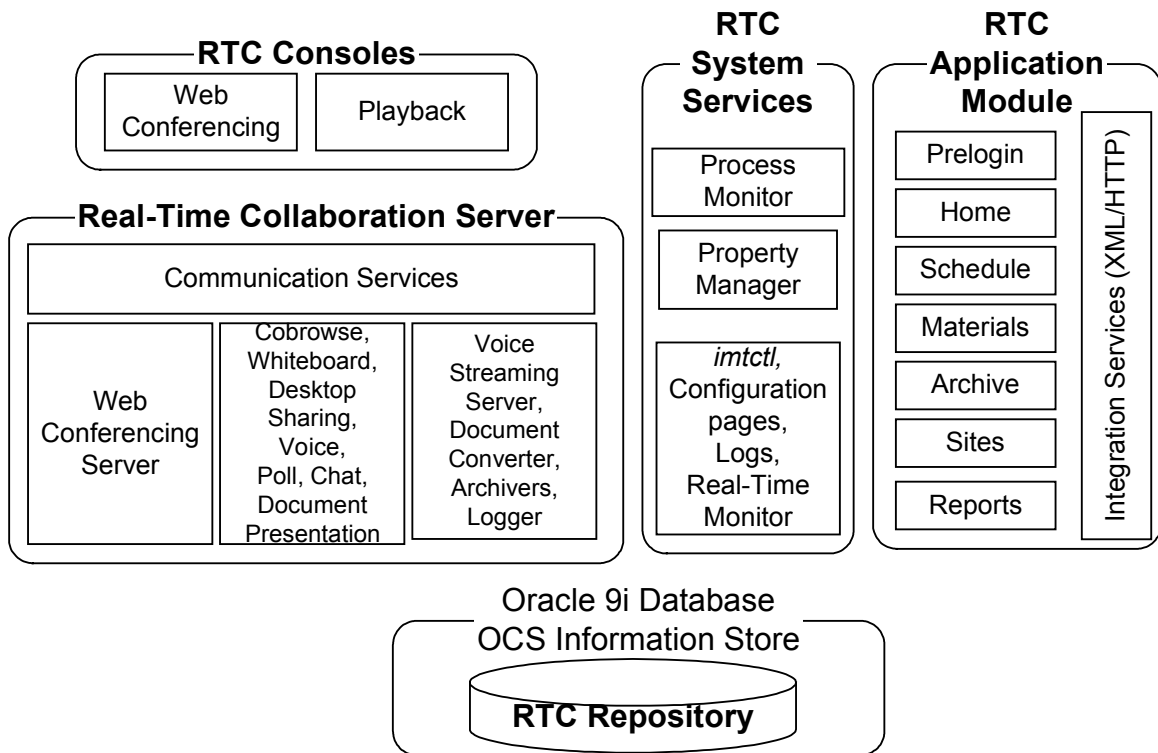


Figure 1 Real-Time Collaboration (RTC) System Software Architecture

Figure 2 illustrates the deployment of the Oracle Web Conferencing system on four nodes. It has one mid-tier, Node #1, which hosts the Real-Time Collaboration Server and the RTC Application Module software. Node #2, the RTC instance, runs the Real-Time Collaboration Voice Streaming Server and Document Conversion Server and is connected to the PSTN network via a T1/ E1 line. The RTC Repository resides on Node #3, and the Oracle Application Server Infrastructure resides on Node #4. As indicated in Figure

2, both the RTC instances are in the DMZ (Node #1, Node #2), and the RTC Repository and the Oracle Application Server Infrastructure system are in the intranet.

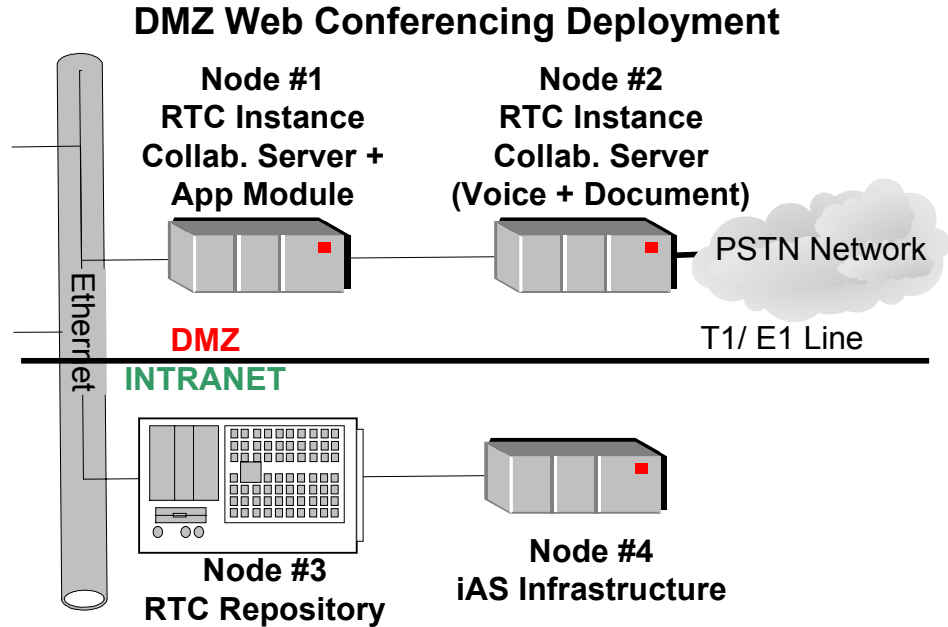


Figure 2 Basic RTC System Deployment Configuration

2.1 CONSOLES

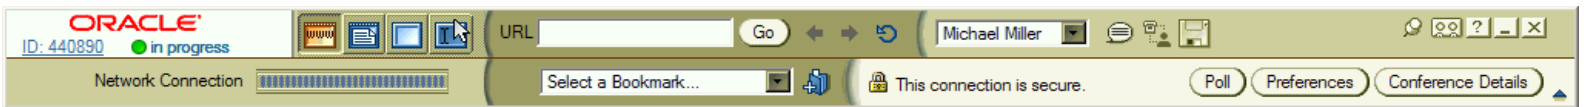


Figure 3 Web Conferencing Console

The RTC Web Conferencing Console has two taskbars, as shown in Figure 3. The first task bar row integrates the cobrowse, document presentation, whiteboard, desktop sharing, chat, voice streaming, screen shot save and recording features into a single taskbar. In the second row, there are indicators for network latency to the Web Conferencing Console from the Real-Time Collaboration Server, bookmarks, SSL/non-SSL encryption, polling, and runtime conference preferences. Network bandwidth (or

latency) is indicated using three colors: Green for best bandwidth (least latency), yellow for moderate, and red for poor bandwidth.



Figure 4 Playback Console

The RTC Playback Console (Figure 4) interprets the archives and plays the data and voice streams either directly from the RTC Application or from the local downloaded copy of the Archive. The RTC Playback Console buffers the content streamed from the RTC Application. It provides VCR-type user interface controls.

2.2 SERVER

The RTC Server is the set of the software components (see Figure 1) that work together at runtime during a Web conference. It consists of the following:

- Communication services (Mod_RTC, RTC MX),
- Web Conferencing Server and its associated extensions, which support different modes of collaboration (Cobrowse, Document Presentation, Whiteboard, Desktop Sharing, Chat, Poll, etc.),
- Voice Streaming Server and Document Conversion Server

Communication Services:

The RTC Communication Services consists of two components: mod_RTC and RTC Multiplexer (RTC MX). The mod_RTC is the Apache module that intercepts the RTC connection requests. It facilitates the connection between the Internet clients and the RTC Web Conferencing Server hosting the conference session. The direct TCP/IP connection from the Apache Web Server (Oracle HTTP Server) is handed off to the RTC MX by the mod_RTC. This way, the attendee's console in a conference has direct TCP/IP connection to the RTC_MX.

The RTC MX listens at fixed TCP/IP ports. The RTC Server can be configured with one or more RTC MX components working on a given machine. Each of the RTC MX on a given machine is connected to all of the RTC Web Conferencing Servers on that machine. Intranet attendees (RTC consoles) establish a direct TCP/IP connection to one of the RTC MXs via the TCP/IP ports exchanged when a user joins a conference. Extranet users are connected to the RTC MX via the mod_RTC. The RTC MX component acts as TCP/IP socket concentrators to the attendees and has one network

connection to the corresponding Web Conferencing Servers.

All the TCP/IP connections used by the RTC MX traverse the standard firewall configurations.

Web Conferencing Server

The RTC Web Conferencing Server is a multi-threaded server that hosts one or more Web conferences concurrently. It is an extensible, real-time collaboration server that offers Cobrowse, Document Presentation, Desktop Sharing, Chat, Voice Streaming and Polling collaboration modes. All attendees are connected via the RTC MX to the Web Conferencing Server. As the Web Conferencing Server receives new content in any of the collaboration modes active in the Web conference, it materializes the new data and broadcasts the new data to the conference attendees. The conference roster management (host, attendees, presenter) is controlled by the Web Conferencing Server.

The Real-Time Collaboration Server and Web Conferencing Console determine the network bandwidth at regular intervals based on the packet round trips between the server and consoles. The RTC system has a patented “packet collapsing” algorithm, which enables RTC consoles on a slow to poor network bandwidth to receive larger packets less frequently (several small packets collapsed into one) from the server. The Web Conferencing Server determines which packets can be dropped without adversely affecting the user experience and “repacks” several packets into one. This way it reduces the amount of data to be transmitted, providing a good conference experience for users connected via slow networks.

Voice Streaming Server and Document Conversion Server

The RTC Voice Streaming Server gets the voice stream from the PSTN network and converts it to GSM codec (13 Kbps) and sends it to the RTC Web Conferencing Server that is hosting the Web conference. Then the RTC Web Conferencing Server synchronizes the voice and data streams and broadcasts the combined stream to the attendees. The RTC Document Conversion Server converts all Microsoft Office documents to PNG and HTML to be used in Document Presentation mode. It is optional to deploy RTC voice and document services. If these services are not deployed, then the users in a conference cannot use voice streaming and Document Presentation mode.

2.3 APPLICATION MODULE

The Application Module software provides the Web user interface to the system and programmatic interface via the Web Services (XML over

HTTP). It facilitates the concepts of “sites,” whereby an enterprise can have a single shared instance across the sales, marketing, support, and training and development organizations. RTC User Management and roles (RTC User and RTC Administrator) and the mapping of Oracle Internet Directory (OID) users to RTC users are done in this component. All existing LDAP users can be integrated into the system via Oracle Internet Directory component of Oracle Collaboration Suite.

The Web pages component of the Application Module can be augmented to support newer tasks. It does page layout (using model-view-controller architecture of Oracle Application Server UIX beans), menu rendering using table-driven data and manages several Web sessions in a single J2EE container (OC4J). The rendering of Web pages is layered on JAVA business objects, which manipulate the data in the RTC Repository via JDBC. It does connection pooling and multiplexes a single database connection to process several user requests.

RTC Integration Services facilitate programmatic interfaces to the JAVA business objects. These services enable third-party desktop or Web applications to create, join, and play back conference archives without using the Web Conferencing Web pages.

2.4 SYSTEM SERVICES

The Real-Time Collaboration System is a distributed system with several software components working together to offer real-time conferences and playback services. It is the RTC system services software that enables all the disparate software components to work together, ensures 24X7 reliability and the ability to administer and manage the distributed system. RTC System Services consist of the Property Manager, Process Monitor, Availability Services, and a set of runtime services and Web pages. These system services components make it easy to configure, manage, and monitor the system.

The Property Manager component provides the facility to configure any single RTC component and its configuration parameters (e.g., listening ports, debug flags). It is easy to configure any component on any node in this distributed system using the `imctl` command line interface to the Property Manager. This property management software provides a hierarchy of components in the system. Figure 5 shows this concept of the RTC system, RTC cluster, RTC instance, and RTC components (RTC Multiplexer, RTC Web Conferencing Server, `mod_RTC`, etc.). A set of RTC Components that share the same Oracle home on a node is an RTC Instance. One or more RTC Instances in a given location belong to the RTC cluster. There could be one or more RTC clusters in an RTC system.

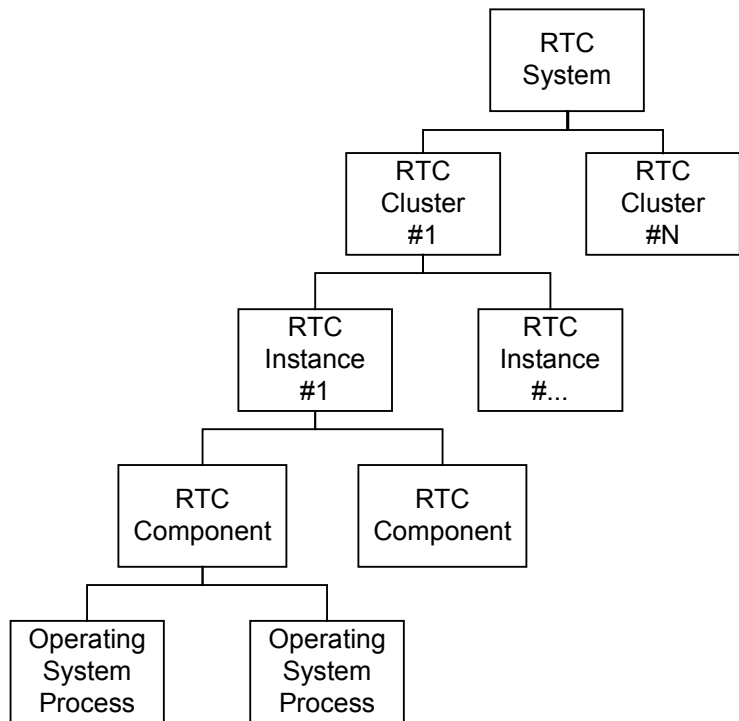


Figure 5 RTC System, Cluster, Instance, Component hierarchy

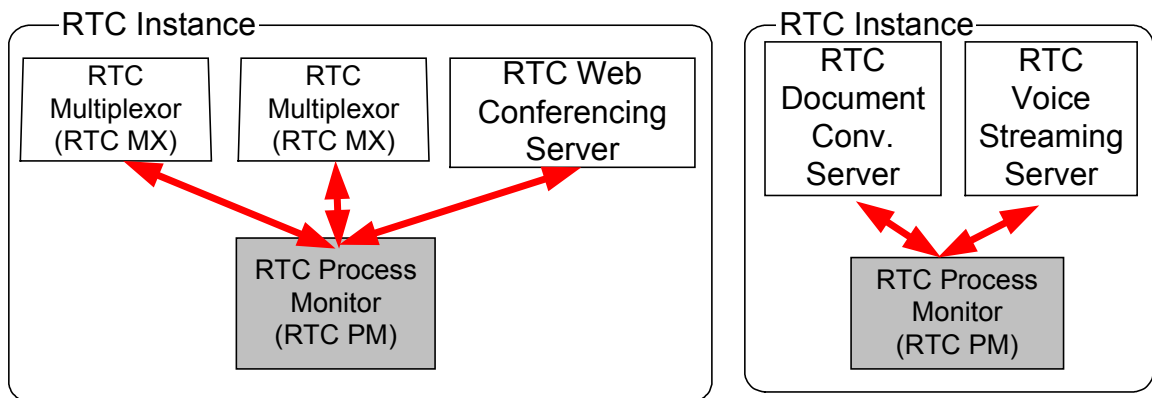


Figure 6 RTC Process Monitor Keeps RTC Components Running

The RTC Process Monitor, as shown in Figure 6, checks the health of all the RTC components in an RTC instance at regular intervals. If one of the components does not respond, the process monitor restarts it. In addition, there are sets of service availability tests provided as part of the RTC system

services. Oracle Enterprise Manager Daemons or any third-party management framework (Site Scope, HP Openview, IBM Tivoli, etc.) use these service availability tests for monitoring.

2.5 REPOSITORY

The RTC Repository is the Real-Time Collaboration system schema that takes advantage of the Oracle 9i Database. It is a set of database tables, indices, materialized views, etc., that are accessed by the RTC Application Module, RTC System Services, and the RTC Server. All the standard database tools that work with Oracle 9i databases can be used to monitor the RTC Repository.

3. RUN-TIME EXECUTION

3.1 WEB CONFERENCE

The hosts of a Web conference log in to the RTC Application Module either using the Web pages or via a 3rd party Web application integrated using RTC Integration Services. Once the host starts a Web conference, the RTC application module gets the current load on all the Web Conferencing Server processes and it then assigns the least loaded server to handle the Web conference. Following that, the RTC Web Conferencing Console gets the conference ID# and establishes a direct TCP/IP connection via one of the RTC Multiplexer components to the RTC Web Conferencing Server. If the host is using a system outside the firewall, then the mod_RTC component in Apache facilitates hand-off of the TCP/IP socket to the RTC MX. This way, the host gets a direct connection to the RTC Server. Similarly, attendees download the Web Conferencing Console and establish a direct connection to the RTC Web Conferencing Server hosting the conference as indicated in Figure 7.

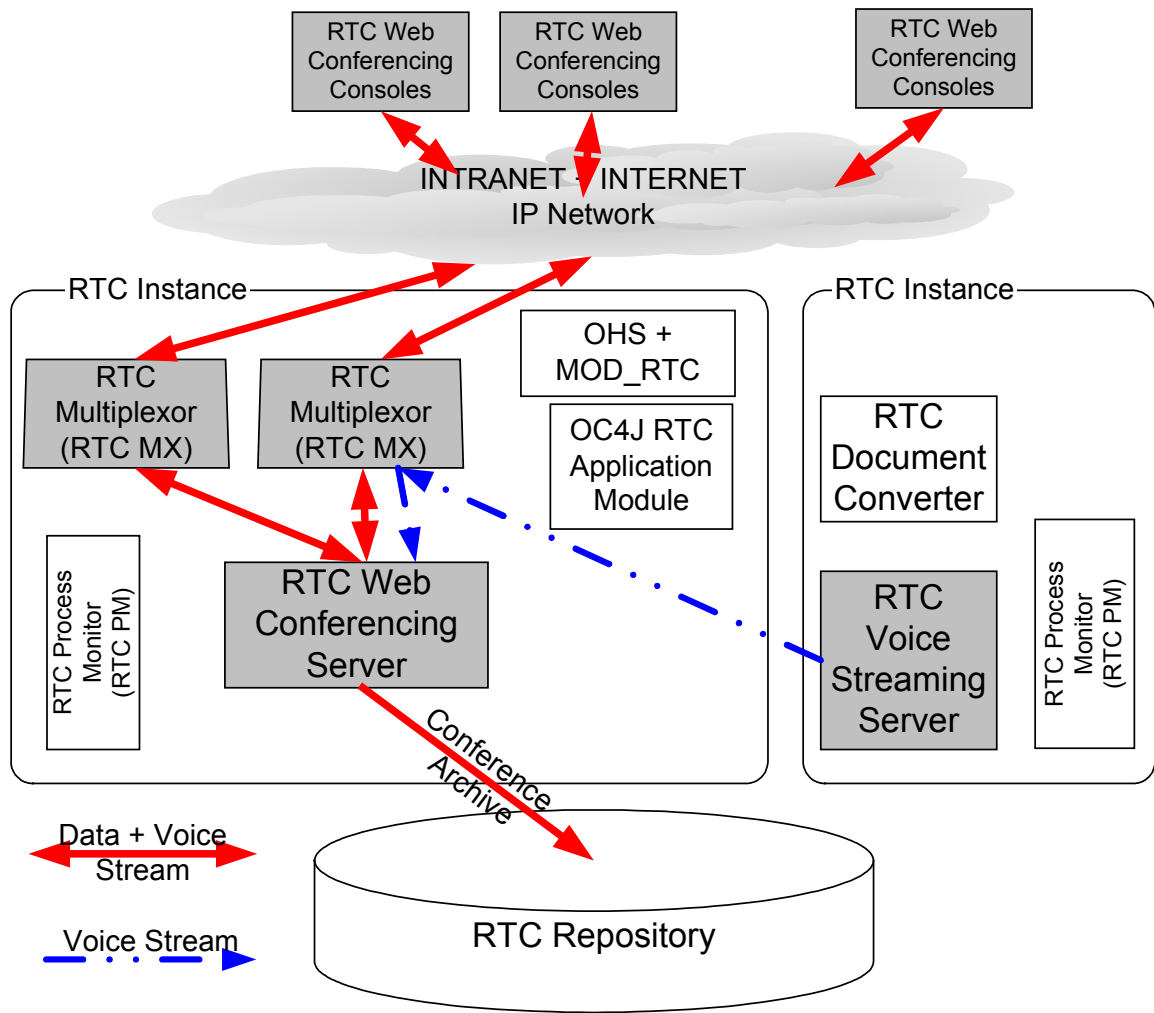


Figure 7 RTC Server Sends Data + Voice Streams to the RTC Web Conferencing Consoles

In Figure 7, as new data gets materialized, the presenter's Web Conferencing Console sends it to the Web Conferencing Server. The Web Conferencing Console optimizes the capture only to the changes and sends them to the server. If voice is streamed in the Web conference, then the RTC Web Conferencing Server gets the voice stream from the RTC Voice Streaming Server. It then synchronizes the data and voice stream and sends them to the attendees' Web Conferencing Consoles as shown in Figure 7. During the conference, the moderation (presenter and attendees) is controlled by the Web Conferencing Server. The Web Conferencing Server uses packet collapsing technology to give each user the best possible real-time conference experience. Attendees with moderate/ poor network bandwidth get less number of packets transmitted from the server (selectively). In case of recorded conferences, the Web Conferencing Server,

asynchronously writes the data back to the RTC Repository without causing delay for the conference runtime performance.

3.2 PLAYBACK

When a recorded conference is played back, the RTC Playback Console gets the archive content directly from the RTC Application Module. It buffers the content to offer a better user experience. All the security restrictions for attending the conference apply to play back archives. If any user was not allowed to attend a conference, then the conference archive is also not accessible to the user.

For users' convenience, the RTC archives can be downloaded to any PC for offline playbacks.

3.3. SECURITY, PRIVACY, AND AUDITING

Conferences and Archives

The Web Conferencing Application provides support for public and restricted conferences. In the case of public conferences, everyone who has the conference ID# can attend the conference. For restricted conferences, only invited attendees, who are also registered users of the system, are allowed. The Web Conferencing application authenticates the hosts or participants before allowing them to join the conference. Both public and restricted conferences can be password protected (enforced by the RTC Application Module). In all cases, only registered users of the system can create conferences.

SSL and Non-SSL Conferences

RTC consoles communicate with the RTC Server using a proprietary communication protocol. These messages can be encrypted using SSL for greater privacy. The default option is non-SSL conference traffic between the consoles and server.

Privacy and Auditing

The RTC system administrators cannot see the actual content of any conference unless they are invited to it. They can see only system-related information (number of attendees, Microsoft Windows platforms, network latency, etc.) at the system, instance, and conference level.

All the RTC conferences can be archived. Oracle Corporation recommends that each organization have a policy for archiving based on their business needs and the regulations by which they are bound (NASD, HIPPA, Sarbanes-Oxley, etc.).

3.3 SCALABILITY ARCHITECTURE

As shown in Figure 8, the RTC Server architecture can have multiple RTC MX components operating in an RTC instance (on a node). This allows the server to take advantage of multiple network cards and also to provide high-availability on the machine. Each Web Conferencing Server is a multi-threaded process that takes advantage of multi-CPU systems. Thus, the RTC Server architecture is designed to harness all the network capacity, CPUs, and memory available on any SMP node. A typical configuration of the RTC Server on an SMP machine is shown in Figure 8. The RTC Property Manager (component of the RTC System Services) makes it easy to support different configurations for these RTC systems.

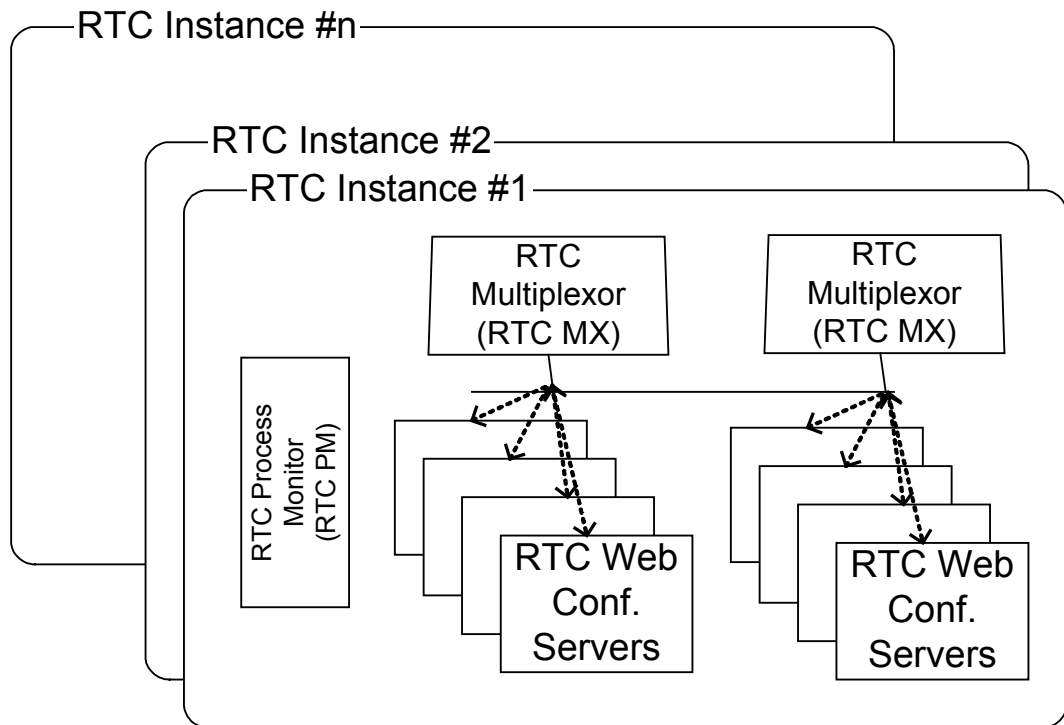


Figure 8 Architecture that Scales Up and Out

As shown in Figure 8, each of the RTC Instances is bound to a physical node. By clustering several of these machines together, one can easily scale up for large capacity workloads.

3.4. HIGH AVAILABILITY FOR CONFERENCES

The RTC System does not have a single point of failure. Even when things fail, there are techniques in place to make the failure absolutely transparent to conference attendees or to limit the scope of the failure as follows:

- If the RTC MX goes down, all the clients and servers will lose their connections to the failed RTC MX component. Both the RTC

Consoles and Web Conferencing Server have the automatic reconnect feature, which facilitates the resumption of the connections via the other RTC MX components active in the node/instance. The RTC Process Monitor immediately brings up the RTC MX component, which is ready for connections.

- If the RTC Web Conferencing Server goes down, only the conferences residing in the current server are affected and the rest of the conferences active in other Web Conferencing Servers in the instance are unaffected.
- The Web Conferencing Servers are restarted a few times every day due to the complicated nature of JAVA garbage collection and other resource usages. Each Web Conferencing Server schedules itself for restart as soon as all the conferences active within the server come to an end.
- If the RTC Process Monitor goes down, the Oracle Application Server Process Monitor (OPMN) brings it back. This does not affect any conferences active in the system.

The above features in the RTC Server architecture offer the highest availability for each conference and the best reliability for the entire RTC system.

4. INTEGRATION SERVICES AND SITES

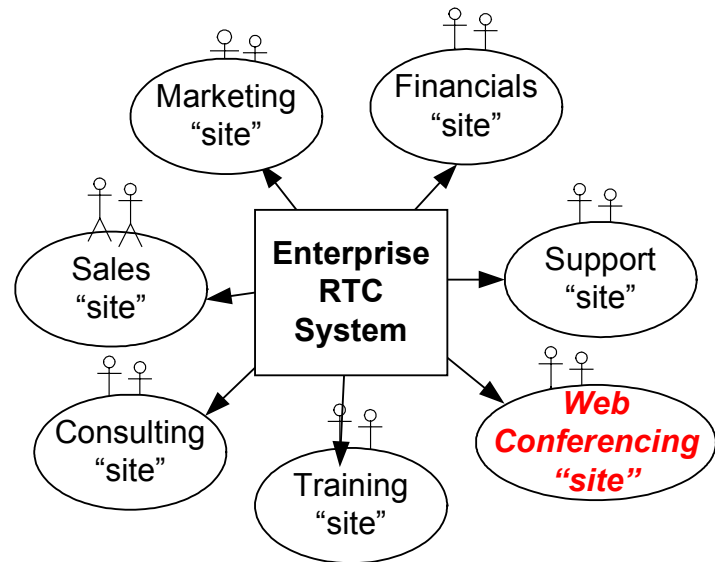


Figure 9 RTC Sites for Various Lines of Business

The RTC administrators can create “sites” for each organization to integrate with the RTC system using their own Web or desktop applications. After the creation of the site the RTC Integration Services will be used to pass the

requests and responses to the RTC Application Module. For example, as shown in Figure 9, a site for the sales organization is created in the RTC system. Subsequently, requests from users of the sales applications to create and join Web conferences are routed to the RTC system. Users of the sales application are unaware that they are using the RTC system. Also, if the sales site is created such that every conference starts in Desktop Sharing mode with SSL encryption and recorded, then the RTC system starts every conference in this mode. This technology feature set offers the enterprise a single shared instance to be used by *all* the lines of business.

The RTC Integration Services use the Web Services paradigm of sending the requests data in XML over HTTP. This facilitates easy integration into all the Web applications used at an enterprise.

5. SAMPLE DEPLOYMENTS

DMZ Deployment

As shown in Figure 2, it is typical to deploy the mid-tiers in the DMZ and the RTC Repository and the Oracle Application Server Infrastructure system in the intranet. All the RTC instances on the mid-tiers will be running the RTC Application Module in HTTPS. It is necessary to have the RTC instances in the DMZ to facilitate the connections to users outside the enterprise.

Geographical Deployments

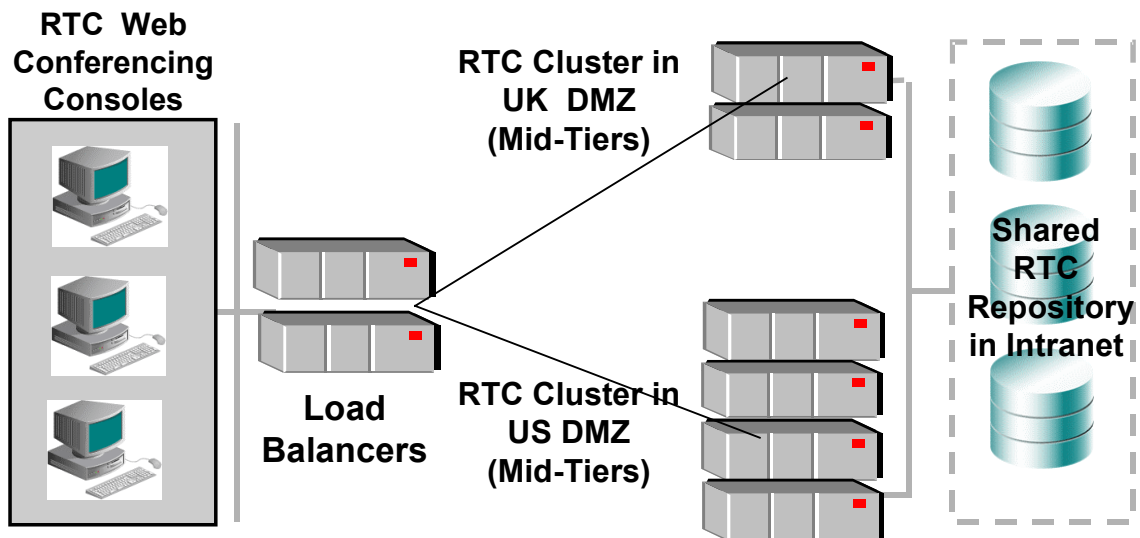


Figure 10 Geographical Clusters in UK and US

In a large enterprise with a WAN spread across different parts of the world, it makes sense to deploy a geographically dispersed solution to reduce the network latencies for conference attendees. Figure 10 shows a case where the RTC system has two RTC clusters, one in the United Kingdom (UK) and another in the United States (US). All the users of the system from Europe are connected to the RTC cluster in the UK and the network traffic for conferences are localized to that part of the network.

6. ROADMAP

Real-time presence and instant messaging will be added to the Real-Time Collaboration feature set. In addition, the Web Conferencing features will be augmented to provide more advanced whiteboard, full-duplex voice streaming, improved multi-cast support, and more. These features make the Oracle RTC system the platform of choice for real-time collaboration for enterprises of all sizes. For more details, please refer to the Oracle Real-Time Collaboration Statement of Direction collateral within the Oracle Collaboration Suite.

CONCLUSION

Oracle Web Conferencing is the real-time collaboration component in the Oracle Collaboration Suite. This functionality is offered by the Oracle Real-Time Collaboration (RTC) system. The RTC Software architecture has five components: Consoles, Server, Application Module, System Services, and the Repository.

The RTC Web Conferencing Console facilitates support for cobrowse, desktop sharing, whiteboard, chat, poll, document presentation, and voice streaming in a single conference. The RTC Playback Console can play back the synchronized data and voice stream from the RTC Repository, or from the locally downloaded conference archives. The RTC Server architecture uses a multi-threaded process and a process pool model to scale up on SMP machines. Several of these machines can be clustered together to scale out for larger workloads. In addition, the RTC Server offers superior real-time performance to all attendees in a Web conference through a packet collapsing algorithm for users on different network bandwidths.

The RTC system is a distributed system with software components working together on different machines bound together by the RTC System Services. The property management component makes it easy to configure and change the settings of an RTC system that consists of several RTC clusters (geographically distributed). RTC Process Monitor, RTC Server architecture with no single point of failure and automatic client and server reconnect features provide high availability and 24 X 7 reliability to the RTC System.

The Oracle RTC system offers a secure environment by providing restricted conferences for invited attendees only, the conference archives with the same restrictions as the conferences and SSL encryption. It offers integration services (XML/ HTTP) to facilitate integration into all the portals, Web, and desktop applications in use at the enterprise. Every line of business gets its own view of the system using the RTC sites feature, where by a single instance for the enterprise can be shared by all the organizations.

Oracle Web Conferencing system can integrate with any LDAP system for user management, requires no changes to firewall configurations, and consumes least amount of network bandwidth, supports and scales to the Web conferencing needs of *all* enterprises.

Oracle Real-Time Collaboration System Architecture Overview
October 2003
Authors: Ramu V. Sunkara
Copyright © Oracle Corporation 2003
All Rights Reserved Printed in the U.S.A.

This document is provided for informational purposes only and the information herein is subject to change without notice. Please report any errors herein to Oracle Corporation. Oracle Corporation does not provide any warranties covering and specifically disclaims any liability in connection with this document.

Oracle is a registered trademark and, Oracle Real-Time Collaboration and Oracle Web Conferencing are trademarks of Oracle Corporation.

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
1+650.506.7000
Fax 1+650.506.7200

Copyright © Oracle Corporation 2000
All Rights Reserved