

# Oracle Video Server System Technical Overview

*An Oracle White Paper*  
*March 1998*

# Oracle Video Server System Technical Overview

## INTRODUCTION

The Oracle Video Server system is an integrated technology for streaming digital video from a server to multiple clients. This system's benefits include real-time delivery, minimal client storage requirements, elimination of bulky media, VCR-style control over content playback, and support for value-added services such as near video-on-demand. This white paper identifies the components of this system and explains how the components interact.

## WHAT IS STREAMED DIGITAL VIDEO?

Streamed digital video works much like the digital video you may have already used; video is recorded in an analog format or directly to a digital format. If the original is in an analog format such as videotape, it is encoded (compressed) into a digital format that can be stored on disk and decoded for playback.

Streamed video, however, is fed from a server with large storage and delivery capabilities to a client that decodes and displays the streamed video as it arrives. This removes the need for the video to physically reside either on a playback medium (such as a laser disc) or locally on the machine that displays the video.

Streamed video provides several advantages over such traditional media as film and videotape. Like the video on a laser disc, streamed digital video does not degrade or "wear out" from repeated usage. It remains in original condition until it is explicitly deleted, modified, or overwritten.

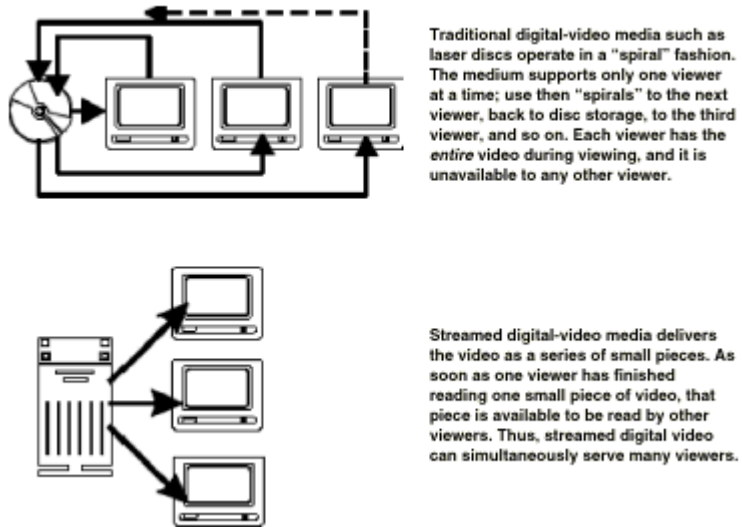
Streaming video over a network also results in a high degree of portability. Unlike the laser disc, streamed digital video does not require that you physically move the storage medium with the display mechanism. The client machine need only have network access to the server machine where the videos are stored.

Digital video can also be viewed as non-streaming, downloadable files. However, such files retain many disadvantages of traditional storage media:

- You must download the entire file before you can view any of the file. In contrast, you can begin viewing streamed video as soon as the beginning of the stream reaches the client machine.
- The local machine must provide storage for the entire downloaded video file. Streamed video requires the client machine to provide only enough capacity to store and display the video portion currently being delivered and viewed (and perhaps a small buffer to support error-recovery).

More than this, streamed digital video is scalable. Many people can view the same video—not just multiple copies of the same movie title—at very nearly the same time. Figure 1 shows the difference in these delivery mechanisms.

**Figure 1: Traditional digital video vs. streamed digital video**



This important feature of streamed digital video and video server systems is called video-on-demand. It enables your clients to watch what they want to when they want to, choosing any title that has been digitized and stored for delivery.

Another way to make streamed digital video available to clients is to schedule broadcasts at regular intervals. An administrator can schedule delivery of a specific video at a specific time on a specific channel. Near Video-on-Demand (NVOD) is an application of scheduling in which the administrator schedules delivery of a video to begin automatically at regularly scheduled intervals on a different channel. For example, you might schedule a training video to begin at 9:00 a.m. and every 60 minutes thereafter, until 8 p.m. If users miss the 9:00 showing, they can wait until 10:00 a.m. and watch the appropriate channel. Other users could watch when their schedules permit.

Video-on-demand and scheduled video are possible because the data is stored as a series of stripes, typically of 32 KB or 64 KB each, with the stripes distributed over multiple disks. A single stripe might represent several seconds of video. As soon as the stripe has been read by one client, that stripe is available for another client. This allows many different parts of the file to be read simultaneously, enabling a single copy of a video to serve many users concurrently.

Striping also enables a sophisticated backup-and-recovery system known as RAID (redundant array of inexpensive disks). With RAID, video can still be delivered during a disk failure. RAID needs to recover only a single stripe from the failed disk, rather than an entire video. This reduces the amount of storage required for duplication, allowing more space to be devoted to video content.

Even after encoding, digital video files are still quite large. Depending on the number of titles you want to make available, storage on a local machine may not be practical. A video server machine offers the storage and throughput capacities needed for timely and reliable delivery of digital video. In turn, this means you need not find storage for every video you want to view on your desktop computer—which would leave little room for anything else!

## **ORACLE VIDEO SERVER SYSTEM**

The Oracle Video Server (OVS) system is a unique implementation of video server technology for networked computers that store, manage, deliver, and display digital multimedia data on demand—including real-time, full-screen video and high-fidelity audio. OVS is supported on a variety of server platforms and can scale to serve many thousands of concurrent users.

### **OVS System Architecture**

The OVS system is fully compliant with Oracle's Network Computing Architecture (NCA), which consists of three tiers:

- Client
- Application Server
- Data Server

The NCA also defines three layers:

- Server Layer: provides the basic functionality for each tier
- Cartridge Layer: provides programmatic functionality for the Server Layer of each tier
- Communication Layer: provides communication between servers, cartridges, and tiers

#### **Tier 1: Client**

There are two client components provided as part of the OVS system.

- **Oracle Video Client:** The Oracle Video Client (OVC) is responsible for obtaining, decoding, and displaying the video stored on the server. The OVC provides multiple interfaces that allow you to incorporate streaming audio and video from OVS into your own applications.
- **Oracle Video Server Manager:** The Oracle Video Server Manager (VSM) is a Java application that provides a management environment for OVS services, clients, and content.

Both of these components are discussed in more detail later in this white paper.

### **Tier 2: Application Server**

The Application Server tier consists of the Oracle Video Server and the VSM application logic. The OVS functions as the application server in this environment, receiving and processing requests for digital video, then delivering the digital video content to the client device. The VSM application uses both HTTP and CORBA (Common Object Request Broker Architecture) as the network layer.

### **Tier 3: Data Server**

The Data Server tier contains the target services the client needs in order to access application-specific data. In the OVS system environment, the Data Server tier consists of the Oracle database (optional). A database is a reliable repository for storage of structured data—that is, the *metadata*. The OVS system can use the Oracle database to write and query database tables associated with OVS processes, including *logical content*, clips, and schedules.

## **ORACLE VIDEO SERVER**

The OVS functions as the application server, receiving and processing requests for digital video, then delivering the digital video content to the client device. The OVS includes the following components:

**Oracle Media Net (OMN)** - OMN is a networking infrastructure that enables the OVS and its clients to communicate in a distributed computing environment. Using heterogeneous network protocols, OMN enables connectionless communication among OVS system components running on different platforms. OMN is Oracle's implementation of CORBA. It allows services programmed in different languages and distributed over different computers to communicate without concern for each other's location or the details in transporting and converting data among them.

OMN abstracts the network for the application developer, and the network protocols underlying OMN are transparent to the components using it. You can create an application once and deploy it on many different platforms without changing its OMN calls. OMN also automatically bridges connections across network types, making multiple protocols in the same network transparent and presenting a reliable message service, even if the underlying network protocols use unreliable messaging.

**Media Data Store (MDS)** - Oracle developed the patented MDS for scalability. Distributed reads take advantage of SCSI technology; reads from multiple disks return data almost simultaneously. The MDS directory server controls access to MDS files and manages their layout on disk. Processes that read or write MDS files must first gain access through the MDS directory server. For example, when a video pump is instructed to play a video file, it sends a message to the MDS directory server, which returns a small structure describing

the layout of the file on disk. With this structure, the video pump accesses the file's contents directly. This direct access prevents the MDS directory server from becoming an I/O bottleneck.

**Video Pump** - The video pump reads video files from the MDS and delivers them to the network in real time. When a client requests video, the video pump receives a message from the stream service, reads the appropriate portion of the file from the MDS, and sends video data to the client.

**Stream Service** - The stream service handles requests from clients for video and audio content. When the stream service receives a request, it obtains necessary information to meet the request. If the request is for a logical content title, rather than for a tag file, the stream service contacts the content service for a list of tag files necessary to meet the request. The stream service reads the tag files associated with the request. Each tag file describes which portion of a requested file must be delivered to meet the request.

- The *broadcast data service* manages broadcast scheduling information in the database. This service reads broadcast schedules from the database and writes them to memory where they can be read by the scheduler service.
- The *scheduler service* tracks the current time and a list of scheduled events and notifies the appropriate service, such as the Near Video-on-Demand (NVOD) exporter service, when a broadcast event should occur.
- The *NVOD exporter service* plays a scheduled logical content title on a scheduled channel when notified by the schedule service. You can use this service to handle any scheduled video, including NVOD, pay-per-view, or regular TV broadcasting.

### **OVS Scalability**

OVS has been proven in applications requiring thousands of concurrent high-bandwidth accesses to both long and short duration video streams. These applications have been integrated into the IT infrastructure to support billing, ad insertion, and content management systems.

OVS provides scalability across multiple dimensions:

- **Concurrent Streams** - If demand for a server (such as an OVS service) increases, you can run additional instances of the server to better handle the load. Oracle Media Net can distribute server access request across the available instances to balance the load. If one instance stops, Oracle's Media Net redistributes new requests across the remaining instances.
- **Bit rates** - In networked media systems, quality improves linearly as bandwidth increases. OVS enables scalability across the widest range of bit-rates. A single architecture spans rates from 28.8Kb/sec modems to ATM, 100 Mbps LANs, and digital satellite delivery systems.

- **Hardware platforms** - OVS is truly multiplatform; from a single-processor server to the latest Symmetric Multi-Processor (SMP) clusters to Massively Parallel (MPP) media servers such as nCube. OVS is optimized for each architecture by Oracle product lines and by leveraging our strategic partnerships with the hardware vendors.

### **File Availability with the MDS**

OVS stores physical content in the MDS, a file system optimized for storing and delivering video in real time. The MDS can tolerate disk errors and varying disk latency while providing real-time access. The MDS can also deliver real-time video while you are administering the MDS.

The MDS enforces write consistency on all MDS files:

- Files can be written by only one client at a time.
- Files cannot be renamed, removed, truncated, or locked into read-only mode while being written.
- If a process fails while writing to a file, the MDS makes the file available to be written by other processes after a few minutes.

MDS files are stored in MDS volumes, which are named collections of disks. Each MDS volume stores files in a flat namespace. The MDS provides quick access to a file given an identifier, such as a filename, but MDS is not meant to be a general-purpose file system with a hierarchy of subdirectories. See the “Logical Content” section for information on searching and browsing the file system.

MDS provides uninterrupted file access to multiple clients using the following technologies:

- **Disk Striping** - Striping divides a file into pieces and stores each piece on a different disk. As a result, the file is distributed across many disks, rather than concentrated on one. This spreads requests among the disks and improves performance when the file is accessed by concurrent clients. Each file is striped across all the striped volume’s disks. Striping is transparent to the user and a striped file appears logically as one contiguous sequence of bytes.
- **RAID (Redundant Arrays of Inexpensive Disks) Protection** - RAID protection means storing data redundantly so it remains accessible in the event of a disk failure. The *software* RAID protection implemented for the MDS is similar to hardware-based RAID protection but also guarantees on-time delivery even if a disk or controller fails. RAID is similar to disk mirroring, but uses less space for redundant data.

The MDS was written for high availability. If a disk in a RAID set (or a controller in a multi-controller RAID set) fails, MDS rebuilds the data from

duplicate information in the RAID set. This is all done while streaming to the client machines, so the users do not even know there has been a disk failure.

### **ORACLE VIDEO CLIENT (OVC)**

The Oracle Video Client (OVC) is responsible for obtaining, decoding, and displaying the video stored on the server.

In addition, the OVC can do more than simply request and display digital video. Its true power lies in allowing you to incorporate streaming audio and video from OVS into your own applications. Such customized interactive video applications can include computer-based training (CBT), interactive kiosks, corporate repositories, and Web sites. A broad range of integration tools and sample applications give application developers virtually unlimited flexibility.

**Application Integration Tools** -The OVC software contains a variety of components, each extensible and customizable. Developers can deploy the base video client on various machines, using these “extensions” as the means of integrating the client’s functionality:

- Java native classes and a Java player applet, to provide consistency and portability across a variety of operating platforms, such as Microsoft Windows and various “flavors” of UNIX. Use of a Java-based video client also eliminates the need for an HTML browser on the client.
- Netscape™ plug-in, for use with compatible Internet browsers, such as Netscape Navigator and Microsoft Internet Explorer. This allows integration of “in-line” video within a web HTML page without launching a separate media player application window. You can also use Java and JavaScript and with Netscape’s LiveConnect interface to customize control of the Oracle Video Web Plug-in. For example, you can add clickable buttons or icons to play, pause, and seek video and to display lists of available content.
- ActiveX & OCX custom controls, to integrate the video client into applications designed to operate in compliance with 32-bit multimedia applications, such as Microsoft Visual Basic, Developer 2000/Oracle Forms, and Oracle Power Objects. Such applications can start, stop, and seek locations within video streams from the OVS.

OVC also provides a rich set of CORBA-compliant APIs that allow developers to create their own client. These APIs also provide fine-grained control over the media stream itself, which is needed by specialized applications such as Media Asset Management and Authoring Tools.

**Other OVC Features** - In addition to a wide range of application integration tools, OVC supplies other functionality that improves overall quality and ease-of-use. OVC software applications can receive and display video content using Microsoft’s ActiveMovie component architecture. ActiveMovie “calls” the

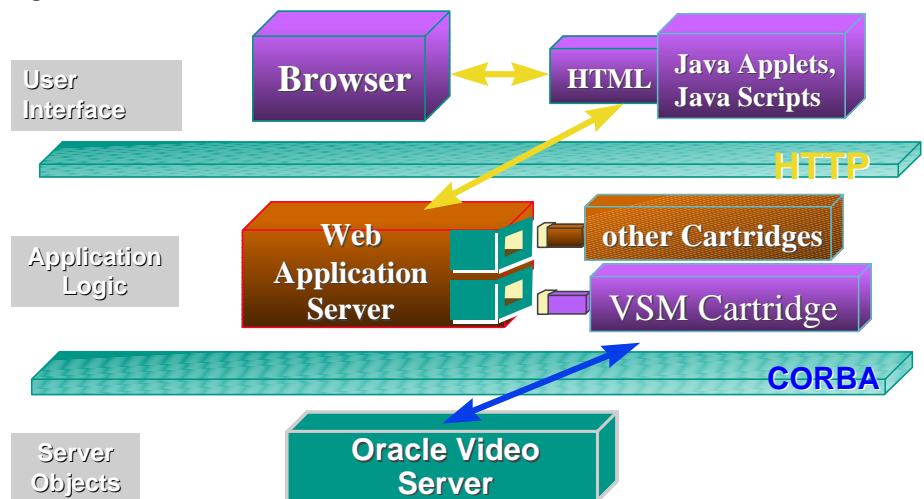
appropriate video decoder without programmatic or user intervention, allowing developers to concentrate on building applications and not on adding application logic for various encoding formats.

The decoder requires smooth, reliable video delivery, but the network may sometimes deliver packets late or out of order. OVC ensures proper decoder operation by re-ordering packets properly and buffering a small amount of data to provide the decoder a consistent stream of ordered video data.

### VIDEO SERVER MANAGER (VSM)

VSM is a web browser-based application that monitors and manages the OVS and its clients. It is implemented as an NCA cartridge. VSM monitors OVS processes and streams and enables management of MDS content. The VSM Client is a set of Java classes that can be run as an applet or application. It is packaged into a single JAR file for portability between operating systems.

Figure 2: VSM Architecture



VSM enables the administrator to perform a variety of tasks, including:

- starting and stopping instances of the OVS
- viewing the status of critical OVS services
- monitoring and managing OVS functions and components
- MDS disks
- HSM tertiary storage tapes
- logical content
- broadcast schedules (including near video-on-demand support)
- clients of the OVS

- real-time feeds
- scheduling and playing videos

## PHYSICAL VS. LOGICAL CONTENT

### Physical Content

Physical video files are created from actual video through encoding by a third-party software or hardware codec (compressor/decompressor). Various codecs conform to international standards such as Moving Picture Experts Group (MPEG) or H.263. Other proprietary codecs support fractals or wavelets. These formats compress video frames so they use less space. The way in which an encoder multiplexes, or puts together, video and audio is called a *container* or *mux format*.

The OVS can stream a variety of compression formats, including:

- MPEG-1 or MPEG-2
- Iterated Systems ClearVideo
- Radius CinePak
- Intel Indeo

The OVS can store and retrieve a variety of containers such as:

- Audio Visual Interleaved (AVI)
- MPEG-2 Packetized Elementary Stream (PES) or Transport Stream Packets
- MPEG-1 Packets or Packs
- RKF (Raw Key Frame) Container Formats

The OVS can also deliver video files in any container format meeting the raw keyframe criteria:

- stateless - all the data for displaying the picture in a video frame is contained entirely in that frame, rather than in any previous frames.
- contiguous - all the data for a frame is stored together in the video file; no other data is mixed with it

### Logical Content

Logical content is an abstract view of the video and audio files contained in the MDS. It allows the creation of a unified stream from a group of clips, physical content, tag files, and content files.

Although your video application may play a single movie or recorded speech from beginning to end, you may also want to include previews of other movies, content, or upcoming events. Each video segments is known as a *clip*.

The OVS enables content creators to assemble clips and label them with a single title. Because these collections of clips do not need to be assembled into a single physical file, they are called *logical content*. As beads can be strung together in different combinations to create a piece of jewelry, so can logical content clips be “strung together” to create a unified stream of digital video. This unified stream eliminates delays between clips and supports looping of the entire sequence or the last clip.

The *content service* maintains information that maps logical content to physical content. A client application can query the content service for a list of logical content or to resolve logical content to physical tag files. The Content Resolver interface to the Oracle database dynamically creates logical content. An optional static version does not require a database.

### **REAL-TIME FEEDS**

While the majority of digital video is *stored*, there is a growing demand for live or *real-time* video. In the real-time model, content is broadcast as it is occurring. For example, breaking news or an interactive training course may need to be delivered in real time.

Both stored and real-time video have their strengths and weaknesses. Stored video allows viewers to start watching whenever their schedule allows. Real-time video provides information as it occurs. This adds the possibility of support for two-way interactivity.

Unique OVS technology called Real Time Feeds (RTF) marries the best of the stored and real time models. OVS exposes a server-side API that digital encoder vendors can use to record, encode, and begin streaming digital video *as that event occurs*. News channels, stock updates, sports, corporate events, and interactive services can use one-step encoding (video is encoded as it is recorded) and RTF to supply the latest information via digital video. The viewer sees the event nearly as it happens, the only delay being the minimal time necessary to encode and deliver each moment of the event. Viewers can also pause, restart, rewind, and fast forward the video up to the current “live” point, enabling customized “instant replay” on an individual basis.

**Instant Replay for every user** - RTF allows OVS to act as a “sink” as well as a “source” of video. The system administrator can define a buffer store a specified amount of the feed. RTF automatically *wraps*, adding new content and deleting the oldest content. Content becomes viewable about 15-25 seconds after the event occurs. Clients can rewind up to the length of the buffer for this feed. As a result, the client has complete control over “instant replays”.

RTF supports continuous feeds direct from encoder to MDS, with history from a range of popular encoder products such as Optivision, Lucent, Nuko, Minerva and Futuretel. This allows encoders to provide both video and metadata, including format, bit rate, and tag information for storage in the Oracle database.

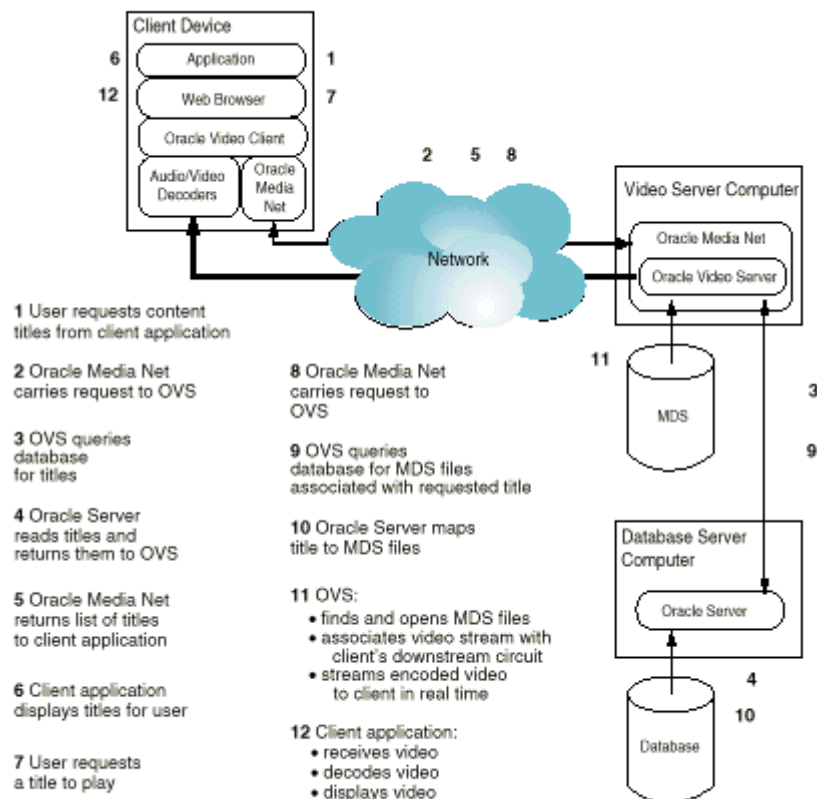
**A VCR in the network** - Since the content in the rotary buffer can be indexed and tagged, users can point-and-click to view previously recorded video. If you can't be home to watch Seinfeld, instead of programming your home VCR, consider an on-line, Electronic Programming Guide (EPG). Users could click [Seinfeld/9 p.m./Thursday/Channel 4](#) and watch their favorite shows at their leisure.

These unique capabilities of OVS will create entirely new classes of broadcast applications in sports, news, financial information, and television network programming.

### THE VIDEO ROUND TRIP

Figure 3 shows the steps the OVS system follows to deliver video. This example uses a symmetric network. Steps 7 through 12 make up a video *round trip*.

**Figure 3: An OVS system round trip**



In Step 7, when the client requests a streamed video file over the network, the request goes to the Stream Service, using the stream ID to identify the file.

In Steps 9 through 11, the Stream Service opens the tag file for this stream ID and reads the bit rate, presentation rate and content type. It compares the values

to the parameters provided by the client. If these parameters do not match, the Stream Service refuses the client request. If these parameters match, the Stream Service passes the stream id and file name to the Video Pump and instructs it to begin streaming the video file to the client.

The Stream Service does not touch content files, and the Video Pump does not touch tag files.

If the client requested that the video file be streamed from some point other than the beginning, the Stream Service looks in the tag file for the appropriate location in the content file. It instructs the Video Pump to begin streaming the content file from that location. In some cases, the Stream Service gives the Video Pump data to prepend to the stream. This data makes the stream a valid video stream. The Video Pump then opens the file, prepends the data buffer if necessary, and begins streaming the data to the client.

### **NETWORK TRANSPARENCY**

OVS provides a rich infrastructure for video networking applications. It transparently supports almost any physical datalink, including ATM, HFC, Ethernet, FDDI, ISDN, Cable modems, xDSL and analog or digital modems. OMN supplies the abstraction for control messages between clients and servers, while Internet Protocol (IP) handles video delivery over different networking topologies.

The System Administrator can specify services and capabilities via configuration utilities that indicate:

- which networking protocols (for example, TCP, UDP, ATM) are available from a given service
- aggregate bandwidth available from network interface cards
- specific bit-rates for a given “channel”
- specific range of client addresses from which a service can be allocated

These parameters are built upon a hierarchical model, allowing clients to make requests for video and brokering those requests across a range of services.

### **SUMMARY**

The Oracle Video Server system is an integrated technology for streaming digital video from a server to multiple clients. The major components of the system are the Oracle Video Server (OVS), the Oracle Video Client (OVC) and the Video Server Manager (VSM).

- The OVS is comprised of several components, including the Media Data Store (MDS), a patented file system that has been optimized for storing and delivering video to multiple clients in real time.

- The OVC requests and displays video from the server. It also includes a number of application integration tools that enable developers to create customized applications that use the OVS system.
- The VSM is a web browser-based application that monitors and manages the OVS and its clients.

Oracle Video Server System Technical Overview  
March 1998  
Author: Daryl Rosen  
Copyright © Oracle Corporation 1998  
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 Enabling the Information Age, Oracle Video Server, and Network Computers are trademarks of Oracle Corporation.

All other company and product names mentioned are used for identification purposes only and may be trademarks of their respective owners.

**ORACLE**

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

Worldwide Inquiries:  
415.506.7000  
Fax 415.506.7200  
<http://www.oracle.com/>

Copyright © Oracle Corporation 1998  
All Rights Reserved