Mohamad Sarfraz -- email: "Sarfraz, Mohammed" <sarframx@Exchange.Aramco.com.sa>

White Paper

Video Streaming in Saudi Aramco

Video Streaming Project Team
Computer Applications Department

TABLE OF CONTENTS

1. Introduction

Computer Applications Department in partnership with other departments in IT has been involved in providing video streaming services for all IT customers, providing an alternative means of communication. This method of communication has been widely adopted as an effective means of communication primarily by Public Relations Department, Management Training Department, SAMSO, SAP training, Loss Prevention and number of other organizations.

The increased interest in this form of communication has also translated into demand for higher quality service including guaranteed delivery. While the current infrastructure was evolved based on gradual adoption of video streaming, it is not capable of handling this sudden increase in demand.

This white paper provides an overview of the services being offered via video streaming and also provides a list of areas that need further immediate attention to improve the quality of the service. Being a service organization, it is imperative that appropriate steps are taken to improve the quality of service for our customers.

2. Current Situation

The current video streaming infrastructure is based on a central distribution model (see figure 1). Tower Computing Center (TCC) currently hosts the production video streaming servers. The current infrastructure based on a central distribution model is close to saturation point and has a capacity of handling about 2,400/1,200 concurrent users, based on encoding rates of 100kbs and 250kbs respectively . In most of the recent video streaming events, close to 5000 - 6000 users have been trying to access the live event. This number of users is above and beyond what current infrastructure can handle. This has also resulted in degradation of the quality of service as network congestion on a critical segment usually has a domino effect.

2.1. High-Level Overview of Video Streaming Services

This Section provides a high level overview and configuration of the services that are offered over the video streaming network. Different organizations within IT work together to provide these services. CAD provides a single point of contact for all customers.

The video streaming service is currently offered using Unicast communication over the TCP/IP network. The services that have been provided by to customers are as follows:

1. **Video Library** – It is an on-demand collection of close to 1800 videos. The contents are indexed for search. There is an increase in demand for this service and there is a backlog of about 1000 videos to be made available. Key customers for this service are PRD, Loss Prevention, SAP and Management Training.
2. **Live Broadcast** – This service is offered in partnership with Media Productions from Public Relations. It is usually used for executive speeches, streaming of conferences, etc. It has recently being expanded to perform broadcasts from remote locations by making use of VSAT.
3. **Video Conferencing Bridging** – This service was established by CAD in partnership with the Video Conferencing Group in Services department to allow desktop access to the video conferencing network. The service allows broadcasting of VC sessions and is currently being utilized by SAMSO to allow their staff to view regular sessions from the Mayo Clinic in United States.

4. **Live broadcasts from Media Production Studio** – A live video channel has been configured with video servers in TCC to allow live broadcasts to take place directly from studios in BLDG 3030.
5. **Media Lab** – CAD has setup a Media Lab to allow it to provide services related to create, encode, edit digital content and other related services.

Most of these services are running in parallel i.e. while there is a live event going on customers might be simultaneously accessing on-demand video on video library while another set of customers may be tapping into a live video conferencing event.

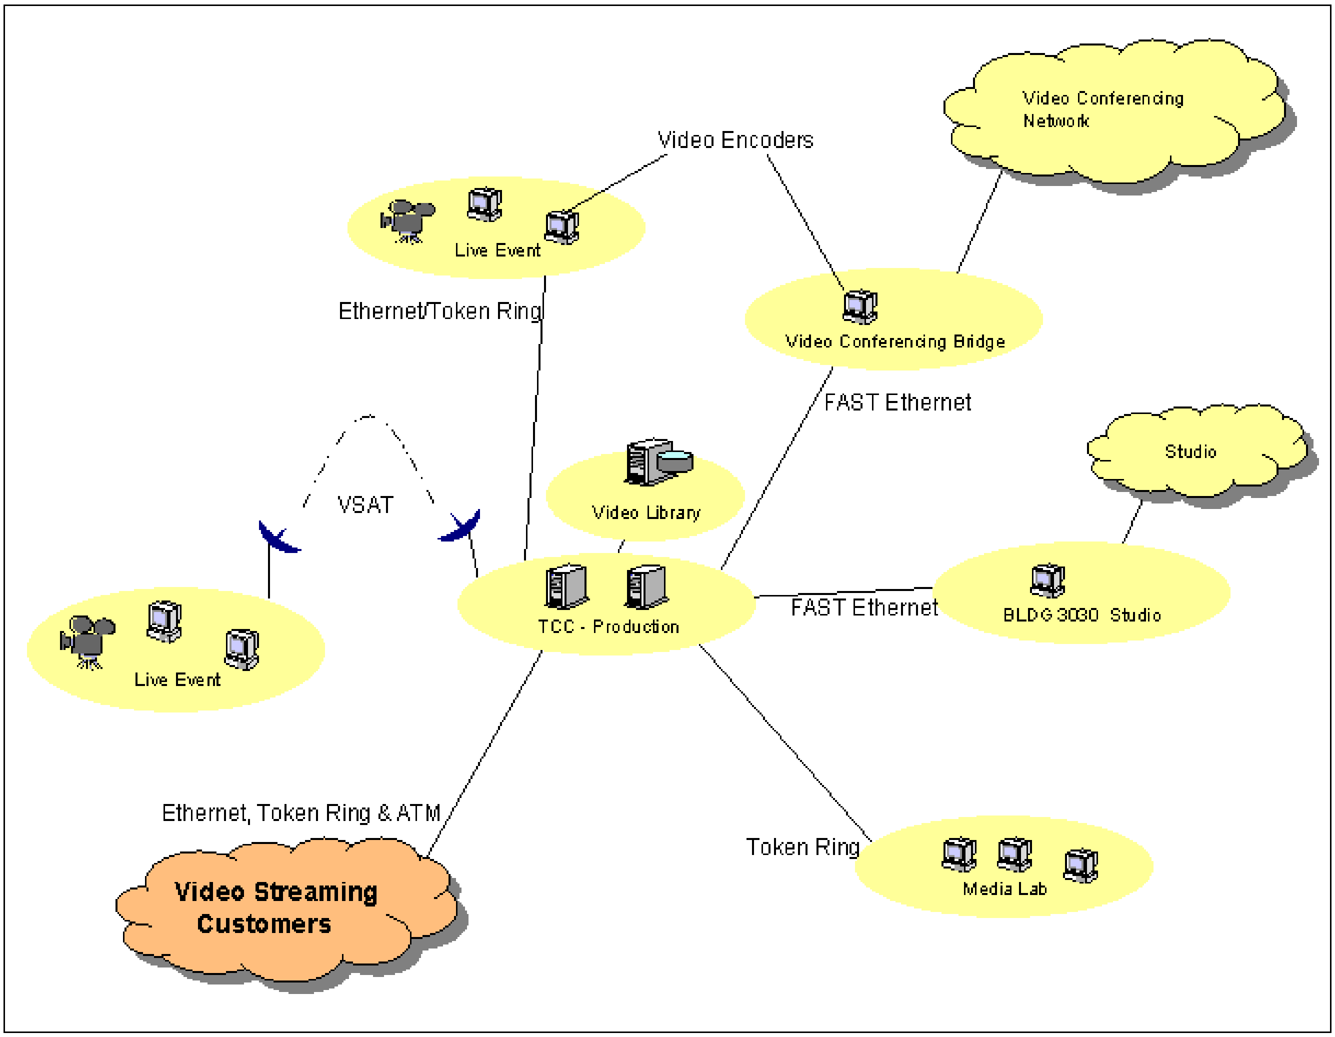

Figure 1 Video Streaming Services

2.2. Infrastructure

This section describes the current video streaming infrastructure that is located in TCC. The infrastructure is based on a centralized model. It was implemented couple of years ago starting with a single server and has been evolved gradually based on the increase in demand.

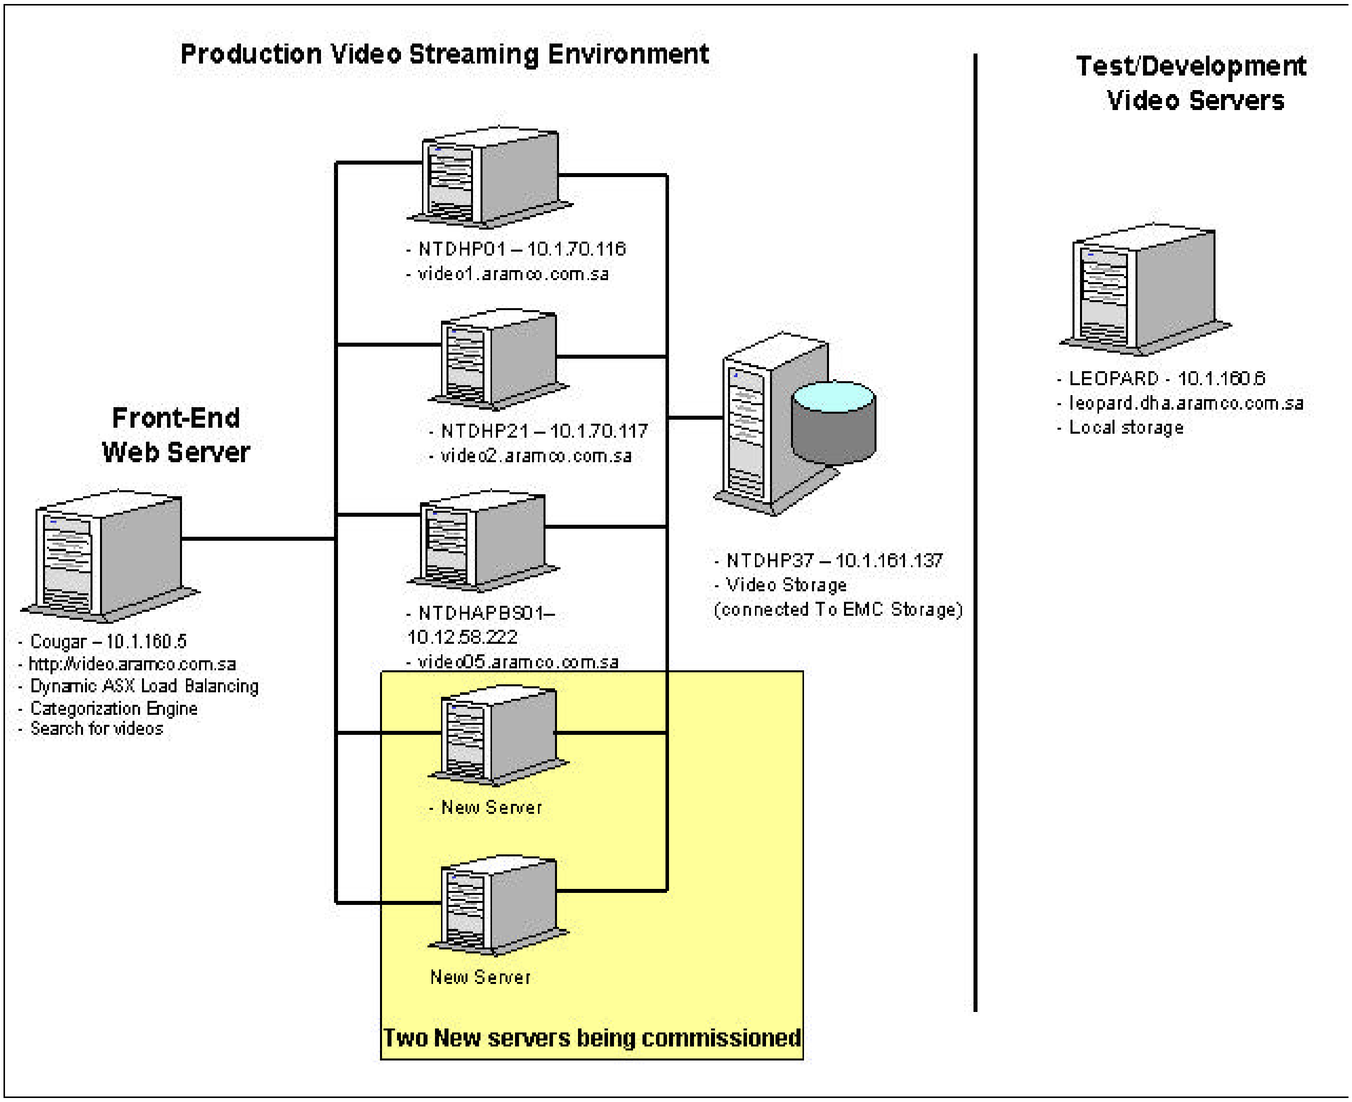

Figure 2 Infrastructure in TCC

It consists of three video servers connected to shared storage on EMC. The servers are load balanced using dynamic ASX load balancing. While the infrastructure handles the load for both live events and on-demand content (made available via video library). For popular live events when the demand is high additional servers from test and development environment are temporarily configured to handle additional load. All the servers are connected to 155 Mbps ATM segment and have been distributed on different segments to distribute load and network traffic.

3. Limitations of Current Infrastructure

The current infrastructure is based on a central distribution model using Unicast distribution. This method of distribution was an obvious choice for the video streaming team couple of years ago as it required least amount of change to existing network infrastructure and was capable of serving the demand at the time. However, the centralized distribution method is at a saturation point due to number of factors, most prominent factor being the network limitation. As shown in the table below there is a limit to the number of Unicast sessions that can be established on a given segment. Support for number of concurrent users can be increased by reducing the encoding bit-rate, however customer expectation for higher quality videos usually makes it difficult. Taking into account the fact that network is shared by number of other applications i.e. mail, SAP, internet etc current infrastructure is capable of handling 2,400/1,200 concurrent users based on streaming at 100kbs/240kbs respectively.

Based on the different network configurations that Saudi Aramco has given below is the list of maximum number of concurrent users that can view the events on-line.

Type Of Network Interface	Bit Rate	Max Number of clients (theoretical limit assuming that no other traffic on network exists)
ATM (622 mbps)	250k	2488
ATM (155 mbps)	250k	620
Fast Ethernet (100 mbps)	250k	400
Ethernet (10 mbps)	250k	40
Token Ring (16 mbps)	250k	64

3.1. Unicast vs Multicast

Media traffic can be classified based on different parameters. The three methods of video content distribution are:

- *Unicast* - a separate copy of the data is sent from the source to each client that requests the stream.
- *Broadcast* - a single copy of the data is sent to all clients on the network.
- *Multicast* - a single copy of the data is sent to those clients who request it.

The current streaming setup only makes use of Unicast distribution where each client regardless of the service requested (live event, video library, video conferencing bridge) gets their own copy of the stream. Broadcast is not recommended as it creates unnecessary congestion on the network and the routers are configured to only allow broadcast on local segments. Multicasting is more efficient in its bandwidth usage because multiple copies of data are not sent across the network. Data is not sent to clients who do not want it. The user is simply instructing the network card on the computer to listen to a particular IP address for the multicast. The client does not have to be identified to the computer originating the multicast. Any number of computers can receive a multicast transmission without bandwidth saturation.

This section outlines the limitations of Unicast technologies when used for video streaming of large corporate events. Given below is a chart that outlines how network bandwidth utilization grows as number of client's accessing the video stream increases.

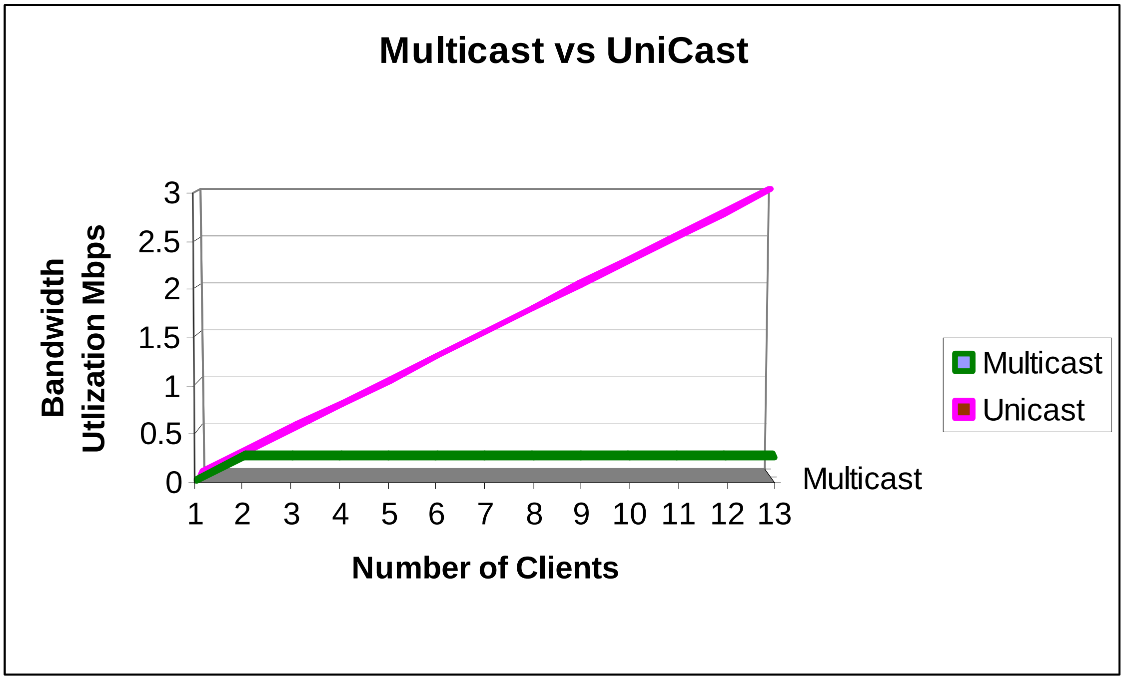

Based this chart as number of users tapping into a stream from a given segment increase it creates a congestion on that particular segment. This impact is not just limited to local segments as number of users from different segments access the servers increase it also impacts the back bone.

Server side logs show that on large events like Hawiyah gas plant inauguration and Grievance policy event close to 5000-6000 thousand users tried to tap into the event within a very short period of time, which resulted in both network congestion on various segments and also servers not being able to serve large number of customers.

3.2. Case Study of a Recent Live Event – Grievance Policy

On January 28, a live broadcast of grievance policy was conducted with Personnel Department. The event was to be broadcasted to limited number of conference rooms and selected desktop users, however customer included the link to video event using company wide email announcement. The event was conducted from LIP. While the servers were configured to handle maximum of 2,400 users only.

This case illustrates how network congestion on one part of network resulted in degradation of the service for whole event. Video encoders were sending data from LIP conference room directly to production servers in TCC, while at the same time large number of desktop users in LIP were trying to connect to production servers in TCC. As load on network segment from LIP to TCC increased due to large number of users from LIP trying to view the event. The load on network resulted in video streaming servers in TCC were not able to get consistent stream, this resulted in degradation of service for the majority of customers who were viewing it from non LIP segments.

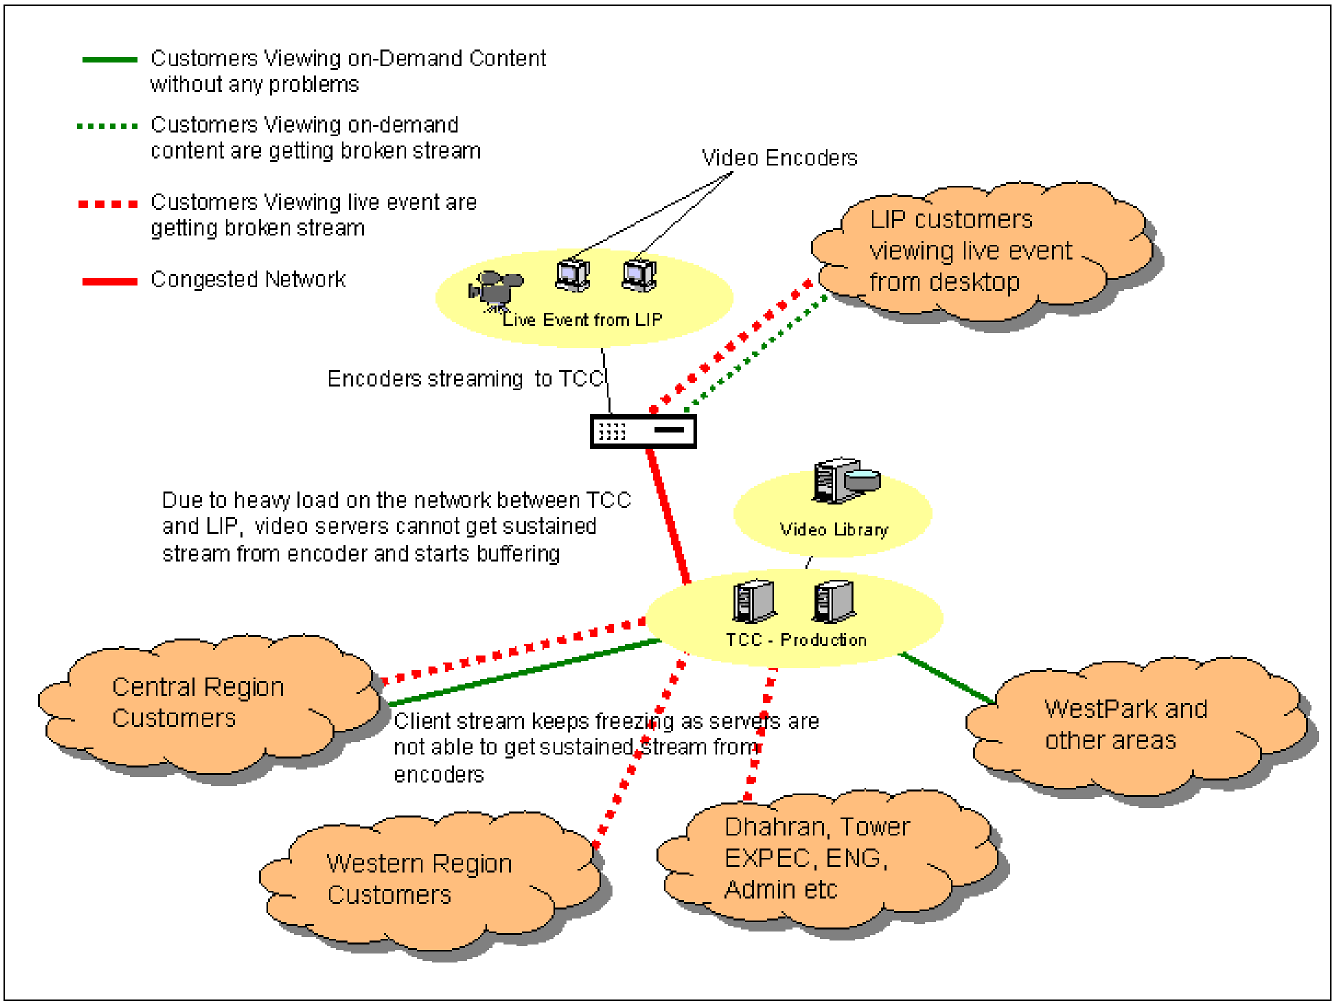

Figure 1 Case Study: Grievance Policy Live Event

If there was a distribution/caching server available for LIP customers so that all the requests from LIP were directed to local distribution server. It would have relieved the network segment between LIP and TCC resulting in better service for large majority of customer who were viewing the event from other segments.

4. Future Directions

It is recommended that current infrastructure be expanded from the current limit of 2,400/1,200 concurrent users to at least 10,000 concurrent users. Perform transition from central model to distributed model while maintaining central control. While multicasting to desktop from centralized location is an ideal situation, taking into account the fact that not all network devices are multi-cast aware and reconfiguration of current network to provide multicasting to the desktop will require considerable amount of redesign effort, It is recommended that a combination of multicast and Unicast distribution approach be used to offer service to larger audience.

It is proposed that architecture similar to the Internet Multicast Backbone (Mbone) is implemented within Saudi Aramco network. The MBone is used for audio and video multicasts of Internet Engineering Task Force (IETF) meetings, the National Aeronautics and Space Administration, and the United States House of Representatives and Senate among others.

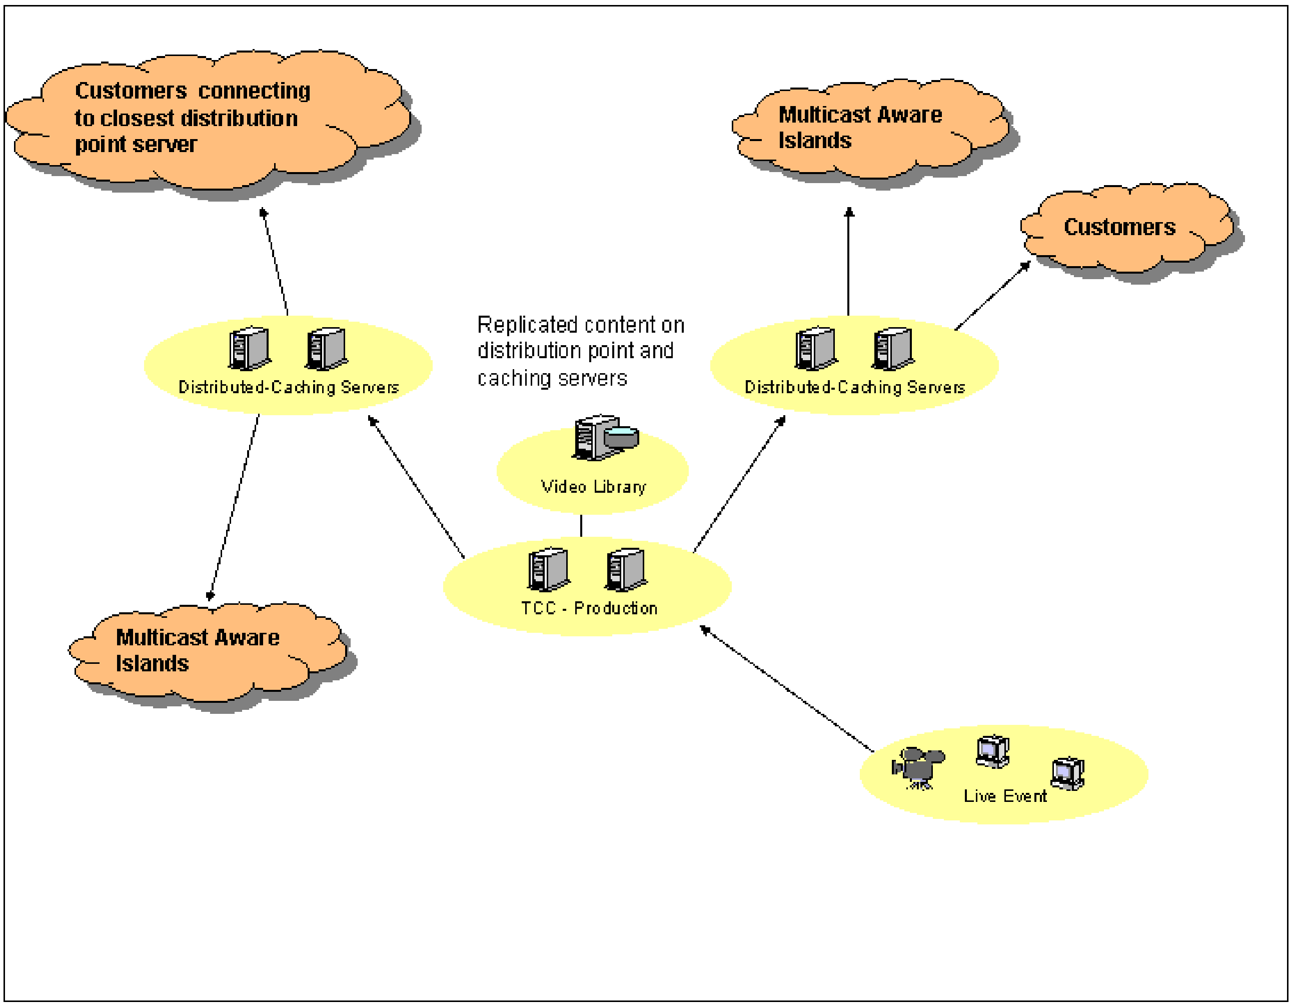

Figure 2 Hub-Spoke Distribution Model

Establish a series of multicast-enabled islands, collections of contiguous networks, connected together using tunnels. Since multicast-enabled islands will only address the broadcast of live events and streaming of video conferencing sessions, It is also proposed that distribution and caching servers are deployed company wide using hub-spoke replication model. These distributed servers will have replicated content from central distribution servers. By enabling context sensitive dynamic load distribution all the traffic for on-demand and live content will be redirected to local distribution servers. Thereby reducing load on network backbone and central distribution servers.