



## **Calculating your Bandwidth and Software License Needs for Flash Media Server 2**

by Chris Hock

January 19, 2007

Copyright © 2002-2007 Adobe Systems Incorporated. All rights reserved.

The information contained in this document represents the current view of Adobe on the issue discussed as of the date of publication. Because Adobe must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Adobe, and Adobe cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for information purposes only. ADOBE MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Adobe may have patents, patent applications, trademark, copyright or other intellectual property rights covering the subject matter of this document. Except as expressly provided in any written license agreement from Adobe, the furnishing of this document does not give you any license to these patents, trademarks, copyrights or other intellectual property.

Adobe and Flash are either trademarks or registered trademarks of Adobe, Inc. in the United States and/or other countries. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Adobe Systems Incorporated  
601 Townsend Street  
San Francisco, CA 94103  
415-252-2000

## Contents

<b>Where to Start? .....</b>	<b>1</b>
<b>Flash Communication Server Editions .....</b>	<b>1</b>
Bandwidth vs. Connections .....	1
Data-only Applications.....	2
Audio-Video Applications .....	2
<b>Calculating Bandwidth .....</b>	<b>2</b>
Prerequisites .....	2
Bandwidth Calculations .....	3
Example 1 – One to Many .....	3
Example 2 – Many to Many.....	4
Bandwidth Considerations .....	5
Client Connection Speed .....	6
Application Design Matters .....	6
<b>Server Performance .....</b>	<b>7</b>
Recommended Server Hardware .....	7
Test Findings for Various Applications and Client Connection Speeds .....	8
Data Applications .....	8
Audio and Video Applications .....	8
<b>Conclusion .....</b>	<b>11</b>
<b>Appendix A - Estimating Number of Simultaneous Users .....</b>	<b>12</b>
What are Simultaneous Users?.....	12
How to Estimate Number of Simultaneous Users.....	12



Flash Media Server 2 provides a very powerful development environment that enables you to deliver compelling, interactive audio and video applications. These applications can be as simple as a one-to-many streamed video or as complex as a many-to-many video conferencing system. Because of this, determining the number of users that can effectively use your application, how much bandwidth you will need, what software licenses this equates to, and how much server hardware you should buy may not be apparently straight forward. This whitepaper is designed to address these specific issues and help you determine the amount of software and hardware you will need for your situation.

## Where to Start?

Most people who are evaluating the use of the Flash Media Server for their applications will reach a point where they have questions along the lines of the following: “how many users can connect to the Flash Media Server?”, “should I buy a Personal Edition or an Origin and Edge Edition?”, “how many servers will I need?”, etc.

In order to answer these questions (and more), it’s important that you understand three things:

- The capacity ceilings for the various editions of the Flash Media Server
- The total bandwidth that your application will require
- For heavily trafficked web sites and applications, you will also need to know performance characteristics of the Flash Media Server so you can make informed decisions about when it makes sense to add additional server hardware

## Flash Media Server Editions

To understand the various Editions of the Flash Media Server, their capacity limits, and how you can add more capacity, please read the Pricing and Editions page on Adobe.com. The URL to this specific page is:

<http://www.adobe.com/products/flashmediaserver/productinfo/pricing/>

## Bandwidth vs. Connections

As is documented on the Pricing and Editions page, the Flash Media Server contains two capacity ceilings: *a peak bandwidth per second* and a *number of simultaneous connections (users)* limit. The limit you hit first will depend on your application.

### ***Data-only Applications***

If your application only contains shared data (e.g., a text chat, multi-player game, or data streaming) and no audio or video, you will almost certainly reach the *number of simultaneous connections (users)* limit before you reach the *peak bandwidth per second* limit. In this case, calculating the number of simultaneous users your application can support is relatively straight forward - it is equal to the maximum number of simultaneous users the edition(s) you purchased can accommodate.

For example if you purchased the Professional Edition of Flash Media Server 2, you can have 2,500 simultaneous users using your applications. (Note: this assumes that each user uses exactly one connection and you do not have them connected to two different applications or the same application twice.)

### ***Audio-Video Applications***

If your application contains audio and/or video – as most Flash Media Server based applications do – you will almost certainly reach the *peak bandwidth per second* limit before you reach the *number of simultaneous connections (users)* limit. In order to calculate how many users you can accommodate with one Professional Edition in this case, you will need to calculate the amount of peak bandwidth you expect to flow through your system.

This is discussed in more details in the next section.

## **Calculating Bandwidth**

Developers that incorporate audio and/or video in their applications are well advised to develop a bandwidth strategy prior to building their applications. The purpose of the bandwidth strategy is to insure that the total amount of bandwidth served by the Flash Media Server and consumed by each client is sufficient to deliver the desired quality of service, is in line with the physical hardware limitations and software licenses you have, and can accommodate various levels of traffic. Like a good functional specification, a bandwidth strategy should highlight any problem areas upfront before systems are purchased and the application developed.

If you have not developed your bandwidth strategy yet however, this section will provide you with a method for estimating your bandwidth, and hence software license needs for the Flash Adobe Server.

### **Prerequisites**

The amount of peak bandwidth used by the Flash Media Server depends on a number of factors such as the bandwidth the users connect at, the bandwidth in your pipeline, the bandwidth of the audio and video streams, and of course, the type of application you develop (is it one-way streaming, or 2-way/multi-way chat application).

In order to estimate the bandwidth that your Flash Media Server application will use, you will need the following information:

- The number of simultaneous users that will be using your application at peak load (see Appendix A for tips on estimating this figure)
- The connection rate at which you expect your users to connect (e.g., what percent of your users will be on a 56k dial-up connection? On Cable/DSL? On LAN?)
- The number of audio/video streams that your application will have per user
- The limitations that you have (if any) on the server hardware and internet connection
- The target encode rate of the audio and video streams in your application

At this point you should also have a relatively firm idea of the functionality of your application and how users will be interacting with the application.

## Bandwidth Calculations

This section documents the methods for estimating bandwidth for two different types of applications – a one-to-many application such as a Video on Demand system and a many-to-many application such as a Video Conferencing application. For each example, a simple formula for calculating bandwidth is shown for both the server bandwidth (which is required to determine the number of software licenses you need) and the client bandwidth (which should be calculated as part of your bandwidth strategy to insure good quality of service to each client).

### *Example 1 - One to Many*

In this example, the application is a one-way video on demand system that streams recorded video at the user's request. In this example, the Flash Media Server will serve only one stream to each user.


#### Bandwidth Calculations

- Calculating *server* bandwidth needs ( $BW_s$ ):
 
$$BW_s = N * S$$
  - $N$  = Number of simultaneous users (subscribers)
  - $S$  = Stream encoded at constant kbps

- Calculating *client* bandwidth needs ( $BW_c$ ):

$$BW_c = S$$

- $S$  = A/V content encoded at constant kbps



### Example Calculation

Calculate the server bandwidth needed to stream video encoded at 100 kbps to 100 simultaneous users:

$$10 \text{ Mbps} = 100 \text{ kbps} * 100$$

This calculation can help you determine how many users you can get on a Professional Edition of the Macromedia Flash Media Server 2 under the various profiles. For example, using the 1,000 concurrent connections / 40Mbps profile, you would be able to get 400 simultaneous viewers assuming the stream was 100kbps.

Of course, if you want to stream to more than 400 users, you could either lower the quality of the video or buy and “stack” more Professional Edition licenses as explained on the Pricing and Editions page:

<http://www.adobe.com/software/flashcom/productinfo/editions/>

### Adjusting the Calculation for Multiple Bit rates

The calculations above assumed content encoded at constant bit rate. Most often, however, you will want to vary the bit rate of the content to suit the viewing audience. This will affect your bandwidth needs at both the client and server.

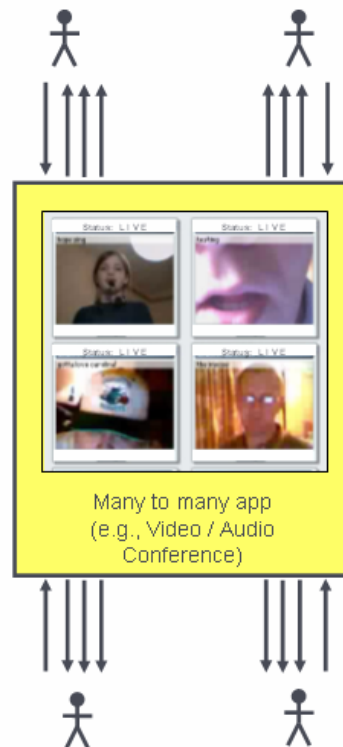
For example, let’s say that we estimate that half of the 100 simultaneous users were going to connect on a 56k dial-up modem and the other half would connect with DSL/Cable in the example above. Further, let’s say that while the video encoded at 100 kbps was appropriate for the DSL/Cable viewers, we wanted to encode a separate video at 38 kbps for the dial-up modem users.

In this case, the total bandwidth of the system is lowered to 6.9 Mbps:

$$6.9 \text{ Mbps} = (100 \text{ kbps} * 50) + (38 \text{ kbps} * 50)$$

### Example 2 - Many to Many

In this example, a number of people connect to a “conference room”. While in the conference room, each person will be broadcasting their own audio and video (e.g., via a webcam) while receiving the audio and video broadcasts from the others in the room. Calculating the bandwidth estimates here is very similar to the example above. However, the number of streams increases exponentially. In this case, the Media Server needs to serve X squared number of streams where X is the number of simultaneous users in a room.



For example, if there were 4 people in a room, person #1 would send (publish) one stream and receive three other streams for a total of four streams. Likewise, person #2, #3, and #4 would also consume four streams each. Therefore the total number of streams that the Media Server serves in this case is 16 – 4 people x 4 streams each, or thought about in a slightly different manner, 4 publishers of streams and 4 consumers or subscribers of streams.

### Bandwidth Calculations

- Calculating *server* bandwidth needs ( $BW_s$ ):

$$BW_s = (P * N) * S$$

→ P = Number of Publishers

→ N = Number of Subscribers

→ S = Stream encoded at a constant kbps

- Calculating *client* bandwidth needs ( $BW_c$ ):

$$BW_c = P * S$$

→ P = Number of Publishers

→ S = Stream encoded at a constant kbps

### Example Calculation

Server bandwidth needed for 4 person webcam chat @ 100 kbps streams:

$$1.6 \text{ Mbps} = (4 * 4) * 100 \text{ kbps}$$

Client bandwidth needed for 4 person webcam chat @ 100 kbps streams:

$$400 \text{ kbps} = 4 * 100 \text{ kbps}$$

So what is the answer to the “how many users can I get on a Professional Edition of the Flash Media Server 2?” in this example? If each room video conference generates bandwidth of 1.6 Mbps, then you can add 25 simultaneous rooms of four simultaneous people conferencing ( $40 \text{ Mbps} / 1.6 \text{ Mbps} = 25$ ) for a total of 100 (25 rooms x 4 people each) total participants on the 1,000conn/40Mbps Pro Edition Profile.

### No Limit on Number of Streams

The example above illustrates another important aspect of the Flash Media Server – there is no limit on the number of *streams* that can be served. Each user of your application (or said another way, one *connection*) can publish or subscribe to an unlimited number of streams as long as the total *peak bandwidth per second* limit is not reached.

### Bandwidth Considerations

Conducting calculations such as those above in your bandwidth strategy document will help you identify important bandwidth considerations for your application. Let’s explore two important considerations that are highlighted in this last many-to-many example.

### ***Client Connection Speed***

This should be evident, but worth stating here since we're on the topic. In the last many-to-many calculation example, note that even though it was a relatively simple application with a relatively small video stream, each client consumed 400 kbps of information. This is a lot of bandwidth – in fact, too much if your user is on dial-up modem or, in some cases, even on DSL. Users connecting to this application with connections less than 400 kbps will most likely encounter pauses for rebuffering and other poor quality effects.

Your bandwidth strategy document should help you deliver the correct amount of information to each client. Some other things to keep in mind concerning client connections:

- Save bandwidth for network overhead – A client's bandwidth will not only need to carry data from your application, but it will also need to carry traffic for normal network overhead. Reserve about 20% of a client's bandwidth for overhead when calculating the amount of data to send to each client.
- Some connections do not have the same upload and download speed. If your application requires interaction from the client (e.g., sending webcam video), make sure application takes their upload limitations into account.

You can find good tools for testing upload and download connection speeds on the internet. One useful test is here: <http://www.dslreports.com/tools>

### ***Application Design Matters***

In the last example, we estimated that would be able to add 100 total simultaneous participants to one Professional Edition (25 rooms x 4 people each).

What if we were to change the design of the application and instead of allowing four people to connect to one conference room, we only allowed two people to connect to a particular room? Or we only have one big conference room and everyone connects?

#### **Example 2a - Many rooms, few people**

In the scenario where we have many rooms, but limit each room to only two participants, the server bandwidth drops to 0.4 Mbps per room:

$$0.4 \text{ Mbps} = (2 * 2) * 100 \text{ kbps}$$

And we can now accommodate 400 total simultaneous participants across 100 rooms on one Professional Edition with the 1,000/40Mbps Profile. Note also that the client bandwidth needs drops to 200 kbps which opens the app up to more possible users – still not a good application for dial-up users, but most DSL users should now be able to use this application.

#### **Example 2b - One room, many people**

What about calculating the case where everyone joins one big conference room? Here, entering 40 Mbps as server bandwidth, assuming the number of publishers = number of subscribers = N, and solving for N:

$$40 \text{ Mbps} = (\mathbf{N} * \mathbf{N}) * 100 \text{ kbps}$$

You see that you can get 20 simultaneous participants on one Professional Edition with this application (assuming that each participant has a 2 Mbps internet connection).

These examples illustrate that the design of the application plays a very important role in bandwidth utilization and, ultimately, in the end-user experience.

By now, you should have a good idea of how to calculate the server bandwidth requirements for your application. What if your estimates require the server to serve up 200 Mbps of audio/video? Is it realistic to push this much a/v bandwidth on one physically server machine? Two servers? Three servers?

## Server Performance

Once you have your bandwidth strategy in place, you're going to need to determine what and how much server hardware you will need to run your application. Applications that involve video and audio tend to be very CPU intensive. Once the server's CPU becomes stressed, video and audio packets may be lost, or pauses for rebuffering may occur, and your quality of service may degrade. Therefore, it's important to understand the load that your servers can effectively handle so you can determine the appropriate number of servers you'll need to meet the traffic demands of your application.

So, how much traffic can you put on a Flash Media Server before the CPU starts to become stressed? A couple of factors play a very important role in determining this - the server hardware obviously, the type of application you develop (is it one-way streaming, or 2-way/multi-way chat application), and the connection at which users connect. Some of our findings and comments are during testing of the Flash Media Server are listed below:

### Recommended Server Hardware

The number of CPUs and amount of RAM are important factors in determining the server performance. For servers which will be used in deployment, we believe the best price/performance server hardware for running Macromedia Flash Communication Server 1.5 is the following:

- Dual Intel® Xeon™ Processors, at least 2.0GHz
- At least 1.0 GB RAM
- 1 Gigabit Ethernet card

While there is some gain in performance by having more CPUs per server, the increased cost of 4+ CPU servers may not generally be justified.

## **Test Findings for Various Applications and Client Connection Speeds**

The type of application is very important because applications use CPU resources differently. A simple text chat application, for example, will most likely use very little CPU resources. You could add tens of thousands of users to your application before you start seeing any significant CPU stress on the server. Applications that use video, however, may start to show CPU stress at hundreds of users.

The next sections below provide a general summary of our test results for a number of different Flash Media Server applications. The test results listed below were based on a server with a dual Pentium 3 1.1 Ghz processor, 1 GB RAM, Windows 2000 Server, and a 1 Gigabit Ethernet card. The server was connected to clients on an isolated network.

### ***Data Applications***

In our testing of a simple data sharing application, we have added over 17,000 clients to a data sharing application without seeing any significant CPU stress. While we did not test this application for its maximum condition, we felt that the maximum condition in this case would mostly likely be constrained by the number of socket connections permitted by the operating system rather than by server CPU stress. The data shared amongst the clients in this specific test used very little processing power. If your application has a high level of data sharing you may find different results.


### ***Audio and Video Applications***

In our testing of audio/video applications, we've found that server performance is very dependent on the type of application, as expected.

To help you try to estimate the amount of load a single server might be able to handle for your application, we've provided a summary of test results for three different types of applications:

- Recorded A/V – the test application served a recorded audio and video file (flv file) to a client upon request. It is characteristic of a Video on Demand type of application.
- Live A/V – the test application served a live audio and video stream to each client when they connected to the application. It is characteristic of a live event broadcast or a live virtual classroom setting where audio and video are being streamed live from one user to many users.
- Live Conference - the test application connected four (4) clients to a conference room. Each client published one live stream of audio and video and received three live audio and video streams from the other clients similar to Example 2 in the Bandwidth Calculation section above. It may be characteristic of a web or video conference or n-way online webcam chat application for a dating or community website.

The figures below provide an overview of the *maximum amount of bandwidth Flash Media Server can handle* and the *maximum bandwidth at which clients can connect* for each of these three applications when run on a dual Pentium 3 1.1 Ghz processor, 1.2 GB RAM, Windows 2000 Server, and a 1 Gigabit Ethernet card. These tests were performed on an Origin Edition license to the Flash Media Server.



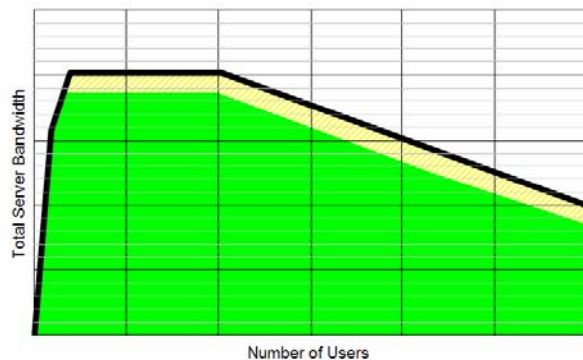
**Figure 1:** Maximum server bandwidth vs. Number of connected users

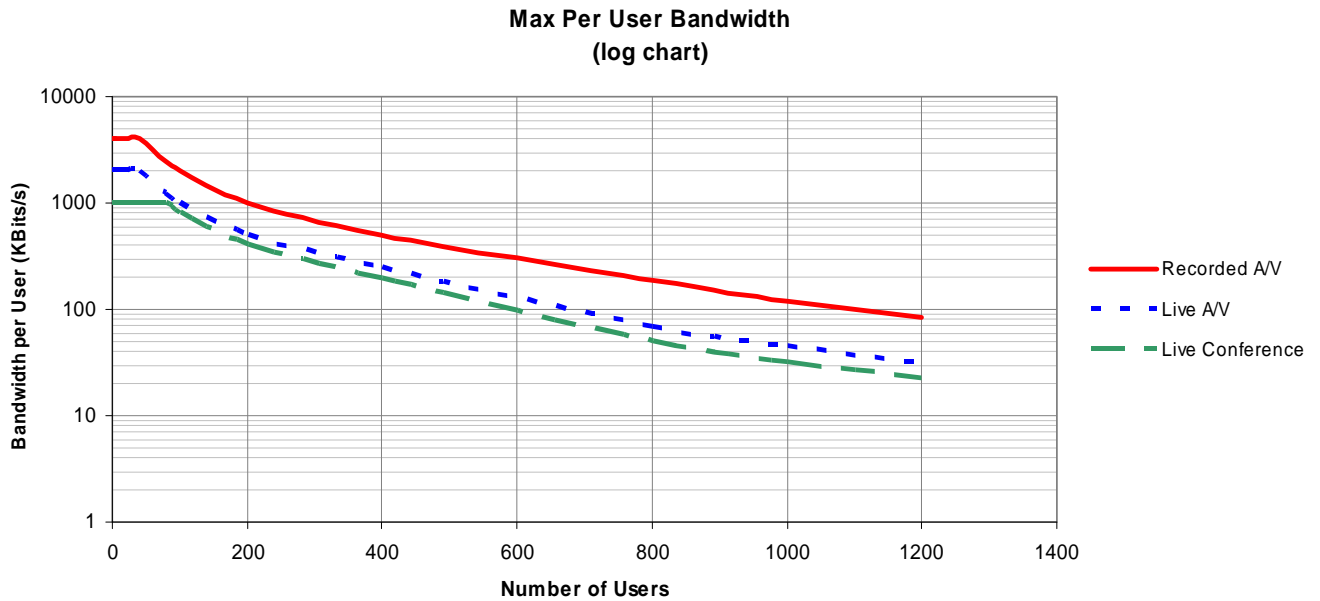
The curves on the Maximum Total Server Bandwidth figure above correspond to the maximum amount of bandwidth per a given number of users that Flash Media Server is capable of handling on our dual P3 test machine.

When developing the bandwidth strategy for your application, you should insure that the total bandwidth served up by the Flash Media Server is below the particular curve corresponding to the type of application you are building.

A depiction of a safe operating zone for a recorded A/V application is shown in the solid green area in the image on the right.

You should also leave a buffer, as indicated by the yellow hatched area in the image to the right, to insure you can accommodate any bandwidth surges.





**Figure 2:** Maximum bandwidth per user vs. Number of connected users

The curves on the Maximum Bandwidth per User figure above correspond to the maximum amount of bandwidth the Flash Media Server will be able to send to clients (measured per client) given the total number of users connected to the application.

Again, when developing the bandwidth strategy for your application, you should insure that the bandwidth per user served up by the Flash Media Server is below the particular curve corresponding to the type of application you are building.

In general, in our testing with applications that use video and audio, we've found that, as might be expected, serving recorded audio/video performs better than serving live audio/video.

Furthermore, we found that clients consuming larger amounts of data (high bandwidth clients) are more efficient than clients consuming smaller amounts of data (low bandwidth clients).

## Conclusion

By now you should have all the information you need to estimate the bandwidth utilization of your application, the appropriate Media Server licenses that you will need and some ballpark information about when you can expect to start thinking about adding another server as your traffic grows.

The last, but most important, piece of advice is to be sure you thoroughly test your application yourself. This whitepaper provides a lot of good information, but ultimately the only way to know for sure how what type of server resources your application will take is to test it in real world circumstances.

## Appendix A - Estimating Number of Simultaneous Users

As mentioned in the Calculating Bandwidth section of this whitepaper, accurate knowledge about the *number of simultaneous users that will be using your application at peak load* is one of the things that you will need in order to calculate the bandwidth requirements of your Flash Media Server application.

### What are Simultaneous Users?

Most companies do not have accurate estimate of the number of simultaneous users on a particular portion of their web site. In general, web stat logging software typically does not provide this much granularity. In fact, if your client tells you that they expect a hundred thousand, a half of million, or even a million or more simultaneous users, it is probably time to raise a red flag. Unless you are building the next version of Yahoo's IM, MSN's IM, or AOL's IM, your web site will not see this amount of simultaneous traffic. Chances are your client was talking about the amount of total page views or hits over the course of some time frame.

For this reason, it is important to understand what we are looking for. The *number of simultaneous users that will be using your application at peak load* is the maximum amount of users that will be using your application *at the same time*. It is not the total number of users who will ever use your application. It is not the number of people who have used your application in a single day. The *number of simultaneous users that will be using your application at peak load* will likely be much lower than the total number of users of your application.

### How to Estimate Number of Simultaneous Users

Since this the number of simultaneous users is such an important piece of information for determining the number of Flash Media Server software licenses and the bandwidth capacity you will need for your application, this section walks you through one way of estimating this number when the only information you have is very general web page statistics.


We'll use a hypothetical case: The ACME Company intends to launch a 10 minute marketing video linked off the home page of their web site for 20 days to promote an upcoming product launch. They have stats from their last product launch that showed 150,000 people viewed the promotional materials during that 20 day campaign. In this example, we also know that almost 50% of this traffic occurred in the first 5 days of the campaign.

To find an estimate for the number of peak simultaneous users, we'll try to find the single hour in which the most users viewed the video and then make an assumption that  $1/6^{\text{th}}$  of the visitors in that hour will be watching the video simultaneously. The  $1/6^{\text{th}}$  factor is a simplifying assumption based on the fact that the viewers will be distributed evenly in that peak hour and each will be viewing the video for only 10 minutes of a possible 60 minutes in the hour. Therefore at any one point in the hour, on average, only  $1/6^{\text{th}}$  of all of the viewers in that hour would have been viewing the video simultaneously (given a random distribution).

In the example above, since we don't know how many visitors our site will have during the peak hour, we'll need to make some more assumptions.


First, since we know that about 50% of the traffic will occur during the first five days, let's assume that there is a spike in traffic during the first five days, and then traffic gradually decreases. With these assumptions, and the fact that we expect 150,000 viewers over 20 days with about 50% of those in the first five days, we can use Microsoft Excel to model up an estimate of daily web traffic such as that shown below.

Day	# of viewers	% of traffic
Day 1	10,000	6.7%
Day 2	14,000	9.3%
Day 3	16,000	10.7%
Day 4	16,250	10.8%
Day 5	15,000	10.0%
Day 6	13,000	8.7%
Day 7	10,250	6.8%
Day 8	8,750	5.8%
Day 9	7,500	5.0%
Day 10	6,500	4.3%
Day 11	5,750	3.8%
Day 12	5,000	3.3%
Day 13	4,500	3.0%
Day 14	4,000	2.7%
Day 15	3,500	2.3%
Day 16	3,000	2.0%
Day 17	2,500	1.7%
Day 18	2,000	1.3%
Day 19	1,500	1.0%
Day 20	1,000	0.7%
	150,000	100.0%



In this example, our model would suggest that the day with the largest amount of viewers would be Day 4 with 16,250 viewers. The next step is to find the hour with the greatest amount of expected viewers. Making the assumption that this marketing campaign will be aimed predominately at US audiences, we can use Excel to develop a similar model of estimated traffic per hour during the day such as that shown below.

Time	Viewers	% of traffic
midnight - 1AM PST	163	1%
1AM - 2AM PST	163	1%
2AM - 3AM PST	-	0%
3AM - 4AM PST	163	1%
4AM - 5AM PST	163	1%
5AM - 2AM PST	163	1%
6AM - 7AM PST	488	3%
7AM - 8AM PST	813	5%
8AM - 9AM PST	1,138	7%
9AM - 10AM PST	975	6%
10AM - 11AM PST	1,463	9%
11AM - 12noon PST	1,950	12%
12 noon - 1PM PST	1,788	11%
1PM - 2PM PST	1,138	7%
2PM - 3PM PST	1,138	7%
3PM - 4PM PST	813	5%
4PM - 5PM PST	650	4%
5PM - 6PM PST	325	2%
6PM - 7PM PST	488	3%
7PM - 8PM PST	650	4%
8PM - 9PM PST	325	2%
9PM - 10PM PST	325	2%
10PM - 11PM PST	488	3%
11PM - midnight PST	488	3%
	16,250	100%



The distribution of traffic into the various hours as shown above roughly models the traffic we've seen for certain videos deployed on adobe.com. Using this model with our ACME Company example we've now identified that the single hour during the entire campaign (20 days) with the most viewers will be from 11AM - 12noon PST on Day 4 and that there will be 1,950 viewers at this time.

We can now multiply the 1,950 viewers during the peak hour by 1/6 to get the estimate for the peak number of simultaneous users. In this case, our estimated number of peak simultaneous viewers of this video will be 325.

Just to complete this example and determine the bandwidth and license requirements, let's assume that we expect half of the users to be on dial up connections and the other half on DSL or LAN. Further, we'll assume that we'll offer the dial up users video encoded at 38kbs and offer the broadband users video encoded at 150kpbs. Our total server bandwidth required would be about 33.5 Mbps [ = (.038 \* (325/2)) + (.150 \* (325/2)) ] and should be able to be handled with 1 Professional Edition of Flash Media Server 2 with the 1,000conn/40Mbps Profile configured.