

Building a Performance Model of Streaming Media Applications in Utility Data Center Environment

Ludmila Cherkasova
Hewlett-Packard Laboratories
Internet Systems and Storage Lab
1501 Page Mill Road
Palo Alto, CA 94303, USA
cherkasova@hpl.hp.com

Loren Staley
Hewlett-Packard Co
Always On Infrastructure Solutions
19111 Pruneridge Ave
Cupertino, CA 95014, USA
loren_staley@hp.com

Abstract

Utility Data Center (UDC) provides a flexible, cost-effective infrastructure to support the hosting of applications for Internet services. In order to enable the design of a “utility-aware” streaming media service which automatically requests the necessary resources from UDC infrastructure, we introduce a set of benchmarks for measuring the basic capacities of streaming media systems. The benchmarks allow one to derive the scaling rules of server capacity for delivering media files which are: i) encoded at different bit rates, ii) streamed from memory vs disk. Using an experimental testbed, we show that these scaling rules are non-trivial. In this paper, we develop a workload-aware, media server performance model which is based on a cost function derived from the set of basic benchmark measurements. We validate this performance model by comparing the predicted and measured media server capacities for a set of synthetic workloads.

1 Introduction

A Utility Data Center (UDC) [13, 14] provides a flexible, cost-effective solution by using advanced management software which allows resources to be automatically reassigned in response to changing business and IT requirements. The UDC architecture is an ideal platform to support the efficient hosting of applications for Internet services. Most Internet applications are difficult to manage due to their highly variable service demand. The most typical practice used by current service providers is to significantly over provision the amount of resources needed to support such applications by tailoring the resources to maximum expected load. The UDC infrastructure provides a set of new management capabilities for requesting/releasing the system resources to dynamically provision the application demands.

The area of multimedia services in a networked environment is one of the rapidly expanding fields in today’s technological world. The delivery of continuous media from a central or distributed server cluster to a large number of clients is a challenging and resource intensive task. The evidence from media workload analysis [1, 2] indicates that client demands are highly variable: the “peak-to-mean” ratio may be an order of magnitude or higher. Additionally, media applications are characterized by stringent real-time constraints, as each stream

requires isochronous data playout. These features make media application a perfect candidate to benefit from flexible resource management and resource provisioning in the Utility Data Center infrastructure.

We are pursuing a vision where large-scale hosting infrastructures simultaneously provide “resource-on-demand” capabilities to competing Internet services by using statistical multiplexing in order to dynamically satisfy the requirements of services with highly variable loads. Our ultimate goal is to design and implement a “utility-aware” streaming media service which automatically requests the necessary resources from the UDC infrastructure in order to adapt to variable workload demand. Among the *prerequisites* to a design of such a utility service is the ability to measure the capacity of media servers (nominal and currently in use) in order to evaluate the amount of available system resources for admitting new client requests or to request/release resources as needed.

In this paper, we develop a *workload-aware performance model of streaming media applications* which lays down a foundation of our utility-aware service design. In order to build such a performance model several classic performance questions should be answered:

- how to measure the *basic* capacity of a streaming media server?
- what is the set of *basic* benchmarks exposing the performance limits and main bottlenecks of a media server?
- how to estimate the expected media server capacity for realistic workload if the measured capacities of streaming media server under the basic benchmarks are given?

Commercial media server solutions are distinguished by the number of concurrent streams a server can support [9] without losing a quality of stream, i.e. until real-time constraint of each stream can be met. A standard commercial stress test measures a maximum number of concurrent streams delivered by the server when all the clients are accessing the same file encoded at a certain bit rate, e.g. 500 Kb/s.

However, a multimedia content is typically encoded at different bit rates depending on a type of content and a targeted population of clients and their connection band-

width to the Internet. What are the *scaling rules for server capacity* when delivered media content encoded at different bit rates? For example, if a media server is capable of delivering N concurrent streams encoded at 500 Kb/s, will this server be capable of supporting $2 \times N$ concurrent streams encoded at 250 Kb/s? The other issue with a standard commercial stress test is that all the clients are accessing the same file. Thus another question to answer is: how a media server performance is impacted when different clients retrieve different (unique) files of a media content?

We introduce a simple set of *basic benchmarks* to establish the scaling rules for server capacity when delivered streams are encoded at different bit rates:

- *Single File Benchmark*;
- *Unique Files Benchmark*.

The measurement results reported in the paper show that these scaling rules are non-trivial.

There are different system resources limiting a media server performance under the different basic benchmarks. Under the *Single File Benchmark*, the server performance is CPU bounded, while under the *Unique Files Benchmark*, the server performance is disk-bound. A typical media server traffic presents a combination of some clients accessing the same files (during the same period of time) and some clients retrieving the “unique”, different files. Under the “mixed” workload, neither a CPU nor a disk are the limiting performance system resources, and hence, they can not be used for monitoring of the available server capacity. In order to manage streaming media service efficiently, we need to be able to estimate:

- what is a currently available/used service capacity?
- what fraction of service capacity is required to support a particular stream without a degradation of its quality?

In this paper, we develop a *workload-aware performance model* of media server aiming to estimate the media server capacity for delivering a realistic workload. This model is based on a *cost* function derived from the set of basic benchmark measurements. The cost function defines a *fraction* of system resources needed to support a particular media stream depending on the stream bit rate and type of access (memory file access or disk file access). We validate this performance model by comparing the predicted and measured media server capacities for a set of synthetic workloads.

In summary, the main contribution of the paper is a new unified framework for

- measuring a media service capacity via a set of basic benchmarks;
- deriving the resource requirements (a cost) of a particular stream from the basic benchmarks measurements;
- estimating a currently available/used service capacity under the realistic workloads via a new workload-aware performance model based on a cost function.

The remainder of the paper presents our results in more detail.

2 Experimental Testbed

Figure 1 shows our experimental testbed used for measuring the media system performance:

- a server is rp2450 system with 1-way 550MHz PA 8600 processor, memory – 4 GB, IO Cards – 2 1000SX Ethernet, 4 Ultra2 SCSI port on 2 dual port cards, disk configuration – 4x15K rpm 18.2 GB disks 4-way striped using HP-UX LVM, HFS file system with block size configured to 64KB; operating system – HP-UX 11.0,
- 14 clients machines are rp2450 systems with 2-way 550MHz PA 8600 processors, running HP-UX 11.0 operating system,
- 14 clients machine are connected to a server by a Gbit switch 4108gl via 1Gb/s links,
- a media server software: RealServer 8.0 from RealNetworks[9].




Figure 1: Experimental setup.

The configuration and the system parameters of our experimental setup are specially chosen to avoid some trivial bottlenecks when delivering multimedia applications: such as limiting I/O bandwidth between the server and the storage system, or limiting network bandwidth between the server and the clients. Such system configuration exposes the performance limits specific to application itself, and allows us to derive more general conclusions about its performance. Our experimental setup uses the general purpose components available in Utility Data Center environment.¹

For the load tests, the client emulators are instructed to start a predefined number of streams requesting a certain set of files. The emulator software has a flexibility allowing different clients to request different files. Additionally, different emulators may be configured with their own specific parameters as for the number of streams as for a set of requested files.

¹Typically, UDC uses a SAN-based storage solution. A technique and approach proposed in the paper is not limited to a particular configuration of the experimental testbed used in the study: it allows one to measure a performance of multimedia applications in a very general setting.

3 Set of Basic Benchmarks and Metrics

In this section, we introduce a set of *capacity metrics* and two *basic benchmarks* aimed at determining the performance limits of media server caused by different system performance bottlenecks. Both of the benchmarks introduced below use specially created 20 min video clips encoded at different bit rates with RealProducer G2 from RealNetworks representing the typical needs of the Internet audience:

- 28 Kb/s for analog modem users,
- 56 Kb/s for analog modem and ISDN users,
- 112 Kb/s for dual-ISDN users,
- 256 Kb/s for cable modem users,
- 350 Kb/s for DSL/cable users,
- 500 Kb/s for high-bandwidth users.

The primary objective of the basic benchmarks is to define how many concurrent streams of the *same bit rate* can be supported by a media server without degrading the quality of any streams.

3.1 Capacity Metrics: Server Overload and Degraded Stream Quality

In our experimental setup, we use two main performance metrics to determine whether a server has reached its maximum capacity for the applied workload. One of these metrics is collected at the media server side and the other one - at the client side. These metrics are complementary. Typically, the server side metric reliably reflects the server state and can be used alone to determine whether the server capacity for applied workload is reached.²

Server Side Metric: *Server Overload Metric*. RealServer can be configured to report a set of useful statistics (*ServerStats*) about its current state. The default value for a reporting interval is 3 sec, we use a 10 sec value in our experimental setup. We list here only a few statistics from *ServerStats* which we observe during the stress tests for server monitoring purpose.

- the number of streams currently playing by the server,
- the aggregate bandwidth requirements of currently accepted streams,
- the average bandwidth delivered by the server during the reporting interval,
- the number of packets being sent with violation of the real-time constraints: an alarming information about a server being in *overload state*.

To determine whether the server capacity is reached, we use a *server overload metric* providing the information about the packets sent with violation of “on-time delivery”. These packets are indicative of a server being overloaded and that its available capacity is exceeded.

²In case, when the network bandwidth in the measurement testbed is a limiting performance bottleneck, the combined information from server-side and client-side metrics helps to reveal this bottleneck.

Client Side Metric: *Degraded Stream Quality Metric*. On a client side, we are interested in observing whether a client has entered a *rebuffering state* which means:

- the client has stopped playing because the current “play” buffer is empty,
- the client waits until the “play” buffer is filled to acceptable level to continue play back.

We use a specially written software, called *ClientStats*, to observe the number of *rebuffering events* and the statistics on the average stream bandwidth received by the client. In our experimental setup, where communication infrastructure between the server and the clients is not a limiting resource, the existence of rebuffering events on the client side reflects a degrading quality of the delivered streams as a results of the server being overloaded.

3.2 Single File Benchmark

The *Single File Benchmark* measures a media server capacity when all the clients in the test are accessing the same file. We designed a completely automatic benchmark which runs a sequence of tests with an increasing number of clients requesting the same file from the server for each of the six encoding bit rates. During each test, we collected *ServerStats*, *ClientStats*, general performance information of a server machine and the client machines using *vmstat*. To make sure that none of the video clip data are present in the file buffer cache (memory), and a file is streamed from the disk, a file system is unmounted and mounted back before each test point.

The main goal of the basic benchmarks is to define how many concurrent streams of the *same bit rate* can be supported by the media server without degrading the quality of streams.

While the maximum achievable media server capacity depends on the encoding bit rate of the file, the performance data collected at the server and the clients are similar. Using two representatives, files encoded at 112 Kb/s and 350 Kb/s, we show more details about server performance, its limits and bottlenecks under the *Single File Benchmark*.

For a file encoded at 112 Kb/s, a media server capacity is reached at 1350 concurrent streams, and for a file encoded at 350 Kb/s – at 500 concurrent streams. Thus we observe that the media server capacity scales differently than the encoding bit rates of the underlying files: the difference in encoding bit rates between 112 Kb/s file and 350 Kb/s file is 3.125, while the difference in server capacity for corresponding bit rates is only 2.7. Intuitively, media server incurs an additional overhead while handling a **higher number** of concurrent streams encoded at lower bit rate comparing with handling a **smaller number** of concurrent streams encoded at higher bit rate (even for a case, when the aggregate bandwidth requirements of the delivered streams are the same).

Figures 2 a) and 3 a) show the CPU utilization and the application related part of it for corresponding file bit rates. Under the *Single File Benchmark*, the media




Figure 2: *Single File Benchmark* for a) 112 Kb/s tests: a) Server utilization, b) Server bandwidth, c) Stream bandwidth.




Figure 3: *Single File Benchmark* for a) 350 Kb/s tests: a) Server utilization, b) Server bandwidth, c) Stream bandwidth.




Figure 4: *Single File Benchmark*: a) Server Capacity, b) Maximum Delivered Bandwidth c) Server Capacity Scaling Rules.

server is CPU bounded. The shape of the CPU utilization curve is not linear for both encoding bit rates. During the first half of the server capacity load, the CPU utilization grows linearly and reaches nearly 85%. The second half of the server capacity is achieved by utilizing the remaining 15% of the CPU resources, and finally, when the CPU utilization reaches 100%, it becomes a limiting factor of the media server performance.

Often, the utilization of a bottlenecked resource is used to build a simple normalized performance model of a corresponding system or application. For example, a system capacity achievable at 100% utilization of a bottleneck resource represents a 100% of applied load. Then at 50% utilization of a limiting resource a system can process a 50% of an applied load (i.e. 50% utilization of a bottleneck resource corresponds to 50% of server capacity). The shape of the CPU utilization curve in Figures 2 a) and 3 a) overwrites an existence of such simple performance model for streaming media server (even for a case, when a workload is delivered from memory). We would like to stress this observation in light of a performance model developed in Section 4.

Figures 2 b) and 3 b) show the overall bandwidth in

Mb/s delivered at the server side. It grows linearly, until the server capacity is reached, and after that bandwidth delivered by the server flattens. For a *Single File Benchmark* and a file encoded at 112 Kb/s, the bandwidth delivered by a server peaks at 151.2 Mb/s, and for a file encoded at 350 Kb/s, the maximum delivered bandwidth is 175 Mb/s.

Figures 2 c) and 3 c) present a complementary information from the client side: once a server capacity is reached, the quality of the streams starts to degrade significantly: 10% of extra load causes around 9% of stream bandwidth degradation.

The set of graphs in Figures 2 b) and 3 b) helps to predict the quality of service implications under server overload condition. Typical media server application does not have a “built-in” admission control mechanism. Since a maximum server capacity defines the maximum bandwidth a server can deliver at a particular bit rate, any accepted load in addition to a basic server capacity results in degradation of delivered stream quality: $X\%$ of additional load causes $\frac{100 \times X}{(100 + X)}$ % of degradation in delivered stream bandwidth. Thus, measurements of media server performance under a set of basic workloads provide

useful insights into the specifics of application behavior under overload conditions which can be used in design of an efficient admission control strategy for media servers. Utility-aware services typically require a certain level of quality guarantees. Understanding the relationship between the server overload and QoS of delivered streams is essential for implementing such guarantees.

Figure 4 a) presents the maximum capacity in concurrent streams achievable by the streaming media server across six different encoding bit rates under the *Single File Benchmark*. Figure 4 b) shows the corresponding maximum bandwidth in Mb/s delivered by the media server for different encoding bit rates. The shape of the curve is interesting: for higher encoding bit rates the difference in achievable bandwidth is much less significant than for lower bit rates. Many admission control strategies designed in literature use the “fixed” maximum bandwidth a server is capable to deliver as a main “scheduling” resource for admission of a new stream. Evidently, the amount of bandwidth a server is capable to deliver is variable and depends on the encoding bit rates of current streams in progress.

Figure 4 c) shows the normalized graph reflecting the scaling rules for the media server capacity under the *Single File Benchmark* and different encoding bit rates. In this figure, point (1,1) presents the maximum capacity (360 concurrent streams) achievable by a server when all the clients in the test are accessing the same file encoded at a 500 Kb/s bit rate. Each absolute value for the other encoding bit rate is normalized with respect to it. For example, the maximum capacity achievable by a server under *Single File Benchmark* and a 28 Kb/s encoding bit rate is 3300 concurrent streams and is represented by a point (17.9, 9.2).

Figure 4 c) reflects that the server capacity scaling rules are non-trivial. For example, the difference between the highest and lowest bit rate of media streams used in our experiments is about 18 times. However, the difference in maximum number of concurrent streams a server is capable of supporting for corresponding bit rates is only around 9 times. Figure 4 c) shows that a media server capacity scales non-linearly comparing to the encoding bit rates of underlying files.

In summary, under the *Single File Benchmark* only one stream reads a file from the disk, while all the other streams read the corresponding bytes from the file buffer cache. Thus, practically, this benchmark measures a streaming server capacity when the media content is delivered from memory. However, it is not necessary that the streamed media file completely resides or fits in the memory. In essence, this benchmark exercises the shared access by multiple clients to the same file.

3.3 Unique Files Benchmark

The *Unique File Benchmark* measures the media server capacity when each client in the test is accessing a different (unique) file.

Similarly, for the *Unique File Benchmark*, we designed

a completely automatic benchmark for all of the six encoding bit rates, which runs a sequence of tests with an increasing number of clients, where different clients request different files from the server. For a test point with N different clients, each file from the original basic set of files is replicated N times correspondingly, where different copies of the file are given different file names. We used the LVM features of HP-UX to create a striped volume across the 4 disks in our experimental setup. To make sure that none of the video clip data are present in the file buffer cache (memory), and all the files are streamed from the disk, a file system is unmounted and mounted back before each test point.

As intuitively expected, under the *Unique File Benchmark*, the CPU is not a resource which limits server performance: for all the tests in this study, it is below 45%. The server performance under *Unique Files Benchmark* is disk-bound: this particular bottleneck is hard to measure with the usual performance tools. The maximum bandwidth delivered by a disk depends on the number of concurrent streams it can support with an acceptable level of *jitter*, i.e. without violating on-time delivery constraints.

The measurements reveal a similar server behavior under overload for the *Unique Files Benchmark* as was observed for the *Single File Benchmark* in Section 3.2 in Figures 2 b, c) and 3 b, c). This reflects that a media server has a typical behavior under overload conditions across different workload types.

Figure 5 a) presents the maximum capacity in concurrent streams achievable by the streaming media server across six different encoding bit rates under *Unique File Benchmark*. Figure 5 b) shows the normalized graph reflecting the scaling rules for the media server capacity under the *Unique File Benchmark* and different encoding bit rates. It shows that the capacity scaling rules under *Unique File Benchmark* are non-trivial and are different from the capacity scaling rules under *Single File Benchmark*. Finally, Figure 5 c) compares achievable media server capacity for *Single* vs *Unique File Benchmark*. The media server performance is 2.5-3 times higher under the *Single File Benchmark* than under the *Unique Files Benchmark*. These results quantify the performance benefits for multimedia applications when media streams are delivered from memory as well as suggest that much higher performance can be achieved for workloads which exhibit a high degree of shared client accesses to the same media files.

In summary, under the *Unique Files Benchmark* all the streams read their files from disk. Thus, practically, this benchmark measures a streaming media server capacity when a media content is delivered entirely from the disk. The performance results under such workload type are highly dependent on a choice of a file system and disk configuration. Service providers might use the proposed *Unique Files Benchmark* for performance evaluation of different components available in UDC infrastructure to make the right choices for the best performing system configurations.




Figure 5: a) Server Capacity under *Unique Files Benchmark*, b) Server Capacity Scaling Rules under *Unique Files Benchmark*, c) Comparison of Server Capacity: *Single File vs Unique Files Benchmark*.

4 Workload-Aware Performance Model of Streaming Media Server Capacity

In Section 3, we showed how to measure the basic capacity of a streaming media server using the set of basic benchmarks. The main goal of this section is to answer a question: how to compute the expected media server capacity for realistic workload if the measured capacities of streaming media server under the basic benchmarks are given.

In this section, we develop a workload-aware server capacity model which is based on a *cost* function derived from the set of basic benchmark measurements. Intuitively, the cost function defines a *fraction* of system resources needed to support a particular stream depending on the encoding bit rate of the corresponding file and access type to a corresponding file.

With each stream delivered by a server we associate an access type:

- *memory access* if a stream retrieves a corresponding file (or the corresponding bytes of the file) from the memory;
- *disk access* if a stream retrieves a corresponding file from the disk.

Additionally, we use the following notations:

- $X = X_1, \dots, X_k$ - a set of encoding bit rates of the files used in basic benchmarks,
- $N_{X_i}^{single}$ - the maximum measured server capacity in concurrent streams under the *Single File Benchmark* for a file encoded at X_i Kb/s,
- $N_{X_i}^{unique}$ - the maximum measured server capacity in concurrent streams under *Unique Files Benchmark* for a file encoded at X_i Kb/s,
- $cost_{X_i}^{memory}$ - a value of cost function for a stream with memory access to a file encoded at X_i Kb/s,
- $cost_{X_i}^{disk}$ - a value of cost function for a stream with disk access to a file encoded at X_i Kb/s.

Under the *Unique Files Benchmark*, all the streams have a disk access type, because all the streams read the corresponding files from the disk. Hence each stream requires

a fraction of system resources defined by the $cost_{X_i}^{disk}$ value.

Under the *Single File Benchmark*, one stream has a disk access type (it reads the corresponding file from disk, and hence requires a fraction of system resources defined by the $cost_{X_i}^{disk}$ value) while the rest of the streams retrieve a corresponding file from the memory and, therefore, have a memory access type and require a fraction of system resources defined by the $cost_{X_i}^{memory}$ value.

Let 1 be a media server capacity. The following *capacity equations* describe the maximum server capacity measured under a set of basic benchmarks for each encoding bit rate $X_i \in X$:

$$N_{X_i}^{unique} \times cost_{X_i}^{disk} = 1$$

$$1 \times cost_{X_i}^{disk} + (N_{X_i}^{single} - 1) \times cost_{X_i}^{memory} = 1$$

By solving the equations above, we can derive the corresponding cost function values:

$$cost_{X_i}^{disk} = \frac{1}{N_{X_i}^{unique}}$$

$$cost_{X_i}^{memory} = \frac{N_{X_i}^{unique} - 1}{N_{X_i}^{unique} \times (N_{X_i}^{single} - 1)}$$

Let W be a current workload processed by a media server, where

- $X_w = X_1, \dots, X_{k_w}$ - a set of encoding bit rates of the files used in W ($X_w \subseteq X$),
- $N_{X_{w_i}}^{memory}$ - a number of streams having a memory access type to a subset of files encoded at X_{w_i} Kb/s,
- $N_{X_{w_i}}^{disk}$ - a number of streams having a disk access type to a subset of files encoded at X_{w_i} Kb/s.

Then the applied *Load* to a media server under workload W can be computed by a formula:

$$Load = \sum_{i=1}^{k_w} N_{X_{w_i}}^{memory} \times cost_{X_{w_i}}^{memory} + \sum_{i=1}^{k_w} N_{X_{w_i}}^{disk} \times cost_{X_{w_i}}^{disk}$$

If $Load \geq 1$ then the media server is overloaded and its capacity is exceeded. The difference $Load - 1$ defines




Figure 6: Measured vs Expected Server Capacities for a *single-unique-one-bit-rate* workload: a) $\alpha = 1/3$, b) $\alpha = 1/2$, and c) $\alpha = 2/3$.

the amount of overload or server exceeded capacity. If $Load \leq 1$ then the media server operates within its capacity, and the difference $1 - Load$ defines the amount of available server capacity.

We validated this performance model of media server by comparing the predicted (computed) and measured media server capacities for a set of different synthetic workloads:

- *single-unique-one-bit-rate*,
- *single-six-bit-rates*,
- *unique-six-bit-rates*.

Under *single-unique-one-bit-rate* workload, all the clients are accessing the files encoded at the same bit rate. Let the encoding bit rate be X_i Kb/s. Let the number of clients accessing the unique files and the number of clients accessing the same (single) file be defined as following:

- $\alpha \times N_{X_i}^{unique}$ clients are accessing the unique files, where $\alpha \leq 1$ and $N_{X_i}^{unique}$ is a measured server capacity in concurrent streams under *Unique Files Benchmark* for files encoded at X_i Kb/s,
- $(1 - \alpha) \times N_{X_i}^{single}$ clients are accessing the same file where $N_{X_i}^{single}$ is a measured server capacity in concurrent streams under *Single File Benchmark* for a file encoded at X_i Kb/s.

The number of clients (or concurrent streams) in this workload is designed in a special way: if our performance model of server capacity is correct, then under the *single-unique-one-bit-rate* workload the server maximum capacity should be reached. We call it *expected (computed) server capacity*.

Using our experimental testbed, we ran *single-unique-one-bit-rate mix* workload for $\alpha = 1/3$, $\alpha = 1/2$, and $\alpha = 2/3$ and six bit rates of interest. In the workloads under test, we fixed the number of clients accessing the unique files accordingly to a formula defined above and slightly varied the number of clients accessing the same file to determine experimentally when the server capacity under test is reached using the capacity metrics defined in Section 3.1.

Figure 6 shows the measured vs expected server capacities under *single-unique-one-bit-rate* workload. The measured server capacity closely matches the expected

media server capacity across different encoding bit rates and different values of α (the error is within 1-8%).

Under *single-six-bit-rates* workload an equal number of clients are accessing the six single files encoded at six basic bit rates used in the study, i.e. N clients are accessing the same file encoded at 28 Kb/s, N clients are accessing the same file encoded at 56 Kb/s, etc.

Under *unique-six-bit-rates* workload an equal number of clients are accessing the sets of different (unique) files encoded at six basic bit rates, i.e. N clients are accessing N different files encoded at 28 Kb/s, N clients are accessing N different files encoded at 56 Kb/s, etc.

Figure 7 shows the measured vs expected server capacities under *single-six-bit-rates* and *unique-six-bit-rates* workloads.³ The measured server capacity matches the expected server capacity very well for both workloads (the error is less than 4%).




Figure 7: Server Capacity: Measured vs Expected a) *single-unique-one-bit-rate* workload b) *unique-six-bit-rates* workload.

Thus, the performance model of media server capacity based on a *cost* function derived from the set of basic benchmark measurements closely approximates the expected media server capacity for realistic workloads.

5 Related Work

Delivering multimedia content over Internet faces many challenges: an immature broadband Internet infrastructure, the real-time nature of multimedia content and

³We used these two workloads to observe: what are the rules for degradation of stream quality under server overload conditions for media files streamed at different encoding bit rates? The preliminary results show that high bit rate streams experience degradation of quality first. We are going to quantify these *QoS fairness rules* in our future work.

its sensitivity to congestion conditions in the Internet, high backbone transmission cost, a multiplicity of competing standards for media encoding and delivery. It is commonly recognized that multimedia applications can consume significant amount of server and network resources. Much of the research community efforts have been concentrating on improving media server technology (e.g. [5, 10], building caching infrastructure support (e.g. [11]), and designing scalable streaming protocols (e.g. [4, 12].

In [6] a probe-based algorithm for QoS specification and adaptation is proposed to evaluate the resources available (required) on a client and a server side to support certain QoS guarantees for Video-On-Demand service. In later work [7], as a part of QoS-aware resource management for distributed multimedia applications, the resource broker performs admission control using client QoS profile in which the amount of required resources is specified. Such a profile either is prebuilt off-line or created in on-line manner using testing service proposed in [6]. Our work proceeds in a similar direction and proposes a unified framework for measuring a performance of commercially available media server applications. Our approach and designed performance model allow one to estimate the necessary system resources required to support a particular stream without QoS degradation.

The current trend of outsourcing network services to third parties has brought a set of new challenging problems to the architecture and design of automatic resource management in Internet Data Centers. In [3], authors present *Muse* - a design and implementation of an architecture for resource management in a hosting center with an emphasis on energy as a driving resource management issue for large server clusters. To demonstrate *Muse* functionality and its main principles authors use web server workloads. In paper [8], authors develop a framework for QoS driven dynamic resource allocation in Internet Data Centers. The authors goal is to adaptively migrate servers between the set of hosted applications accordingly to their workload demands.

A set of commercial benchmarks are proposed for measuring a performance of commercial web servers. However, similar benchmarks and the corresponding performance studies are not currently available for commercial (academia/research) media servers. Work performed in this paper presents a step in this direction.

6 Conclusion and Future Work

The Utility Data Center provides a convenient, flexible infrastructure to dynamically support hosting services for Internet applications. Among the *prerequisites* to a design of a utility-aware service is the ability to measure the capacity of media servers (nominal and currently in use) in order to evaluate the amount of available system resources for admitting new client requests or to request/release resources as needed.

A standard commercial stress test used by most of the

media server vendors measures a maximum number of concurrent streams delivered by the server when all the clients are accessing the same file encoded at a certain fixed bit rate. We argue, that such test measurements are obviously not sufficient for media server benchmarking. We also argue, that monitoring of a particular system resource can not be used for evaluating of a currently available/used capacity of media service.

In this work, we introduce a set of basic benchmarks to measure the basic capacities of streaming media systems. The set of basic benchmarks allows one to derive the scaling rules of server capacity for media files encoded at different bit rates and delivered from memory vs disk. Using experimental testbed, we show that these scaling rules are non-trivial. We design a performance model of media server which uses the basic benchmark measurements for estimating a resource requirements of a particular stream in order to compute the server capacity for delivering a realistic workload.

Acknowledgements: Authors would like to thank Wenting Tang and Yun Fu for stimulating discussions, Nic Lyons, Brett Bausk, and Jeff Krennek for their help in media content creation, and John Sontag for his active support of this work.

References

- [1] Almeida, J. M., J. Krueger, D. L. Eager, and M. K. Vernon. Analysis of Educational Media Server Workloads, Proc. 11th Int'l. Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV 2001), June 2001.
- [2] L. Cherkasova, M. Gupta. Characterizing Locality, Evolution, and Life Span of Accesses in Enterprise Media Server Workloads. 12th Int'l. Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV 2002), May 2002.
- [3] J. Chase, D. Anderson, P. Thakar, A. Vahdat, and R. Doyle. Managing Energy and Server Resources in Hosting Centers. Jeffrey Proceedings of the Eighteenth ACM Symposium on Operating Systems Principles (SOSP), October 2001.
- [4] L. Gao, J. Kurose, and D. Towsley. Catching and Selective Catching: Efficient Latency Reduction Techniques for Delivering Continuous Multimedia Streams. Proc. ACM Multimedia'99, Orlando, FL, Nov. 1999.
- [5] D. Eager, M. Vernon, and J. Zahorjan. Optimal and Efficient Merging Schedules for Video-on-Demand Servers. Proc. 7th ACM Multimedia Conf. (Multimedia '99), Orlando, Nov., 1999.
- [6] K. Nahrstedt, A. Hossain, S.-M. Kang. A Probe-based Algorithm for QoS Specification and Adaptation, in Proc. of 4th IFIP International Workshop on QoS (IWQoS'96), Paris, France, March, 1996.
- [7] K. Nahrstedt, H. Chu, S. Narayan. QoS-aware Resource Management for Distributed Multimedia Applications, Journal on High-Speed Networking, Special Issue on Multimedia Networking, vol. 8, num. 3-4, IOS Press, 1998.
- [8] S. Ranjan, J. Rolia, H. Fu, E. Knightly. QoS-Driven Server Migration for Internet Data Centers. Proc. of ACM/IEEE Int'l Workshop on Quality of Service (IWQoS), Miami Beach, FL, May 2002.
- [9] Realsystem Server Product Line Comparison. <http://www.reálnetworks.com/products/servers/comparison.html>
- [10] P. Shenoy, S. Hasan, P. Kulkarni, and K. Ramamritham. Middleware versus Native OS Support: Architectural Considerations for Supporting Multimedia Applications. Proc. IEEE Real-time Technology and Applications Symposium (RTAS'02), San Jose, CA, September 2002.
- [11] S. Sen, J. Rexford, and D. Towsley. Proxy Prefix Caching for Multimedia Streams. In Proc. IEEE INFOCOM, April, 1999.
- [12] H. Tan, H., D. L. Eager, M. K. Vernon, and H. Guo. Quality of Service Evaluations of Multicast Streaming Protocols. Proc. ACM SIGMETRICS 2002 Int'l. Conf. on Measurement and Modeling of Computer Systems (SIGMETRICS 2002), Marina del Rey, CA, June 2002.
- [13] Utility Data Center: Solutions. <http://www.hp.com/solutions1/infrastructure/solutions/utilitydata/index.html>
- [14] Utility Data Center: HP's First Proof Point for Service-Centric Computing. IDC white paper. <http://www.idc.com>