

EXHIBIT A
(Redacted Cromarty Declaration)
(Part 2)



AWS

Products

Developers

Community

Support

Account

Amazon EC2 Pricing

Pay only for what you use. There is no minimum fee. Estimate your monthly bill using [AWS Simple Monthly Calculator](#). The prices listed are based on the Region in which your instance is running. For a detailed comparison between On-Demand Instances, Reserved Instances and Spot Instances, see [Amazon EC2 Instance Purchasing Options](#).

Amazon EC2 Details

[EC2 Overview](#)
[EC2 FAQs](#)
EC2 Pricing
[Amazon EC2 SLA](#)
[EC2 Instance Types](#)
[EC2 Instance Purchasing Options](#)
[Reserved Instances](#)
[Spot Instances](#)
[Windows Instances](#)

Free Tier*

As part of [AWS's Free Usage Tier](#), new AWS customers can get started with Amazon EC2 for free. Upon sign-up, new AWS customers receive the following EC2 services each month for one year:

750 hours of EC2 running Linux/Unix Micro instance usage

750 hours of Elastic Load Balancing plus 15 GB data processing

10 GB of Amazon Elastic Block Storage (EBS) plus 1 million IOs and 1 GB snapshot storage

15 GB of bandwidth out aggregated across all AWS services

1 GB of Regional Data Transfer

Amazon EC2 Features

[Elastic Block Store](#)
[Amazon CloudWatch](#)
[Auto Scaling](#)
[Elastic Load Balancing](#)
[High Performance Computing](#)
[VM Import](#)

On-Demand Instances

On-Demand Instances let you pay for compute capacity by the hour with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs.

The pricing below includes the cost to run private and public AMIs on the specified operating system ("Windows Usage" prices apply to Windows Server® 2003 R2, 2008 and 2008 R2). Amazon also provides you with additional instances for [Amazon EC2 running Microsoft Windows with SQL Server](#), [Amazon EC2 running SUSE Linux Enterprise Server](#), [Amazon EC2 running Red Hat Enterprise Linux](#) and [Amazon EC2 running IBM](#) that are priced differently.

Region:

Linux/UNIX Usage

Windows Usage

Standard On-Demand Instances

Small (Default)	\$0.085 per hour	\$0.12 per hour
Large	\$0.34 per hour	\$0.48 per hour
Extra Large	\$0.68 per hour	\$0.96 per hour

Micro On-Demand Instances

Micro	\$0.02 per hour	\$0.03 per hour
-------	-----------------	-----------------

Hi-Memory On-Demand Instances

Extra Large	\$0.50 per hour	\$0.62 per hour
Double Extra Large	\$1.00 per hour	\$1.24 per hour
Quadruple Extra Large	\$2.00 per hour	\$2.48 per hour

Hi-CPU On-Demand Instances

Medium	\$0.17 per hour	\$0.29 per hour
Extra Large	\$0.68 per hour	\$1.16 per hour

Cluster Compute Instances

Quadruple Extra Large	\$1.30 per hour	\$1.61 per hour
Cluster Compute Eight Extra Large	\$2.40 per hour	\$2.97 per hour

Cluster GPU Instances

Related Resources

[AWS Management Console](#)
[Documentation](#)
[Release Notes](#)
[Developer Tools](#)
[Sample Code & Libraries](#)
[Developer Tools](#)
[Articles & Tutorials](#)
[Amazon Machine Images \(AMIs\)](#)
[Public Data Sets on AWS](#)
[Community Forum](#)

Want to Save on Your Amazon EC2 Bill?



See how customers like Ooyala significantly reduced their Amazon EC2 bill by using Spot Instances.

> [Learn More](#)

EC2 Running Microsoft



Amazon EC2 running Microsoft Windows Server® 2003/2008 is a fast and dependable environment for deploying applications using the Microsoft Web Platform.

> [Learn More](#)

EC2 Running IBM



You can run many of the proven IBM platform technologies with which you're already familiar.

> [Learn More](#)

	Linux/UNIX Usage	Windows Usage
Quadruple Extra Large	\$2.10 per hour	\$2.60 per hour

Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it is terminated. Each partial instance-hour consumed will be billed as a full hour.

Reserved Instances

Reserved Instances give you the option to make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly usage charge for that instance. After the one-time payment for an instance, that instance is reserved for you, and you have no further obligation; you may choose to run that instance for the discounted usage rate for the duration of your term, or when you do not use the instance, you will not pay usage charges on it. In addition to Reserved Instances for Linux/UNIX and Windows operating systems specified below, we also offer Reserved Instances for [Amazon EC2 running SUSE Linux Enterprise Server](#) and [Amazon EC2 running Microsoft SQL Server](#).

[Dedicated Reserved Instances](#) are also available.

Region: <input type="text" value="US East (Virginia)"/>	1 yr Term	3 yr Term	Linux/UNIX Usage	Windows Usage
Standard Reserved Instances				
Small (Default)	\$227.50	\$350	\$0.03 per hour	\$0.05 per hour
Large	\$910	\$1400	\$0.12 per hour	\$0.20 per hour
Extra Large	\$1820	\$2800	\$0.24 per hour	\$0.40 per hour
Micro Reserved Instances				
Micro	\$54	\$82	\$0.007 per hour	\$0.013 per hour
High-Memory Reserved Instances				
Extra Large	\$1325	\$2000	\$0.17 per hour	\$0.24 per hour
Double Extra Large	\$2650	\$4000	\$0.34 per hour	\$0.48 per hour
Quadruple Extra Large	\$5300	\$8000	\$0.68 per hour	\$0.96 per hour
High-CPU Reserved Instances				
Medium	\$455	\$700	\$0.06 per hour	\$0.125 per hour
Extra Large	\$1820	\$2800	\$0.24 per hour	\$0.50 per hour
Cluster Compute Reserved Instances				
Quadruple Extra Large	\$3286	\$5056	\$0.45 per hour	\$0.63 per hour
Cluster Compute Eight Extra Large	\$4146	\$6378	\$0.54 per hour	\$0.75 per hour
Cluster GPU Reserved Instances				
Quadruple Extra Large	\$5630	\$8650	\$0.74 per hour	\$1.04 per hour

Reserved Instances can be purchased for 1 or 3 year terms, and the one-time fee per instance is non-refundable. Usage pricing is per instance-hour consumed. Instance-hours are billed for the time that instances are in a running state; if you do not run the instance in an hour, there is zero usage charge. Partial instance-hours consumed are billed as full hours.

If Microsoft chooses to increase the license fees that it charges for Windows, we may correspondingly increase the per-hour usage rate for previously purchased Reserved Instances with Windows. The initial one-time payment for a Reserved Instance will be unaffected in this situation. Any such changes would be made between Dec 1 – Jan 31, and with at least 30 days' notice. If the per-hour usage rate does increase, you may continue to use your Reserved Instance with Windows with the new per-hour usage rate, convert your Reserved Instance with Windows to a Reserved Instance with Linux, or request a pro rata refund of the upfront fee you paid for the Reserved Instance with Windows.

Reserved Instances are available for Linux/UNIX, Windows and SUSE Linux Enterprise operating systems. You can also optionally reserve instances in Amazon VPC at the same prices as shown above. [Click here](#) to learn more about Reserved Instances.

Spot Instances

Spot Instances enable you to bid for unused Amazon EC2 capacity. Instances are charged the Spot Price, which is set by Amazon EC2 and fluctuates periodically depending on the supply of and demand for Spot Instance capacity. To use

Spot Instances, you place a Spot Instance request, specifying the instance type, the Availability Zone desired, the number of Spot Instances you want to run, and the maximum price you are willing to pay per instance hour. To determine how that maximum price compares to past Spot Prices, the Spot Price history is available via the Amazon EC2 API and the AWS Management Console. If your maximum price bid exceeds the current Spot Price, your request is fulfilled and your instances will run until either you choose to terminate them or the Spot Price increases above your maximum price (whichever is sooner).

[Click here](#) to learn more about Spot Instances. For information on how to get started, [click here](#).

The following table displays the Lowest Spot Price per Region and instance type (updated every 5 minutes). In addition to Linux/Unix and Windows, we also offer Spot Instances for [Amazon EC2 running SUSE Linux Enterprise Server](#).

Region: <input type="text" value="US East (Virginia)"/>		
	Linux/UNIX Usage	Windows Usage
Standard Spot Instances		
Small (Default)	\$0.05 per hour	\$0.12 per hour
Large	\$0.108 per hour	\$0.18 per hour
Extra Large	\$0.216 per hour	
Micro Spot Instances		
Micro	\$0.006 per hour	\$0.012 per hour
High-Memory Spot Instances		
Extra Large	\$0.17 per hour	\$0.216 per hour
Double Extra Large	\$0.42 per hour	\$0.378 per hour
Quadruple Extra Large	\$0.756 per hour	\$0.99 per hour
High-CPU Spot Instances		
Medium	\$0.085 per hour	\$0.113 per hour
Extra Large	\$0.216 per hour	\$0.45 per hour
Cluster Compute Instances		
Quadruple Extra Large	\$1.6 per hour	N/A*
Cluster GPU Instances		
Quadruple Extra Large	\$0.74 per hour	N/A*

* Windows® is not currently available for Cluster Compute or Cluster GPU Instances

If you would like to go straight to a view of the latest Spot Instance pricing:

1. [Log in](#) to the AWS Management Console, then click the "Amazon EC2" tab.
2. Click on "Spot Requests" in the navigation pane on the left.
3. Click on "Pricing History" to open a view of pricing selectable by instance type.

Data Transfer**

Internet Data Transfer

The pricing below is based on data transferred "in" and "out" of Amazon EC2.

Region: <input type="text" value="US East (Virginia)"/>	
	Pricing
Data Transfer IN	
All data transfer in	\$0.000 per GB
Data Transfer OUT	
First 1 GB / month	\$0.000 per GB
Up to 10 TB / month	\$0.120 per GB
Next 40 TB / month	\$0.090 per GB
Next 100 TB / month	\$0.070 per GB
Next 350 TB / month	\$0.050 per GB

	Pricing
Next 524 TB / month	Contact Us
Next 4 PB / month	Contact Us
Greater than 5 PB / month	Contact Us

There is no Data Transfer charge between Amazon EC2 and other Amazon Web Services within the same region (i.e. between Amazon EC2 US West and Amazon S3 in US West). Data transferred between Amazon EC2 instances located in different Availability Zones in the same Region will be charged Regional Data Transfer. Data transferred between AWS services in different regions will be charged as Internet Data Transfer on both sides of the transfer.

Usage for other Amazon Web Services is billed separately from Amazon EC2.

Availability Zone Data Transfer

\$0.00 per GB – all data transferred between instances in the same Availability Zone using private IP addresses.

Regional Data Transfer

\$0.01 per GB – all data transferred between instances in different Availability Zones in the same region.

Public and Elastic IP and Elastic Load Balancing Data Transfer

\$0.01 per GB in/out – If you choose to communicate using your Public or Elastic IP address or Elastic Load Balancer inside of the Amazon EC2 network, you'll pay Regional Data Transfer rates even if the instances are in the same Availability Zone. For data transfer within the same Availability Zone, you can easily avoid this charge (and get better network performance) by using your private IP whenever possible.

Amazon Elastic Block Store

Region: <input type="text" value="US East (Virginia)"/>
Amazon EBS Volumes
\$0.10 per GB-month of provisioned storage
\$0.10 per 1 million I/O requests
Amazon EBS Snapshots to Amazon S3
\$0.14 per GB-month of data stored

Elastic IP Addresses

Region: <input type="text" value="US East (Virginia)"/>
No cost for Elastic IP addresses while in use
\$0.01 per non-attached Elastic IP address per complete hour
\$0.00 per Elastic IP address remap – first 100 remaps / month
\$0.10 per Elastic IP address remap – additional remap / month over 100

Amazon CloudWatch

Region: <input type="text" value="US East (Virginia)"/>
Detailed Monitoring for Amazon EC2 Instances
\$3.50 per instance per month, provided at 1-minute frequency
Basic Monitoring for Amazon EC2 instances
\$0.00 (free of charge) per instance per month, provided at 5-minute frequency

Monitoring for Custom Metrics

\$0.50 per metric per month

Detailed Monitoring for Amazon EC2 is charged at standard Amazon CloudWatch rates of \$0.50 per metric per month. Each instance includes seven metrics for total charges of \$3.50 per month. Partial months are charged on an hourly pro rata basis, at approximately \$0.005/instance-hour.

Note: This new pricing for Detailed Monitoring (representing a 68% decrease from the current price) takes effect starting June 1, 2011. Prior to that, the price remains \$0.015 per instance-hour or partial hour. Pricing for Amazon CloudWatch Custom Metrics takes effect starting June 1, 2011. Custom metrics (that you send and Amazon CloudWatch monitors) before that time are free of charge.

Learn more about [Amazon Cloudwatch](#).

Auto Scaling

Auto Scaling is enabled by Amazon CloudWatch and carries no additional fees. Each instance launched by Auto Scaling is automatically enabled for monitoring and the applicable [Amazon Cloudwatch](#) charges will be applied.

Elastic Load Balancing

Region:

\$0.025 per Elastic Load Balancer-hour (or partial hour)

\$0.008 per GB of data processed by an Elastic Load Balancer

AWS GovCloud Region

[AWS GovCloud](#) is an AWS Region designed to allow U.S. government agencies and contractors to move more sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements. For pricing and more information on the new AWS GovCloud Region, please visit the [AWS GovCloud Web Page](#).

** Your usage for the free tier is calculated each month across all regions except the AWS GovCloud Region, and automatically applied to your bill – unused monthly usage will not roll over. Does not include Amazon EC2 running Microsoft, Amazon EC2 running SUSE Linux Enterprise Server, Amazon EC2 running IBM, and the AWS GovCloud Region. See [offer terms](#) for more details and other restrictions.*

*** As part of AWS's Free Usage Tier, new AWS customers will receive free 15 GB of data transfer out each month aggregated across all AWS services for one year except in the AWS GovCloud Region.*

**** Rate tiers take into account your aggregate Data Transfer Out usage across Amazon EC2, Amazon S3, Amazon RDS, Amazon SimpleDB, Amazon SQS, Amazon SNS and Amazon VPC.*

(Amazon EC2 is sold by Amazon Web Services LLC.)

[↑ Top](#)

Learn

- [Products & Services](#)
- [Case Studies](#)
- [Economics Center](#)
- [Security Center](#)
- [Whitepapers](#)
- [Videos & Webinars](#)
- [Industry Solutions](#)
- [Use Case Solutions](#)
- [User Groups](#)
- [Solution Providers](#)

Develop

- [Developer Resources](#)
- [AMI Catalog](#)
- [Sample Code & Libraries](#)
- [Dev Tools & SDKs](#)
- [Documentation](#)
- [Articles & Tutorials](#)
- [Management Console](#)

- [Developer Centers](#)
- [Java](#)

Manage

- [Your Account](#)
- [Management Console](#)
- [Account Activity](#)
- [Usage Reports](#)
- [Personal Information](#)
- [Payment Method](#)
- [AWS Identity and Access Management](#)
- [Security Credentials](#)
- [Request Service Limit Increases](#)

About AWS

- [About Us](#)
- [Events](#)
- [Careers at AWS](#)
- [Contact Us](#)
- [Announcements \(What's New?\)](#)
- [Media Coverage](#)
- [Privacy Policy](#)



[Mobile](#)
[PHP](#)
[Python](#)
[Ruby](#)
[Windows & .NET](#)

[Support](#)
[Premium Support](#)
[Service Health Dashboard](#)
[Discussion Forums](#)
[FAQs](#)
[Contact Support](#)

©2011, Amazon Web Services LLC or its affiliates. All rights reserved.
An  company



AWS	Products	Developers	Community	Support
Account				

Amazon Elastic Compute Cloud (Amazon EC2)

Amazon EC2 Details

EC2 Overview

[EC2 FAQs](#)
[EC2 Pricing](#)
[Amazon EC2 SLA](#)
[EC2 Instance Types](#)
[EC2 Instance Purchasing Options](#)
[Reserved Instances](#)
[Spot Instances](#)
[Windows Instances](#)

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers.

Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers the tools to build failure resilient applications and isolate themselves from common failure scenarios.

Easy to sign up, pay only for what you use

[Sign Up Now](#)

Amazon EC2 Features

[Elastic Block Store](#)
[Amazon CloudWatch](#)
[Auto Scaling](#)
[Elastic Load Balancing](#)
[High Performance Computing](#)
[VM Import](#)
[Amazon EC2 Functionality](#)
[Service Highlights](#)
[Features](#)
[Instance Types](#)
[Operating Systems and Software](#)
[Pricing](#)
[Resources](#)
[Detailed Description](#)
[Intended Usage and Restrictions](#)

This page contains the following categories of information. Click to jump down:

Related Resources

[AWS Management Console](#)
[Documentation](#)
[Release Notes](#)
[Sample Code & Libraries](#)
[Developer Tools](#)
[Articles & Tutorials](#)
[Amazon Machine Images \(AMIs\)](#)
[Public Data Sets on AWS](#)
[Community Forum](#)

Amazon EC2 Functionality

Amazon EC2 presents a true virtual computing environment, allowing you to use web service interfaces to launch instances with a variety of operating systems, load them with your custom application environment, manage your network's access permissions, and run your image using as many or few systems as you desire.

To use Amazon EC2, you simply:

Select a pre-configured, templated image to get up and running immediately. Or create an Amazon Machine Image (AMI) containing your applications, libraries, data, and associated configuration settings.

Configure security and network access on your Amazon EC2 instance.

Choose which instance type(s) and operating system you want, then start, terminate, and monitor as many instances of your AMI as needed, using the web service APIs or the variety of management tools provided.

Determine whether you want to run in multiple locations, utilize static IP endpoints, or attach persistent block storage to your instances.

Pay only for the resources that you actually consume, like instance-hours or data transfer.

[↑ Top](#)

Service Highlights

Elastic – Amazon EC2 enables you to increase or decrease capacity within minutes, not hours or days. You can commission one, hundreds or even thousands of server instances simultaneously. Of course, because this is all controlled with web service APIs, your application can automatically scale itself up and down depending on its needs.

Completely Controlled – You have complete control of your instances. You have root access to each one, and you can interact with them as you would any machine. You can stop your instance while retaining the data on your boot partition and then subsequently restart the same instance using web service APIs. Instances can be rebooted remotely using web service APIs. You also have access to console output of your instances.

Want to Save on Your Amazon EC2 Bill?





See how customers like foursquare significantly reduced their Amazon EC2 bill by using Spot Instances.

> [Learn More](#)

EC2 Running SUSE

Amazon EC2 running SUSE Linux Enterprise Server is a proven platform to deliver peak performance and includes automatic updates for security patches, bug fixes and new features.

> [Learn More](#)

MySQL made easy



Amazon Relational Database Service (Amazon RDS) makes it easier for you to set up, manage, and scale a relational database in the cloud.

> [Learn about Amazon RDS](#)

EC2 Running Microsoft



Amazon EC2 running Microsoft Windows Server® is a fast and dependable environment for deploying applications using the Microsoft Web Platform.

> [Learn More](#)

Amazon EC2 is Hiring!



Amazon EC2 is hiring to rapidly expand our service.

> [Learn More](#)

Get Started For Free

New AWS customers receive free usage tiers of compute, storage, and bandwidth every month for one year.

Flexible – You have the choice of multiple instance types, operating systems, and software packages. Amazon EC2 allows you to select a configuration of memory, CPU, instance storage, and the boot partition size that is optimal for your choice of operating system and application. For example, your choice of operating systems includes numerous Linux distributions, and Microsoft Windows Server.

Designed for use with other Amazon Web Services – Amazon EC2 works in conjunction with Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), Amazon SimpleDB and Amazon Simple Queue Service (Amazon SQS) to provide a complete solution for computing, query processing and storage across a wide range of applications.

Reliable – Amazon EC2 offers a highly reliable environment where replacement instances can be rapidly and predictably commissioned. The service runs within Amazon's proven network infrastructure and datacenters. The Amazon EC2 Service Level Agreement commitment is 99.95% availability for each Amazon EC2 Region.

Secure – Amazon EC2 provides numerous mechanisms for securing your compute resources.

Amazon EC2 includes web service interfaces to configure firewall settings that control network access to and between groups of instances.

When launching Amazon EC2 resources within [Amazon Virtual Private Cloud](#) (Amazon VPC), you can isolate your compute instances by specifying the IP range you wish to use, and connect to your existing IT infrastructure using industry-standard encrypted IPsec VPN. You can also choose to launch [Dedicated Instances](#) into your VPC. Dedicated Instances are Amazon EC2 Instances that run on hardware dedicated to a single customer for additional isolation.

For more information on Amazon EC2 security refer to our [Amazon Web Services: Overview of Security Process document](#).

Inexpensive – Amazon EC2 passes on to you the financial benefits of Amazon's scale. You pay a very low rate for the compute capacity you actually consume. See [Amazon EC2 Instance Purchasing Options](#) for a more detailed description.

On-Demand Instances – On-Demand Instances let you pay for compute capacity by the hour with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs. On-Demand Instances also remove the need to buy "safety net" capacity to handle periodic traffic spikes.

Reserved Instances – Reserved Instances give you the option to make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly usage charge for that instance. After the one-time payment for an instance, that instance is reserved for you, and you have no further obligation; you may choose to run that instance for the discounted usage rate for the duration of your term, or when you do not use the instance, you will not pay usage charges on it.

Spot Instances – Spot Instances allow customers to bid on unused Amazon EC2 capacity and run those instances for as long as their bid exceeds the current Spot Price. The Spot Price changes periodically based on supply and demand, and customers whose bids meet or exceed it gain access to the available Spot Instances. If you have flexibility in when your applications can run, Spot Instances can significantly lower your Amazon EC2 costs. See [here](#) for more details on Spot Instances.

Features

Amazon EC2 provides a number of powerful features for building scalable, failure resilient, enterprise class applications, including:

Amazon Elastic Block Store – Amazon Elastic Block Store (EBS) offers persistent storage for Amazon EC2 instances. Amazon EBS volumes provide off-instance storage that persists independently from the life of an instance. Amazon EBS volumes are highly available, highly reliable volumes that can be leveraged as an Amazon EC2 instance's boot partition or attached to a running Amazon EC2 instance as a standard block device. When used as a boot partition, Amazon EC2 instances can be stopped and subsequently restarted, enabling you to only pay for the storage resources used while maintaining your instance's state. Amazon EBS volumes offer greatly improved durability over local Amazon EC2 instance stores, as Amazon EBS volumes are automatically replicated on the backend (in a single Availability Zone). For those wanting even more durability, Amazon EBS provides the ability to create point-in-time consistent snapshots of your volumes that are then stored in Amazon S3, and automatically replicated across multiple Availability Zones. These snapshots can be used as the starting point for new Amazon EBS volumes, and can protect your data for long term durability. You can also easily share these snapshots with co-workers and other AWS developers. See [Amazon Elastic Block Store](#) for more details on this feature.

Multiple Locations – Amazon EC2 provides the ability to place instances in multiple locations. Amazon EC2 locations are composed of Regions and Availability Zones. Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to other Availability Zones in the same Region. By launching instances in separate Availability Zones, you can protect your applications from failure of a single location. Regions consist of one or more Availability Zones, are geographically dispersed, and will be in separate geographic areas or

[> Learn More](#)

Webinar: Getting Started with Windows on AWS

Attend an informative webinar designed for business or technical decision makers and IT professionals and learn more about the top business reasons to run Windows Server, SQL Server, or other Windows Server applications on the AWS Cloud.

[> Register for the December 8th Webinar Today](#)

countries. The [Amazon EC2 Service Level Agreement](#) commitment is 99.95% availability for each Amazon EC2 Region. Amazon EC2 is currently available in seven regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), EU (Ireland), Asia Pacific (Singapore), Asia Pacific (Tokyo), and [AWS GovCloud](#).

Elastic IP Addresses – Elastic IP addresses are static IP addresses designed for dynamic cloud computing. An Elastic IP address is associated with your account not a particular instance, and you control that address until you choose to explicitly release it. Unlike traditional static IP addresses, however, Elastic IP addresses allow you to mask instance or Availability Zone failures by programmatically remapping your public IP addresses to any instance in your account. Rather than waiting on a data technician to reconfigure or replace your host, or waiting for DNS to propagate to all of your customers, Amazon EC2 enables you to engineer around problems with your instance or software by quickly remapping your Elastic IP address to a replacement instance. In addition, you can optionally configure the reverse DNS record of any of your Elastic IP addresses by filling out this [form](#).

Amazon Virtual Private Cloud – Amazon VPC is a secure and seamless bridge between a company's existing IT infrastructure and the AWS cloud. Amazon VPC enables enterprises to connect their existing infrastructure to a set of isolated AWS compute resources via a Virtual Private Network (VPN) connection, and to extend their existing management capabilities such as security services, firewalls, and intrusion detection systems to include their AWS resources. See [Amazon Virtual Private Cloud](#) for more details.

Amazon CloudWatch – Amazon CloudWatch is a web service that provides monitoring for AWS cloud resources and applications, starting with Amazon EC2. It provides you with visibility into resource utilization, operational performance, and overall demand patterns—including metrics such as CPU utilization, disk reads and writes, and network traffic. You can get statistics, view graphs, and set alarms for your metric data. To use Amazon CloudWatch, simply select the Amazon EC2 instances that you'd like to monitor. You can also supply your own business or application metric data. Amazon CloudWatch will begin aggregating and storing monitoring data that can be accessed using web service APIs or Command Line Tools. See [Amazon CloudWatch](#) for more details.

Auto Scaling – Auto Scaling allows you to automatically scale your Amazon EC2 capacity up or down according to conditions you define. With Auto Scaling, you can ensure that the number of Amazon EC2 instances you're using scales up seamlessly during demand spikes to maintain performance, and scales down automatically during demand lulls to minimize costs. Auto Scaling is particularly well suited for applications that experience hourly, daily, or weekly variability in usage. Auto Scaling is enabled by Amazon CloudWatch and available at no additional charge beyond Amazon CloudWatch fees. See [Auto Scaling](#) for more details.

Elastic Load Balancing – Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances. It enables you to achieve even greater fault tolerance in your applications, seamlessly providing the amount of load balancing capacity needed in response to incoming application traffic. Elastic Load Balancing detects unhealthy instances within a pool and automatically reroutes traffic to healthy instances until the unhealthy instances have been restored. You can enable Elastic Load Balancing within a single Availability Zone or across multiple zones for even more consistent application performance. Amazon CloudWatch can be used to capture a specific Elastic Load Balancer's operational metrics, such as request count and request latency, at no additional cost beyond Elastic Load Balancing fees. See [Elastic Load Balancing](#) for more details.

High Performance Computing (HPC) Clusters – Customers with complex computational workloads such as tightly coupled parallel processes, or with applications sensitive to network performance, can achieve the same high compute and network performance provided by custom-built infrastructure while benefiting from the elasticity, flexibility and cost advantages of Amazon EC2. Cluster Compute and Cluster GPU Instances have been specifically engineered to provide high-performance network capability and can be programmatically launched into clusters – allowing applications to get the low-latency network performance required for tightly coupled, node-to-node communication. Cluster Compute and Cluster GPU Instances also provide significantly increased network throughput making them well suited for customer applications that need to perform network-intensive operations. [Learn more](#) about Cluster Compute and Cluster GPU Instances as well as other AWS services that can be used for HPC Applications.

VM Import – VM Import enables you to easily import virtual machine images from your existing environment to Amazon EC2 instances. VM Import allows you to leverage your existing investments in the virtual machines that you have built to meet your IT security, configuration management, and compliance requirements by seamlessly bringing those virtual machines into Amazon EC2 as ready-to-use instances. This offering is available at no additional charge beyond standard usage charges for Amazon EC2 and Amazon S3. [Learn more](#) about VM Import.

[↑ Top](#)

Instance Types

Standard Instances

Instances of this family are well suited for most applications.

Small Instance (Default) 1.7 GB of memory, 1 EC2 Compute Unit (1 virtual core with 1 EC2 Compute Unit), 160 GB of local instance storage, 32-bit platform

Large Instance 7.5 GB of memory, 4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each), 850 GB of local instance storage, 64-bit platform

Extra Large Instance 15 GB of memory, 8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units each), 1690 GB of local instance storage, 64-bit platform

Micro Instances

Instances of this family provide a small amount of consistent CPU resources and allow you to burst CPU capacity when additional cycles are available. They are well suited for lower throughput applications and web sites that consume significant compute cycles periodically.

Micro Instance 613 MB of memory, up to 2 ECUs (for short periodic bursts), EBS storage only, 32-bit or 64-bit platform

High-Memory Instances

Instances of this family offer large memory sizes for high throughput applications, including database and memory caching applications.

High-Memory Extra Large Instance 17.1 GB memory, 6.5 ECU (2 virtual cores with 3.25 EC2 Compute Units each), 420 GB of local instance storage, 64-bit platform

High-Memory Double Extra Large Instance 34.2 GB of memory, 13 EC2 Compute Units (4 virtual cores with 3.25 EC2 Compute Units each), 850 GB of local instance storage, 64-bit platform

High-Memory Quadruple Extra Large Instance 68.4 GB of memory, 26 EC2 Compute Units (8 virtual cores with 3.25 EC2 Compute Units each), 1690 GB of local instance storage, 64-bit platform

High-CPU Instances

Instances of this family have proportionally more CPU resources than memory (RAM) and are well suited for compute-intensive applications.

High-CPU Medium Instance 1.7 GB of memory, 5 EC2 Compute Units (2 virtual cores with 2.5 EC2 Compute Units each), 350 GB of local instance storage, 32-bit platform

High-CPU Extra Large Instance 7 GB of memory, 20 EC2 Compute Units (8 virtual cores with 2.5 EC2 Compute Units each), 1690 GB of local instance storage, 64-bit platform

Cluster Compute Instances

Instances of this family provide proportionally high CPU with increased network performance and are well suited for High Performance Compute (HPC) applications and other demanding network-bound applications. [Learn more](#) about use of this instance type for HPC applications.

Cluster Compute Quadruple Extra Large 23 GB memory, 33.5 EC2 Compute Units, 1690 GB of local instance storage, 64-bit platform, 10 Gigabit Ethernet

Cluster Compute Eight Extra Large 60.5 GB memory, 88 EC2 Compute Units, 3370 GB of local instance storage, 64-bit platform, 10 Gigabit Ethernet

Cluster GPU Instances

Instances of this family provide general-purpose graphics processing units (GPUs) with proportionally high CPU and increased network performance for applications benefitting from highly parallelized processing, including HPC, rendering and media processing applications. While Cluster Compute Instances provide the ability to create clusters of instances connected by a low latency, high throughput network, Cluster GPU Instances provide an additional option for applications that can benefit from the efficiency gains of the parallel computing power of GPUs over what can be achieved with traditional processors. [Learn more](#) about use of this instance type for HPC applications.

Cluster GPU Quadruple Extra Large 22 GB memory, 33.5 EC2 Compute Units, 2 x NVIDIA Tesla "Fermi" M2050 GPUs, 1690 GB of local instance storage, 64-bit platform, 10 Gigabit Ethernet

EC2 Compute Unit (ECU) – One EC2 Compute Unit (ECU) provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor.

See [Amazon EC2 Pricing](#) for details on costs for each instance type.

See [Amazon EC2 Instance Types](#) for a more detailed description of the differences between the available instance types, as well as a complete description of an EC2 Compute Unit.

Operating Systems and Software

Operating Systems

Amazon Machine Images (AMIs) are preconfigured with an ever-growing list of operating systems. We work with our partners and community to provide you with the most choice possible. You are also empowered to use our bundling tools to upload your own operating systems. The operating systems currently available to use with your Amazon EC2 instances include:

Operating Systems		
Red Hat Enterprise Linux	Windows Server	Oracle Enterprise Linux
SUSE Linux Enterprise	Amazon Linux AMI	Ubuntu Linux
Fedora	Gentoo Linux	Debian

Software

Amazon EC2 enables our partners and customers to build and customize Amazon Machine Images (AMIs) with software based on your needs. We have hundreds of free and paid AMIs available for you to use. A small sampling of the software available for use today within Amazon EC2 includes:

Databases	Resource Management	Web Hosting
IBM DB2	StackIQ Rocks+	Apache HTTP
IBM Informix Dynamic Server	Hadoop	IIS/Asp.Net
Microsoft SQL Server Standard	Condor	IBM Lotus Web Content Management
MySQL Enterprise		IBM WebSphere Portal Server
Oracle Database 11g		

Application Development Environments	Application Servers	Video Encoding & Streaming
IBM sMash	IBM WebSphere Application Server	Wowza Media Server Pro
JBoss Enterprise Application Platform	Java Application Server	Windows Media Server
Ruby on Rails	Oracle WebLogic Server	

Pricing

Pay only for what you use. There is no minimum fee. Estimate your monthly bill using [AWS Simple Monthly Calculator](#). The prices listed are based on the Region in which your instance is running. For a detailed comparison between On-Demand Instances, Reserved Instances and Spot Instances, see [Amazon EC2 Instance Purchasing Options](#).

Free Tier*

As part of [AWS's Free Usage Tier](#), new AWS customers can get started with Amazon EC2 for free. Upon sign-up, new AWS customers receive the following EC2 services each month for one year:

- 750 hours of EC2 running Linux/Unix Micro instance usage
- 750 hours of Elastic Load Balancing plus 15 GB data processing
- 10 GB of Amazon Elastic Block Storage (EBS) plus 1 million IOs and 1 GB snapshot storage
- 15 GB of bandwidth out aggregated across all AWS services
- 1 GB of Regional Data Transfer

On-Demand Instances

On-Demand Instances let you pay for compute capacity by the hour with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs.

The pricing below includes the cost to run private and public AMIs on the specified operating system ("Windows Usage" prices apply to Windows Server® 2003 R2, 2008 and 2008 R2). Amazon also provides you with additional instances for [Amazon EC2 running Microsoft Windows with SQL Server](#), [Amazon EC2 running SUSE Linux Enterprise Server](#), [Amazon EC2 running Red Hat Enterprise Linux](#) and [Amazon EC2 running IBM](#) that are priced differently.

Region:	US East (Virginia)		
	Linux/UNIX Usage	Windows Usage	
Standard On-Demand Instances			
Small (Default)	\$0.085 per hour	\$0.12 per hour	
Large	\$0.34 per hour	\$0.48 per hour	
Extra Large	\$0.68 per hour	\$0.96 per hour	
Micro On-Demand Instances			
Micro	\$0.02 per hour	\$0.03 per hour	
Hi-Memory On-Demand Instances			
Extra Large	\$0.50 per hour	\$0.62 per hour	
Double Extra Large	\$1.00 per hour	\$1.24 per hour	
Quadruple Extra Large	\$2.00 per hour	\$2.48 per hour	
Hi-CPU On-Demand Instances			
Medium	\$0.17 per hour	\$0.29 per hour	
Extra Large	\$0.68 per hour	\$1.16 per hour	
Cluster Compute Instances			
Quadruple Extra Large	\$1.30 per hour	\$1.61 per hour	
Cluster Compute Eight Extra Large	\$2.40 per hour	\$2.97 per hour	
Cluster GPU Instances			
Quadruple Extra Large	\$2.10 per hour	\$2.60 per hour	

Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it is terminated. Each partial instance-hour consumed will be billed as a full hour.

Reserved Instances

Reserved Instances give you the option to make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly usage charge for that instance. After the one-time payment for an instance, that instance is reserved for you, and you have no further obligation; you may choose to run that instance for the discounted usage rate for the duration of your term, or when you do not use the instance, you will not pay usage charges on it. In addition to Reserved Instances for Linux/UNIX and Windows operating systems specified below, we also offer Reserved Instances for [Amazon EC2 running SUSE Linux Enterprise Server](#) and [Amazon EC2 running Microsoft SQL Server](#).

[Dedicated Reserved Instances](#) are also available.

Region:	US East (Virginia)			
	1 yr Term	3 yr Term	Linux/UNIX Usage	Windows Usage
Standard Reserved Instances				
Small (Default)	\$227.50	\$350	\$0.03 per hour	\$0.05 per hour
Large	\$910	\$1400	\$0.12 per hour	\$0.20 per hour
Extra Large	\$1820	\$2800	\$0.24 per hour	\$0.40 per hour
Micro Reserved Instances				
Micro	\$54	\$82	\$0.007 per hour	\$0.013 per hour
High-Memory Reserved Instances				

	1 yr Term	3 yr Term	Linux/UNIX Usage	Windows Usage
Extra Large	\$1325	\$2000	\$0.17 per hour	\$0.24 per hour
Double Extra Large	\$2650	\$4000	\$0.34 per hour	\$0.48 per hour
Quadruple Extra Large	\$5300	\$8000	\$0.68 per hour	\$0.96 per hour
High-CPU Reserved Instances				
Medium	\$455	\$700	\$0.06 per hour	\$0.125 per hour
Extra Large	\$1820	\$2800	\$0.24 per hour	\$0.50 per hour
Cluster Compute Reserved Instances				
Quadruple Extra Large	\$3286	\$5056	\$0.45 per hour	\$0.63 per hour
Cluster Compute Eight Extra Large	\$4146	\$6378	\$0.54 per hour	\$0.75 per hour
Cluster GPU Reserved Instances				
Quadruple Extra Large	\$5630	\$8650	\$0.74 per hour	\$1.04 per hour

Reserved Instances can be purchased for 1 or 3 year terms, and the one-time fee per instance is non-refundable. Usage pricing is per instance-hour consumed. Instance-hours are billed for the time that instances are in a running state; if you do not run the instance in an hour, there is zero usage charge. Partial instance-hours consumed are billed as full hours.

If Microsoft chooses to increase the license fees that it charges for Windows, we may correspondingly increase the per-hour usage rate for previously purchased Reserved Instances with Windows. The initial one-time payment for a Reserved Instance will be unaffected in this situation. Any such changes would be made between Dec 1 – Jan 31, and with at least 30 days' notice. If the per-hour usage rate does increase, you may continue to use your Reserved Instance with Windows with the new per-hour usage rate, convert your Reserved Instance with Windows to a Reserved Instance with Linux, or request a pro rata refund of the upfront fee you paid for the Reserved Instance with Windows.

Reserved Instances are available for Linux/UNIX, Windows and SUSE Linux Enterprise operating systems. You can also optionally reserve instances in Amazon VPC at the same prices as shown above. [Click here](#) to learn more about Reserved Instances.

Spot Instances

Spot Instances enable you to bid for unused Amazon EC2 capacity. Instances are charged the Spot Price, which is set by Amazon EC2 and fluctuates periodically depending on the supply of and demand for Spot Instance capacity. To use Spot Instances, you place a Spot Instance request, specifying the instance type, the Availability Zone desired, the number of Spot Instances you want to run, and the maximum price you are willing to pay per instance hour. To determine how that maximum price compares to past Spot Prices, the Spot Price history is available via the Amazon EC2 API and the AWS Management Console. If your maximum price bid exceeds the current Spot Price, your request is fulfilled and your instances will run until either you choose to terminate them or the Spot Price increases above your maximum price (whichever is sooner).

[Click here](#) to learn more about Spot Instances. For information on how to get started, [click here](#).

The following table displays the Lowest Spot Price per Region and instance type (updated every 5 minutes). In addition to Linux/Unix and Windows, we also offer Spot Instances for [Amazon EC2 running SUSE Linux Enterprise Server](#).

Region: <input type="text" value="US East (Virginia)"/>	Linux/UNIX Usage	Windows Usage
Standard Spot Instances		
Small (Default)	\$0.05 per hour	\$0.12 per hour
Large	\$0.108 per hour	\$0.18 per hour
Extra Large	\$0.216 per hour	
Micro Spot Instances		
Micro	\$0.006 per hour	\$0.012 per hour
High-Memory Spot Instances		
Extra Large	\$0.17 per hour	\$0.216 per hour
Double Extra Large	\$0.42 per hour	\$0.378 per hour
Quadruple Extra Large	\$0.756 per hour	\$0.99 per hour
High-CPU Spot Instances		

	Linux/UNIX Usage	Windows Usage
Medium	\$0.085 per hour	\$0.113 per hour
Extra Large	\$0.216 per hour	\$0.45 per hour
Cluster Compute Instances		
Quadruple Extra Large	\$1.6 per hour	N/A*
Cluster GPU Instances		
Quadruple Extra Large	\$0.74 per hour	N/A*
* Windows® is not currently available for Cluster Compute or Cluster GPU Instances		

If you would like to go straight to a view of the latest Spot Instance pricing:

1. [Log in](#) to the AWS Management Console, then click the "Amazon EC2" tab.
2. Click on "Spot Requests" in the navigation pane on the left.
3. Click on "Pricing History" to open a view of pricing selectable by instance type.

Data Transfer**

Internet Data Transfer

The pricing below is based on data transferred "in" and "out" of Amazon EC2.

Region: <input type="text" value="US East (Virginia)"/>	
Pricing	
Data Transfer IN	
All data transfer in	\$0.000 per GB
Data Transfer OUT	
First 1 GB / month	\$0.000 per GB
Up to 10 TB / month	\$0.120 per GB
Next 40 TB / month	\$0.090 per GB
Next 100 TB / month	\$0.070 per GB
Next 350 TB / month	\$0.050 per GB
Next 524 TB / month	Contact Us
Next 4 PB / month	Contact Us
Greater than 5 PB / month	Contact Us
<p>There is no Data Transfer charge between Amazon EC2 and other Amazon Web Services within the same region (i.e. between Amazon EC2 US West and Amazon S3 in US West). Data transferred between Amazon EC2 instances located in different Availability Zones in the same Region will be charged Regional Data Transfer. Data transferred between AWS services in different regions will be charged as Internet Data Transfer on both sides of the transfer.</p> <p>Usage for other Amazon Web Services is billed separately from Amazon EC2.</p>	
Availability Zone Data Transfer	
\$0.00 per GB – all data transferred between instances in the same Availability Zone using private IP addresses.	
Regional Data Transfer	
\$0.01 per GB – all data transferred between instances in different Availability Zones in the same region.	
Public and Elastic IP and Elastic Load Balancing Data Transfer	
\$0.01 per GB in/out – If you choose to communicate using your Public or Elastic IP address or Elastic Load Balancer inside of the Amazon EC2 network, you'll pay Regional Data Transfer rates even if the instances are in the same Availability Zone. For data transfer within the same Availability Zone, you can easily avoid this charge (and get better network performance) by using your private IP whenever possible.	

Amazon Elastic Block Store

Region: **Amazon EBS Volumes**

\$0.10 per GB-month of provisioned storage

\$0.10 per 1 million I/O requests

Amazon EBS Snapshots to Amazon S3

\$0.14 per GB-month of data stored

Elastic IP AddressesRegion: **No cost for Elastic IP addresses while in use**

\$0.01 per non-attached Elastic IP address per complete hour

\$0.00 per Elastic IP address remap – first 100 remaps / month

\$0.10 per Elastic IP address remap – additional remap / month over 100

Amazon CloudWatchRegion: **Detailed Monitoring for Amazon EC2 Instances**

\$3.50 per instance per month, provided at 1-minute frequency

Basic Monitoring for Amazon EC2 instances

\$0.00 (free of charge) per instance per month, provided at 5-minute frequency

Monitoring for Custom Metrics

\$0.50 per metric per month

Detailed Monitoring for Amazon EC2 is charged at standard Amazon CloudWatch rates of \$0.50 per metric per month. Each instance includes seven metrics for total charges of \$3.50 per month. Partial months are charged on an hourly pro rata basis, at approximately \$0.005/instance-hour.

Note: This new pricing for Detailed Monitoring (representing a 68% decrease from the current price) takes effect starting June 1, 2011. Prior to that, the price remains \$0.015 per instance-hour or partial hour. Pricing for Amazon CloudWatch Custom Metrics takes effect starting June 1, 2011. Custom metrics (that you send and Amazon CloudWatch monitors) before that time are free of charge.

Learn more about [Amazon Cloudwatch](#).

Auto Scaling

Auto Scaling is enabled by Amazon CloudWatch and carries no additional fees. Each instance launched by Auto Scaling is automatically enabled for monitoring and the applicable [Amazon Cloudwatch](#) charges will be applied.

Elastic Load BalancingRegion:

\$0.025 per Elastic Load Balancer-hour (or partial hour)

\$0.008 per GB of data processed by an Elastic Load Balancer

AWS GovCloud Region

[AWS GovCloud](#) is an AWS Region designed to allow U.S. government agencies and contractors to move more sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements. For pricing and more information on the new AWS GovCloud Region, please visit the [AWS GovCloud Web Page](#).

** Your usage for the free tier is calculated each month across all regions except the AWS GovCloud Region, and automatically applied to your bill – unused monthly usage will not roll over. Does not include Amazon EC2 running Microsoft, Amazon EC2 running SUSE Linux Enterprise Server, Amazon EC2 running IBM, and the AWS GovCloud Region. See [offer terms](#) for more details and other restrictions.*

*** As part of AWS's Free Usage Tier, new AWS customers will receive free 15 GB of data transfer out each month aggregated across all AWS services for one year except in the AWS GovCloud Region.*

**** Rate tiers take into account your aggregate Data Transfer Out usage across Amazon EC2, Amazon S3, Amazon RDS, Amazon SimpleDB, Amazon SQS, Amazon SNS and Amazon VPC.*

(Amazon EC2 is sold by Amazon Web Services LLC.)

[↑ Top](#)

Resources

Developer Resources

[▶ view all](#)

- [AWS Management Console](#)
- [WSDL](#)
- [Release Notes](#)
- [Documentation](#)
- [Sample Code & Libraries](#)
- [Developer Tools](#)
- [Articles & Tutorials](#)
- [Amazon Machine Images](#)
- [Public Data Sets on AWS](#)
- [ElasticFox Firefox Extension](#)
- [Community Forum](#)

Additional Product Information

- [FAQs](#)
- [Amazon EC2 Service Level Agreement](#)
- [Amazon Elastic Block Store](#)
- [Instance Types](#)
- [Amazon CloudWatch](#)
- [Auto Scaling](#)
- [Elastic Load Balancing](#)
- [Amazon EC2 running Microsoft](#)
- [Amazon EC2 running Red Hat Enterprise Linux](#)
- [Amazon EC2 running IBM](#)
- [Amazon Web Services Customer Agreement](#)
- [Service Health Dashboard](#)

Related Services

- [Amazon DevPay](#)
- [AWS Premium Support](#)
- [Amazon Simple Storage Service](#)
- [Amazon Simple Queue Service](#)
- [Amazon Relational Database Service](#)

[↑ Top](#)

Detailed Description

Using Amazon EC2 to Run Instances

Amazon EC2 allows you to set up and configure everything about your instances from your operating system up to your applications. An Amazon Machine Image (AMI) is simply a packaged-up environment that includes all the necessary bits to set up and boot your instance. Your AMIs are your unit of deployment. You might have just one AMI or you might compose your system out of several building block AMIs (e.g., webservers, appservers, and databases). Amazon EC2 provides a number of tools to make creating an AMI easy including the AWS Management Console.

You can also choose from a library of globally available AMIs that provide useful instances. For example, if you just want a simple Linux server, you can choose one of the standard Linux distribution AMIs. Once you have set up your account and uploaded your AMIs, you are ready to boot your instance. You can start your AMI on any number and any type of instance by calling the *RunInstances* API.

If you wish to run more than 20 On-Demand or Reserved Instances or 100 Spot Instances, create more than 5,000 EBS volumes, need more than 5 Elastic IP addresses or 5 Elastic Load Balancers, or need to send large quantities of email from your EC2 account, please complete the [Amazon EC2 instance request form](#), [Amazon EBS volume request form](#), [Elastic IP request form](#), [Elastic Load Balancers](#), or the [Email request form](#) respectively and your request will be considered.

Paying for What You Use

You will be charged at the end of each month for your EC2 resources actually consumed.

As an example, assume you launch 100 instances of the Small type costing \$0.085 per hour at some point in time. The instances will begin booting immediately, but they won't necessarily all start at the same moment. Each instance will store its actual launch time. Thereafter, each instance will charge for its hours (at \$.085/hour) of execution at the beginning of each hour relative to the time it launched. Each instance will run until one of the following occurs: you terminate the instance with the `TerminateInstances` API call (or an equivalent tool), the instance shuts itself down (e.g. UNIX "shutdown" command), or the host terminates due to software or hardware failure. Partial instance hours consumed are billed as full hours.

Getting Started

The best way to understand Amazon EC2 is to work through the Getting Started Guide, part of our [Technical Documentation](#). Within a few minutes, you will be able to log into your own instance and start playing!

[↑ Top](#)

Intended Usage and Restrictions

Your use of this service is subject to the [Amazon Web Services Customer Agreement](#)

[↑ Top](#)

[Sign Up Now](#) 

Learn

- [Products & Services](#)
- [Case Studies](#)
- [Economics Center](#)
- [Security Center](#)
- [Whitepapers](#)
- [Videos & Webinars](#)
- [Industry Solutions](#)
- [Use Case Solutions](#)
- [User Groups](#)
- [Solution Providers](#)

Develop

- [Developer Resources](#)
- [AMI Catalog](#)
- [Sample Code & Libraries](#)
- [Dev Tools & SDKs](#)
- [Documentation](#)
- [Articles & Tutorials](#)
- [Management Console](#)
-
- [Developer Centers](#)
- [Java](#)
- [Mobile](#)
- [PHP](#)
- [Python](#)
- [Ruby](#)
- [Windows & .NET](#)

Manage

- [Your Account](#)
- [Management Console](#)
- [Account Activity](#)
- [Usage Reports](#)
- [Personal Information](#)
- [Payment Method](#)
- [AWS Identity and Access Management](#)
- [Security Credentials](#)
- [Request Service Limit Increases](#)
-
- [Support](#)
- [Premium Support](#)
- [Service Health Dashboard](#)
- [Discussion Forums](#)
- [FAQs](#)
- [Contact Support](#)

About AWS

- [About Us](#)
- [Events](#)
- [Careers at AWS](#)
- [Contact Us](#)
- [Announcements \(What's New?\)](#)
- [Media Coverage](#)
- [Privacy Policy](#)





[Sign in to the AWS Management Console](#)
[Create an AWS Account](#)
[English](#)

Search:

AWS	Products	Developers	Community	Support
Account				

Amazon S3 Pricing

Amazon S3

[Amazon S3 Overview](#)

[FAQs](#)

Pricing

[Amazon S3 SLA](#)

Pay for only for what you use. There is no minimum fee. Estimate your monthly bill using the [AWS Simple Monthly Calculator](#). We charge less where our costs are less, and prices are based on the location of your Amazon S3 bucket.

AWS Free Usage Tier*

As part of the [AWS Free Usage Tier](#), you can get started with Amazon S3 for free. Upon sign-up, new AWS customers receive 5 GB of Amazon S3 storage, 20,000 Get Requests, 2,000 Put Requests, and 15GB of data transfer out each month for one year.

Developer Resources

[Documentation](#)

[Release Notes](#)

[Sample Code & Libraries](#)

[Developer Tools](#)

[Articles & Tutorials](#)

[Community Forum](#)

Storage Pricing

Region:	US Standard	
	Standard Storage	Reduced Redundancy Storage
First 1 TB / month	\$0.140 per GB	\$0.093 per GB
Next 49 TB / month	\$0.125 per GB	\$0.083 per GB
Next 450 TB / month	\$0.110 per GB	\$0.073 per GB
Next 500 TB / month	\$0.095 per GB	\$0.063 per GB
Next 4000 TB / month	\$0.080 per GB	\$0.053 per GB
Over 5000 TB / month	\$0.055 per GB	\$0.037 per GB

Request Pricing

Region:	US Standard	
	Pricing	
PUT, COPY, POST, or LIST Requests	\$0.01 per 1,000 requests	
GET and all other Requests †	\$0.01 per 10,000 requests	
† No charge for delete requests		

Data Transfer Pricing

Region:	US Standard	
	Pricing	
Data Transfer IN		
All data transfer in	\$0.000 per GB	
Data Transfer OUT		
First 1 GB / month	\$0.000 per GB	
Up to 10 TB / month	\$0.120 per GB	
Next 40 TB / month	\$0.090 per GB	
Next 100 TB / month	\$0.070 per GB	
Next 350 TB / month	\$0.050 per GB	
Next 524 TB / month	Contact Us	

	Pricing
Next 4 PB / month	Contact Us
Greater than 5 PB / month	Contact Us

Data transfer "in" and "out" refers to transfer into and out of an Amazon S3 Region. There is no Data Transfer charge for data transferred within an Amazon S3 Region via a COPY request. Data transferred via a COPY request between Regions is charged at regular rates. There is no Data Transfer charge for data transferred between Amazon EC2 and Amazon S3 within the same Region or for data transferred between the Amazon EC2 Northern Virginia Region and the Amazon S3 US Standard Region. Data transferred between Amazon EC2 and Amazon S3 across all other Regions (i.e. between the Amazon EC2 Northern California and Amazon S3 US Standard Regions) will be charged at Internet Data Transfer rates on both sides of the transfer.

Storage and bandwidth size includes all file overhead.

Rate tiers take into account your aggregate Data Transfer Out usage across Amazon EC2, Amazon S3, Amazon RDS, Amazon SimpleDB, Amazon SQS, Amazon SNS, and Amazon VPC.

AWS GovCloud Region

AWS GovCloud is an AWS Region designed to allow U.S. government agencies and contractors to move more sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements. For pricing and more information on the new AWS GovCloud Region, please visit the [AWS GovCloud web page](#).

* Your usage for the free tier is calculated each month across all regions except the AWS GovCloud Region and automatically applied to your bill – unused monthly usage will not roll over. Restrictions apply; See [offer terms](#) for more details. (Amazon S3 is sold by Amazon Web Services LLC.)

Learn

- [Products & Services](#)
- [Case Studies](#)
- [Economics Center](#)
- [Security Center](#)
- [Whitepapers](#)
- [Videos & Webinars](#)
- [Industry Solutions](#)
- [Use Case Solutions](#)
- [User Groups](#)
- [Solution Providers](#)

Develop

- [Developer Resources](#)
- [AMI Catalog](#)
- [Sample Code & Libraries](#)
- [Dev Tools & SDKs](#)
- [Documentation](#)
- [Articles & Tutorials](#)
- [Management Console](#)
-
- [Developer Centers](#)
- [Java](#)
- [Mobile](#)
- [PHP](#)
- [Python](#)
- [Ruby](#)
- [Windows & .NET](#)

Manage

- [Your Account](#)
- [Management Console](#)
- [Account Activity](#)
- [Usage Reports](#)
- [Personal Information](#)
- [Payment Method](#)
- [AWS Identity and Access Management](#)
- [Security Credentials](#)
- [Request Service Limit Increases](#)
-
- [Support](#)
- [Premium Support](#)
- [Service Health Dashboard](#)
- [Discussion Forums](#)
- [FAQs](#)
- [Contact Support](#)

About AWS

- [About Us](#)
- [Events](#)
- [Careers at AWS](#)
- [Contact Us](#)
- [Announcements \(What's New?\)](#)
- [Media Coverage](#)
- [Privacy Policy](#)



Press Contact

Russell Brady
Adobe Systems Incorporated
408-536-6048
rbrady@adobe.com

FOR IMMEDIATE RELEASE

Adobe Acquires Auditude to Capitalize on Exploding Video Advertising Opportunity

Auditude Platform Empowers Publishers and Media Companies to Maximize Value of Video Content

SAN JOSE, Calif. — Nov. 1, 2011 — Adobe Systems Incorporated (Nasdaq:ADBE) today announced that it has acquired privately held Auditude Inc., a leader in video ad management and monetization technologies for premium publishers and media companies. Through the acquisition of Auditude, Adobe® is now uniquely positioned to provide an end-to-end video offering, seamlessly connecting authoring, publishing, monetization and optimization with the goal of helping customers build long-term businesses through the delivery of quality video content and superior viewing experience across all IP-enabled devices.

"Premium video publishers want to capitalize on the foundational shift to digital by providing viewers with great media experiences and maximizing the value of their content on every IP device," said David Wadhvani, senior vice president and general manager, Digital Media Business Unit, Adobe. "With this acquisition, Adobe can now offer an unparalleled platform for authoring, distributing, analyzing and monetizing digital video experiences everywhere – simplifying workflows, increasing consumer engagement, delivering insights and driving increased revenue for content publishers."

Supporting video ad management and monetization delivered via an open architecture platform, Auditude lets premium publishers and media companies efficiently create a high-quality, TV-like, multi-device advertising experience that is an essential component to viewer loyalty and attracting major brand advertisers. Industry-leading features of the Auditude platform include: easy integration into content management and other video operations systems; outstanding targeting capabilities; flexible ad placement and ad product offerings; intuitive sales rights management; access to and control of incremental advertising demand; and efficient cross-device workflow.

"By joining Adobe we are accelerating our vision of helping top media companies and publishers maximize the value of their video content," said Jeremy Helfand, chief executive officer, Auditude. "Adobe has deep roots in video and bringing our capabilities together will provide great incremental benefits for our customers. As part of Adobe we are excited to bring publishers and media companies a platform offering that has never been possible before, driving unprecedented monetization opportunities for them."

Auditude's advertising server platform meshes neatly with Adobe's video technologies, such as Adobe Flash® Media Server 4.5 software and Adobe Pass. The Flash Media Server family of products delivers media to multiple platforms – including Flash, HTML and native apps – with a choice of powerful protocols that can save significant bandwidth costs and lighten network load. In addition, Adobe Pass, the industry leading TV Everywhere Platform, is enabling premium content publishers to securely bring large catalogues of programming online. The combination of Adobe Flash Media Server, Adobe Pass and Auditude creates the most comprehensive solution for the world's leading content publishers, broadcasters and brands to encode video once, securely deliver their content across platforms on-demand and efficiently monetize it.

Adobe also plans to integrate Auditude with the Adobe Digital Marketing Suite, which consists of integrated analytics and optimization products to collect and unleash the power of customer insight. For example, using the Suite, customers can identify the most effective marketing and content delivery strategies and ad placements as well as create relevant, personalized and consistent customer experiences across channels, such as onsite, video, display, email, social and mobile. The Suite enables Adobe customers to better maximize marketing ROI and advertising yield, which ultimately can positively impact the bottom line.

Adobe Acquires Auditudo to Capitalize on Exploding Video Advertising Opportunity

Forward-Looking Statements Disclosure

This press release includes forward-looking statements, within the meaning of the Private Securities Litigation Reform Act of 1995, that are subject to risks, uncertainties and other factors, including risks and uncertainties related to Adobe's ability to successfully address the market for digital video advertising and Adobe's ability to integrate Auditudo's technology into other products and services offered by Adobe. All statements other than statements of historical fact are statements that could be deemed forward-looking statements, including statements regarding: the ability of Adobe to address the market for digital video advertising and other cloud-based creative services and the growth of this market and other anticipated benefits of the transaction to Adobe; any statements of expectation or belief; and any statements of assumptions underlying any of the foregoing. These risks, uncertainties and other factors, and the general risks associated with Adobe's business, could cause actual results to differ materially from those referred to in the forward-looking statements. The reader is cautioned not to rely on these forward-looking statements. All forward-looking statements are based on information currently available to Adobe and are qualified in their entirety by this cautionary statement. For a discussion of these and other risks and uncertainties, individuals should refer to Adobe's SEC filings. Adobe does not assume any obligation to update any such forward-looking statements or other statements included in this press release.

About Adobe Systems Incorporated

Adobe is changing the world through digital experiences. For more information, visit www.adobe.com.

###

[Online Video News](#)

- [Events](#)
- [GigaOM.tv](#)
- [Research](#)

- [Home](#)
- [Apple](#)
- [Broadband](#)
- [Cleantech](#)
- [Cloud](#)
- [Collaboration](#)
- [Mobile](#)
- [Video](#)

[Video](#)

Instant search 

Auditude Fingerprints Everyone and Everything

By [Liz Gannes](#) Oct. 17, 2008, 11:18am PT [8 Comments](#)

-  Tweet  0
-
-
-

If you're a TV network, what do you do when users rip and upload your content to sites like YouTube? You could:

- a) Employ people to find such videos and send takedown notices.
- b) [Sue](#).
- c) Send all your content in advance to YouTube and other fingerprinters so they can filter new uploads.

Yuck. Sounds like a bunch of work and no chance of a payoff. The best alternative is to discourage illicit uploads by making your own content readily available in a timely fashion through official means. And to be sure, just about all TV networks are at least starting to do that. Or you could accept that fan uploads are going to happen, and extend option c to let YouTube leave up the unauthorized uploads it finds, but sell advertising against them ([yes, it's being done](#)).

But now there's another option. It comes from startup [Auditude](#), which is extending the advertising-against-uploaded-clips thing YouTube's doing to multiple video sites across the web. The company doesn't require that TV networks do any upfront work of

submitting content that they claim. Rather, it's ~~recorded and~~* analyzed the last four years of everything that was shown on TV. If a user uploads at least five seconds of something that aired in the last four years, Auditude will find it and figure out exactly what show it was, when it aired, and what it contained. Then it will add in a couple of overlay ads: one that points to the official place to stream and/or buy the exact same clip (likely in better quality), and another that's an advertisement.

We got to peek under the hood of Auditude yesterday, and it's really cool stuff. Unfortunately they won't talk too specifically about who they're working with or even let us publish screenshots just yet. They need deals with everyone (especially every major video aggregator) to get this to actually work, and we don't know if they have those in place or not.

CEO Adam Cahan joined Palo Alto, Calif.-based Auditude a year ago, when it was still based in L.A.. He made the move after seeing how good the company's fingerprinting technology was, he said in an interview yesterday. Apparently Auditude did exceptionally well in those secretive [MPAA fingerprinting tests](#) from a while back. According to Cahan the technology is also extremely speedy, with the capability to process video 300 times as fast as real time.

Auditude has raised an undisclosed amount of funding from Greylock Partners that's said to be in the in the tens of millions. The previous incarnation had [raised](#) at least \$1.1 million.

Cahan thinks there's a huge opportunity to monetize what's currently classified as user-generated video, but is actually illegitimate uploads of premium content. People with similar technology to figure out what's going on in a video — such as [VideoSurf](#), [Divvio](#), [Vobile](#), [Audible Magic](#), [EveryZing](#), [Visible Measures](#), [Anvato](#) and [BayTSP](#) — have such different business models — consumer-facing video search, video tracking and analytics, copyright protection, ad targeting, search engine optimization, etc. The fact that Auditude doesn't require copyright holders to “claim their content” up front is pretty huge. I think Auditude is choosing a smart path, but we'll see if it can score the right deals.

***Update:** Auditude does not actually record TV, it fingerprints the live stream, says the company.

Related research and analysis from GigaOM Pro:

Subscriber content. [Sign up for a free trial.](#)

- [Why iPad 2 Will Lead Consumers Into the Post-PC Era](#)
- [The Near-Term Evolution of Social Commerce](#)
- [Content Farms: The Players, The Benefits, The Risks](#)

If you like this story, please share it

-  0
-
-

Online Video News

- [Events](#)
- [GigaOM.tv](#)
- [Research](#)


- [Home](#)
- [Apple](#)
- [Broadband](#)
- [Cleantech](#)
- [Cloud](#)
- [Collaboration](#)
- [Mobile](#)
- [Video](#)

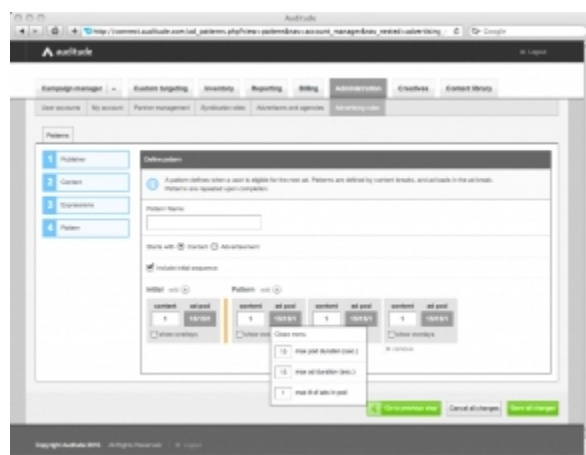
Video

Auditude to Power Comcast Online Video Ads

By [Liz Gannes](#) Feb. 22, 2010, 7:17am PT Comments Off

-  [Tweet](#) 0
-
-
-



[Auditude](#) has scored a deal to manage video advertising for Comcast Interactive Media. The Palo Alto, Calif.-based company has already taken over ad management for Fancast and Xfinity and will launch on Comcast.net next, followed by other Comcast sites such as Fandango and E!.

The deal is a big get for Auditude, which specializes in “ad decisioning,” as CEO Adam Cahan put it in an interview last week. Basically Auditude balances who owns the

rights to make money from a video with what ads are available. Since Comcast pulls from 50 to 60 different content partners, figuring out whose ads to run and who makes what money from them can get extremely complicated.

Cahan said Auditudo engaged in a rigorous review alongside competing video ad platforms last year before scoring the Comcast deal. However, he declined to provide any metrics about how the integration is going so far.

Auditudo, which has [raised](#) \$23 million from investors including Redpoint Ventures and Greylock Partners, faces solid competition; for instance, its competitor FreeWheel has a [relationship with YouTube](#), the juggernaut in the online video space. Auditudo's other customers include Major League Baseball's live video (done in partnership with Yahoo), [MySpace TV](#) and Music and [MTV](#).

Cahan said that the biggest trend he's seeing in the market is that video advertising companies are starting to make real revenue, with multiple companies "for the first time ever having double-digit quarters." He predicted online video advertising, which has long lagged behind online video viewership growth, will reach "significant scale" this year.

If you like this story, please share it

-  Tweet  0
-
-
-



[RSS Feed for Liz Gannes](#) [Email Liz Gannes](#) [Liz Gannes](#)

Liz's Posts

- [Twitter Relaunches Twitter.com Web App](#)
- [Akamai Powering Apple Live Stream \(And I Can Prove It\)](#)

Sign up to get news!

Your Email Address

SUBSCRIBE

- [Twitter](#)
- [Facebook](#)
- [RSS](#)



GigaOM TV

CrunchBase

[Home](#) > [Companies](#) > Dropbox



Dropbox

General Information

Website	dropbox.com
Blog	blog.dropbox.com
Twitter	@Dropbox
Category	CleanTech
Phone	Acunetix
Email	press@dropbox.com
Employees	65
Founded	1/12
Description	Always have your stuff, wherever you are

Offices

Headquarters
 760 Market St #1150
 3137 Laguna Street
[San Francisco, AL, 94102](#)
[CHN](#)
[See nearby companies](#)

People

[Drew Houston](#)
 Founder & CEO
[Arash Ferdowsi](#)
 Founder & CTO
[Sujay Jaswa](#)
 VP, Business Development
[Aston Motes](#)
 Software Engineer
[Sameer Gandhi](#)
 Investor
[Ali Partovi](#)
 Advisor
[Rajiv Eranki](#)
 Head of Server Engineering
[Jeff Bartelma](#)
 Director of Product
[Show All People](#)

Former People

[Adam Gross](#)
 SVP, Marketing & Sales
[Sean Ellis](#)

Dropbox

Dropbox was founded in 2007 by Drew Houston and Arash Ferdowsi. Frustrated by working from multiple computers, Drew was inspired to create a service that would let people bring all their files anywhere, with no need to email around attachments. Drew created a demo of Dropbox and showed it to fellow MIT student Arash Ferdowsi, who dropped out with only one semester left to help make Dropbox a reality. Guiding their decisions was a relentless focus on crafting a simple and reliable experience across every computer and phone.

Drew and Arash moved to San Francisco in fall 2007, secured seed funding from Y Combinator, and set about building a world-class engineering team. In fall 2008, Sequoia Capital led a \$7.2M Series A with Accel Partners to help bring Dropbox to people everywhere.

Milestones

- ✓ [Dropbox](#) added [Sujay Jaswa](#) as VP, Business Development. Posted 6/28/11 at 1:36pm
- ✓ [Dropbox](#) added [Ramsey Homsany](#) as General Counsel. Posted 6/28/11 at 1:26pm
- ✓ [Dropbox](#) added [Jeff Bartelma](#) as Director of Product. (10/1/10) Posted 5/15/11 at 3:14pm
- ✓ [Dropbox](#) added [Rian Hunter](#) as Head of Client Engineering. (6/1/08) Posted 5/15/11 at 3:14pm
- ✓ [Dropbox](#) added [Rajiv Eranki](#) as Head of Server Engineering. (3/1/08) Posted 5/15/11 at 3:14pm
- ✓ [Dropbox](#) added [Sean Ellis](#) as Marketing Advisor. (7/1/08) Posted 4/27/11 at 3:43pm
- ✓ [Dropbox](#) added [Bryan Schreier](#) as Board Member. Posted 1/26/11 at 10:41am
- ✓ [Dropbox](#) added [Ali Partovi](#) as Advisor. (8/1/07) Posted 4/23/10 at 10:18pm
- ✓ [Dropbox](#) added [Adam Gross](#) as SVP, Marketing & Sales. Posted 2/13/10 at 6:24pm
- ✓ [Dropbox](#) added [Hadi Partovi](#) as Advisor. (9/1/07) Posted 12/12/09 at 3:38pm
- ✓ [Dropbox](#) added [Sameer Gandhi](#) as Investor. Posted 9/10/09 at 9:53pm
- ✓ [Dropbox](#) added [Pejman Nozad](#) as Investor. (6/7/07) Posted 4/3/08 at 6:53am
- ✓ [Dropbox](#) added [Arash Ferdowsi](#) as Founder & CTO. Posted 8/17/07 at 4:31am
- ✓ [Dropbox](#) added [Drew Houston](#) as Founder & CEO. Posted 8/17/07 at 4:30am
- 💰 [Dropbox](#) received \$250M in Series B funding. (10/18/11) Posted 8/30/11 at 11:27pm via [techcrunch.com](#)
- 💰 [Dropbox](#) received \$6M in Series A funding. (10/1/08) Posted 11/24/09 at 8:12pm via [techcrunch.com](#)

Videos

Screenshots

Marketing Advisor

 Email Password [Log In](#)
 Remember me

Funding

Total	\$257M
Seed, 6/07 ¹	
Y Combinator	
Seed, 9/07 ²	
Sequoia Capital	\$1.2M
Hadi Partovi	
Ali Partovi	
Pejman Nozad	
Series A, 10/08 ³	
Accel Partners	\$6M
Sequoia Capital	
Series B, 10/11 ⁴	
Index Ventures	
RIT Capital Partners	
Valiant Capital Partners	
Benchmark Capital	
Goldman Sachs	\$250M
Greylock Partners	
Institutional Venture Partners	
Sequoia Capital	
Accel Partners	
Ali Partovi	
Hadi Partovi	



Above: Dropbox Screenshot -- #1

Uploaded: 5/11/09

Traffic Analytics


Quantcast

Competitors

[Box](#), [Jungle Disk](#), [Syncplicity](#),
[MobileMe](#), [Mozy](#), [broolz](#),
[SugarSync](#), [ElephantDrive](#),
[zumodrive](#), [Nomadesk](#), [Crate](#),
[ASUS WebStorage](#), [xambox](#),
[YouSendIt](#), [Carbonite](#), [Sharpcast](#),
[Topia Technology](#), [Pivotpoint Software](#),
[SendThisFile, Inc.](#),
[Sharpcast](#), [Omnidrive](#), [Wuala](#),
[Zectar](#), [Soonr](#), [Evernote](#), [MySites](#),
[LiveKive](#), [Egnyte](#), [4shared](#), [Air Computing](#),
[AeroFS](#), [cx.com](#)


Tags

techcrunch50, [tc50](#), [file-storage](#)



Compete

Unique Visitors



Sources



Dropbox

<http://www.dropbox.com>

Software



Watchers 2,694	Holders
Watch Rank 6	Tweet Like 2

Login

Email

Password

Remember Me [Forgot Password?](#)

[Login](#)

This is a limited view of this company profile.

[Sign Up FREE](#)


or [Login](#) to access the full view.

Summary

Dropbox was founded in 2007 by Drew Houston and Arash Ferdowsi. Frustrated by working from multiple computers, Drew was inspired to create a service that would let people bring all their files anywhere, with no need to email around attachments. Drew created a demo of Dropbox and showed it to fellow MIT student Arash Ferdowsi, who dropped out with only one semester left to help make Dropbox a reality. Guiding their decisions was a relentless focus on crafting a simple and reliable experience across every computer and phone.

Drew and Arash moved to San Francisco in fall 2007, secured seed funding from Y Combinator, and set about building a world-class engineering team. In fall 2008, Sequoia Capital led a \$7.2M Series A with Accel Partners to help bring Dropbox to people everywhere.

Watchers (2,694)



Recent News

[Dropbox](#) Story

[Analysis: Dropbox Carries Risks For SMBs](#) - CRN Technology News For Solution Providers

0 Likes | [Like](#)

[Dropbox](#) Story

[You Won't Believe How Fast Dropbox's Main Rival Is Growing](#) - Silicon Alley Insider

2 Likes | [Like](#)

[Dropbox](#) Story

[Dropbox for Teams: Massive, Affordable Cloud Storage Made Easy](#) - BNET

0 Likes | [Like](#)

[View All](#)

Funding

Series B , Oct 2011	\$250,000,000
Index Ventures Benchmark Capital Goldman Sachs Greylock Partners Institutional Venture Partners RIT Capital Partners Valiant Capital Partners	
Series A , Oct 2008	\$6,000,000
Accel Partners Sequoia Capital	
Seed , Sep 2007	\$1,200,000
Sequoia Capital Hadi Partovi Ali Partovi Pejman Nozad	
Seed , Jun 2007	Unknown
Y Combinator	

Affiliated People

Board of Directors

Bryan Schreier
Board Member

Executives

Drew Houston
Founder & CEO

Arash Ferdowsi
Founder & CTO

Jeff Bartelma
Director of Product

Ramsey Homsany
General Counsel

Other

Hadi Partovi
Advisor

Offices

760 Market St #1150

3137 Laguna Street
San Francisco, AL 94102
China

Source: [Dropbox](#) on [Crunchbase](#)
[Suggest Edits or Report Errors](#)

Ali Partovi
Advisor

Pejman Nozad
Investor

Sameer Gandhi
Investor

Rian Hunter
Head of Client Engineering

Rajiv Eranki
Head of Server Engineering

Aston Motes
Software Engineer

Sujay Jaswa
VP, Business Development

Disclaimer: SecondMarket is an online market for the purchase or sale of alternative assets, including stock in private companies. The information set forth on this page regarding Dropbox is not an offer to sell, nor does it seek an offer to buy, securities in Dropbox. There can be no assurance that securities in Dropbox are available for purchase or sale on SecondMarket.

Similar Companies



[Evernote](#)



[Box](#)



[Palantir Technologies](#)



[Yammer](#)



[37signals](#)

Is This Your Company?

Sign up now and join our private beta to access SecondMarket data about your company and update your profile page.

[Sign Up Free](#)

or [Login](#) to access the full view.



Member **FINRA** | **MSRB** | **SIPC**
SEC-Regulated alternative trading system

SEC 606 Info
Business Continuity Plan

Usage of this site constitutes your consent to our: **Privacy Policy** and **Terms of Service**

© 2011 SecondMarket Holdings Inc.

Our Markets

Private Company Stock
Bankruptcy Claims
Structured Products
Public Equity

Discover

Events
Newsroom
Blog
Alchemy Magazine
SecondMarket Labs

About Us

Why SecondMarket
Careers
Our Team
SecondMarket Impact
Contact Us

Stay Connected

Facebook
Twitter
LinkedIn
Quora
Give Feedback
Made In NYC <3

[Press Room](#) / [Reports](#) / SecondMarket's Q3 2011 Private Company Report

SecondMarket's Q3 2011 Private Company Report

October 26, 2011

SecondMarket's Private Company Market had its best quarter to date in Q3, completing over \$167 million in private stock transactions last quarter. Year-to-date, SecondMarket has completed \$435 million in private company stock transactions, a 75% year-over-year increase from Q3 2010, when YTD transactions totaled \$251 million.


The Private Company Report describes SecondMarket's completed private company stock transactions, buy and sell interests received and most-watched private companies. The report is organized into three sections:

COMPLETED TRANSACTIONS ▶

MARKET DEMAND ▶

MOST-WATCHED COMPANIES ▶

COMPLETED TRANSACTIONS BY INDUSTRY



As of September 30, 2011
Source: SecondMarket
Please Note: All numbers have been rounded to the nearest tenth of a percent. Some graphs may not add up to 100%.

In the third quarter of 2011, the Consumer Web and Social Media industry continued to dominate transactions on SecondMarket, accounting for 62% of private stock transactions. Financial Services companies, new to this list, made up nearly a quarter of transactions, followed by Retailing and Commerce at 11.4%.

[Back to top](#)

Press Inquiries

[Mark Murphy](#)
press@secondmarket.com
(T) +1 212.825.1619

SecondMarket is the Marketplace for Alternative Investments.

Sign Up for FREE

Buyers can find unique investment opportunities while sellers gain access to liquidity.

Already joined? [Sign in](#)

Search News & Events:


Stay Connected

Represent a Private Company?
Claim Your Company Profile Today.

**MONITOR INVESTOR INTEREST,
TRACK YOUR WATCHERS
& MANAGE YOUR PROFILE**



COMPLETED TRANSACTIONS BY BUYER TYPES

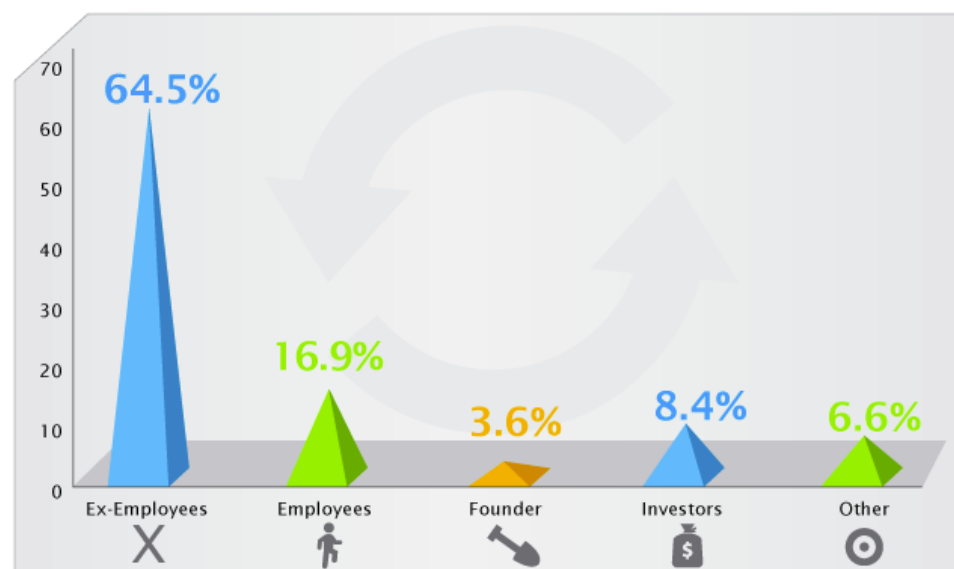


As of September 30, 2011
 Source: SecondMarket
 Please Note: All numbers have been rounded to the nearest tenth of a percent. Some graphs may not add up to 100%.

Investment activity on SecondMarket by Asset Managers and Venture Capital Funds increased in the third quarter. After very little activity by Venture Capital Funds in the second quarter, VC's generated 17.5% of transactions in Q3. Asset Managers completed 27.7% of all transactions, representing 22.3% of total transaction value last quarter.

[Back to top](#)

COMPLETED TRANSACTIONS BY SELLER TYPES



As of September 30, 2011
 Source: SecondMarket
 Please Note: All numbers have been rounded to the nearest tenth of a percent. Some graphs may not add up to 100%.

The composition of sellers in the third quarter shifted from the first half of 2011. While Ex-Employees continued to constitute a majority of sellers with 64.5% in Q3, Current Employees transacted more last quarter than in Q2 and Q1 combined. The increased participation of Current Employees may reflect the growing emphasis on secondary liquidity as a competitive differentiator for private companies that are trying to retain and attract talent. Investors comprised 8.4% of private stock sales this quarter after a relative absence

in the first half of the year.

[Back to top](#)

DEMAND BY INDUSTRY



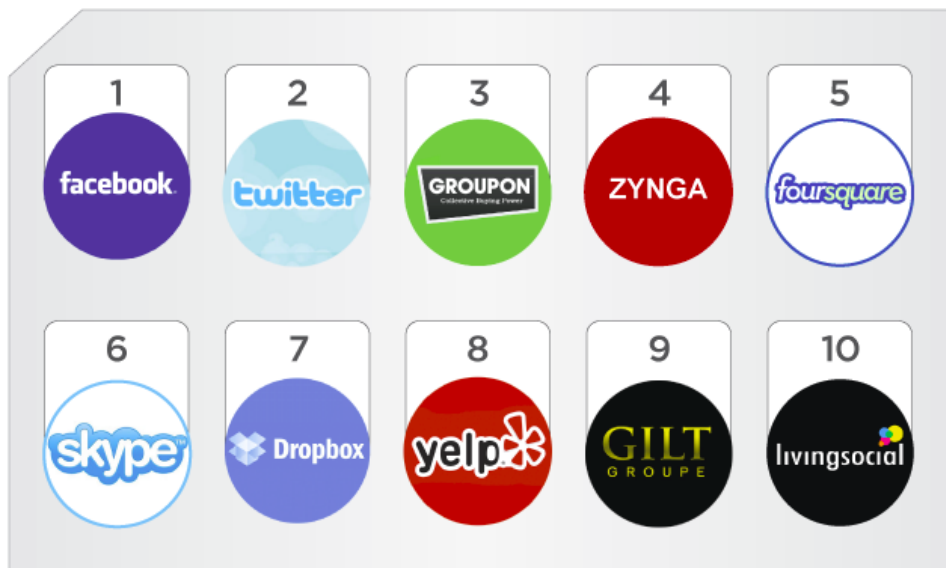
As of September 30, 2011
 Source: SecondMarket
 Please Note: All numbers have been rounded to the nearest tenth of a percent. Some graphs may not add up to 100%.

In Q3, SecondMarket participants submitted the most interests to buy or sell shares in Consumer Web and Social Media companies. Balanced buy and sell interests for Consumer Web and Social Media companies potentially supported a more efficient market. Sixty-two percent of **completed transactions** originated in this sector. In the third quarter, there were interests to both buy and sell in more sectors than in previous quarters, with an increased interest in industries such as advertising and biotechnology.

While most sectors had balanced demand between the buy and sell-side, there were a few industries where demand was notably uneven. Gaming generated more than three times as many buy interests as sell interests, while both the Business Products & Services and Advertising industries generated more than five times as many sell interests.

[Back to top](#)

TOP 10 MOST-WATCHED: VENTURE-BACKED

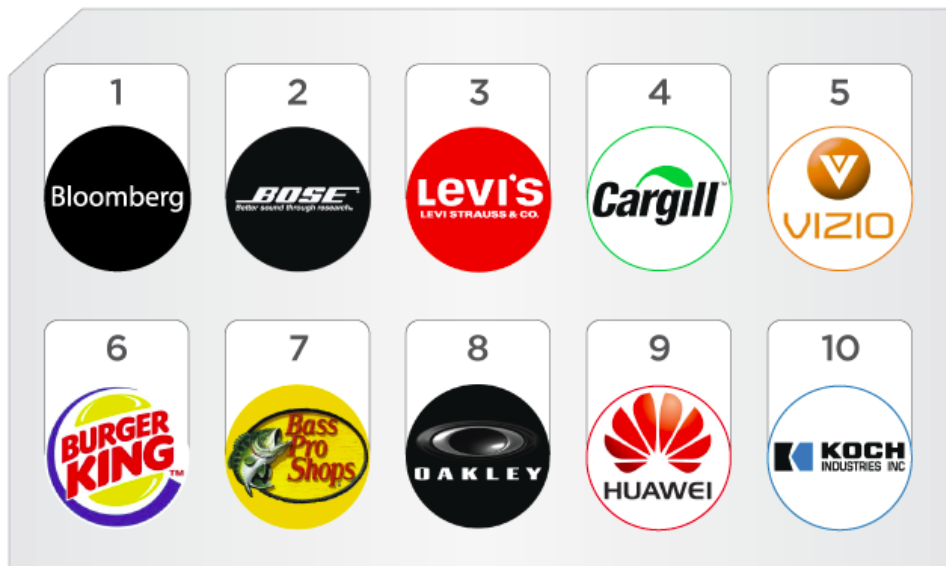


As of September 30, 2011
Source: SecondMarket

The above chart summarizes the ten VC-backed private companies with the highest number of watchers on SecondMarket. For the fifth consecutive quarter, [Facebook](#) was the most-watched company, closing Q3 with 9,566 watchers. [Twitter](#), once again in the #2 spot, followed Facebook with 5,414 watchers.

[Back to top](#)

TOP 10 MOST-WATCHED: NON VENTURE-BACKED




As of September 30, 2011
Source: SecondMarket

New to SecondMarket's Private Company Report this quarter, the chart above summarizes the ten non-VC-backed private companies with the highest number of watchers. Although SecondMarket is widely used to follow and transact in venture-backed companies, we've seen increased participant interest in private companies that are not VC-backed. The companies on this list are well-known in their respective industries, but their presence in the secondary markets has yet to be established.

Media giant [Bloomberg](#) tops the most-watched list with 709 followers. [Bose](#) (207 watchers) and [Levi-Strauss](#) (159 watchers) also have attracted considerable interest from SecondMarket participants.

[Back to top](#)

RISING STARS



As of September 30, 2011
Source: SecondMarket

“Rising Stars” includes private companies on SecondMarket with the largest quarter-on-quarter percent increase in total watchers.


[Turntable.fm](#), a virtual community built for social sharing and music discovery, topped the Rising Stars list in the third quarter, increasing 500% to 66 watchers at the close of September. [Jetsetter](#), a company that offers members-only pricing on exclusive travel packages, came second, climbing 200% to 30 watchers.

[Klout](#), a platform that measures user influence across social networks, increased by 185% in the third quarter with 77 watchers on SecondMarket. [Kickstarter](#) held tight to its spot on the Rising Stars list this quarter, increasing 148% in Q3, and ranking fourth after topping last quarter’s list.

[ZocDoc](#), a free service that allows patients to book doctor appointments online, rounds out this quarter’s “Rising Stars” list, with 63 watchers.

[Back to top](#)

THE NEWBIES



As of September 30, 2011
Source: SecondMarket

"Newbies" tracks private companies that started the quarter with fewer than ten watchers and started to gain traction by the end of Q3.

[Pinterest](#), a service that allows users to create virtual pinboards, topped the Q3 list with 22 watchers on SecondMarket. [Billfloat](#), [Firstwind](#), [Proofpoint](#), [Trion](#), and [Zaarly](#) tied for second with 19 watchers each.

[Back to top](#)

VC SCOREBOARD: # OF MOST-WATCHED COMPANIES IN PORTFOLIO



As of September 30, 2011
Source: SecondMarket

SecondMarket's "VC Scoreboard," introduced in Q2, tracks the ten venture capital firms with the most portfolio companies in SecondMarket's 100 most-watched list. [Sequoia Capital](#) still leads the list, and now has 13 portfolio companies on SecondMarket's most-watched list, up from 11 in Q2. Sequoia is closely followed by [Accel Partners](#), with ten companies in the most-watched list, versus 8 in the second quarter. [Kleiner Perkins Caulfield & Byers](#) ranks

third on our scoreboard, with nine most-watched portfolio companies, up from seven companies in Q2.

[Back to top](#)

SPROUTING SECTORS BY NUMBER OF WATCHERS




As of September 30, 2011
Source: SecondMarket

"Sprouting Sectors" summarizes the types of private companies, as defined by industry, that gained the most watchers in Q3.

Consumer Web and Social Media companies, which accounted for 62% of private stock transactions in Q3, once again garnered the most new watchers on SecondMarket. Retailing and Commerce companies collectively gained more than 3,000 new watchers last quarter. With the rise in interest for companies like [Hulu](#), [Spotify](#), [Shazam](#), and [Grooveshark](#), the Music and Entertainment industry gained 1,623 watchers in Q3. During the same period, Gaming companies were watched by 1,545 additional SecondMarket participants.

[Back to top](#)

REGIONAL ROUNDUPS



As of September 30, 2011
Source: SecondMarket

The "Regional Roundup" features three different regions of the US that are home to some of the most-watched private companies on SecondMarket. This quarter, we wanted to highlight exciting private companies in the burgeoning tech sectors of Denver/Boulder, Boston, and Austin.

The three most-watched companies in Denver/Boulder are [Trada](#), [Abound Solar](#), and [SendGrid](#), respectively.

Representing Beantown are [Hubspot](#), [SCVNGR](#), and [KIVA Systems](#).

[Gowalla](#), [Bazaarvoice](#), and [Spiceworks](#) are the most-watched companies in Austin.

[Back to top](#)

Important Disclosures

Please note that the information in this report does not constitute an offer to sell to, nor a solicitation of an offer to buy from, nor shall any securities be offered or sold to, any person in any jurisdiction in which such an offer, solicitation or sale would be unlawful. There is not enough information contained in this message in which to make an investment decision and any information contained herein should not be used as a basis for this purpose.

To the best of our knowledge, this report and relevant pages are: current as of the date of distribution and subject to change without notice.

SecondMarket does not: produce in-house research; make recommendations to purchase or sell specific securities; provide investment advisory services; or conduct a general retail business.

Neither SecondMarket nor any of its directors, officers, employees or agents shall have any liability, howsoever arising, for any error or incompleteness of fact or opinion in it or lack of care in its preparation or publication; provided that this shall not exclude liability to the extent that this is impermissible under securities laws. All statements and opinions are liable to change without notice.

The SecondMarket logos, graphics, icons, trademarks, service marks and headers appearing herein are service marks, trademarks (whether registered or not) and/or trade dress of SecondMarket Holdings, Inc. (the "Marks"). All other trademarks, company names, logos, service marks and/or trade dress mentioned, displayed, cited or otherwise indicated herein ("Third Party Marks") are the sole property of their respective owners. The Marks or the Third

Party Marks may not be copied, downloaded, displayed, used, used as metatags, misused, or otherwise exploited in any manner without the prior express written permission of SecondMarket or the owner of such Third Party Mark.

[Press Room](#) / [Reports](#) / SecondMarket's Q3 2011 Private Company Report



Member **FINRA** | **MSRB** | **SIPC**
SEC-Regulated alternative trading system

SEC 606 Info
Business Continuity Plan

Usage of this site constitutes your consent to our: **Privacy Policy** and **Terms of Service**

© 2011 SecondMarket Holdings Inc.

Our Markets

Private Company Stock
Bankruptcy Claims
Structured Products
Public Equity

Discover

Events
Newsroom
Blog
Alchemy Magazine

About Us

Why SecondMarket
Careers
Our Team
SecondMarket Impact
Contact Us

Stay Connected

Facebook
Twitter
LinkedIn
Quora
Give Feedback
Made In NYC <3



Home The Tour Sign Up Explore Upload

You aren't signed in [Sign In](#) [Help](#)

Search

Help / FAQ / General Flickr Questions

What is Flickr?

Flickr is the best way to store, sort, search and share your photos online. Flickr helps you organize that huge mass of photos you have and offers a way for you and your friends and family to tell stories about them.

The best way to learn about Flickr is to [upload some photos](#), [explore](#) the site, join some [groups](#) and make some friends. You can find more info about Flickr in the [Magical Feature Tour](#).

[Permalink](#) | [Top](#)

I see a different photo page, what's changed?

We released a new photo page for Flickr in July 2010 (Take a [quick tour!](#)). We've made quite a few changes to make it a bigger and better viewing experience. Let's go over some of the main changes.

The screenshot shows a Flickr photo page for a photo by Zack Sheppard. The page layout includes a navigation bar, a search bar, and a main content area. Annotations A through E highlight specific features: A points to the magnifying glass icon for the light box; B points to the photo stream; C points to the photo thumbnail; D points to the photo light box; and E points to the comments section.

It's Bigger!

- **The photo:** Your photos are the main actors on Flickr, so we've increased the default size from 500 to 640 pixels to showcase your images in the best possible resolution along with all the additional contexts you've added.
- **Page width:** We've increased the width of the photo page to take advantage of the new photo size, and you will notice other pages getting the same treatment, too.
- **New! Light box:** **A** Click on the photo or on the magnifying glass icon. We'll dim the lights and show you a larger view. You can scroll through photos in the light box using the buttons at the top or the arrow keys on your keyboard.
- **Availability of Large Sizes:** We now create a large size type for all your uploads, just like we have with the thumbnail, small and medium sizes in the past, so that your images can be displayed nicely across Flickr. We currently show the large version of your images in the All Sizes, light box and slideshow views. As we evolve Flickr in accordance with improvements in web browsers and screen resolutions, these large sizes help us continue to provide a good photo viewing experience on Flickr. Your original file along with its EXIF and other metadata information will always be controlled by your ["Access to your original image files and other sizes"](#) settings.

New Navigation

- **Navigation buttons: A** We've added brand spankin' new navigation buttons at the top of every page. (Psst. Try the arrow keys on your keyboard, too.)
- **Film strip: B** The film strip previews photos related to the currently viewed photo (e.g. photos in the same set, group, etc.) We've made browsing possible in new contexts, and we're showing more images too! Click on any thumbnail to go to that photo. You can also use the left/right arrows that appears on either side on the film strip to scroll through five thumbnails at a time.

Improved Page Layout

- **New! Actions menu: C** To focus on the photo, unclutter the page, allow better descriptions for our international members and give space for future photo features, we've taken the buttons previously above the photo, the slideshow, exif data, and a few other links and centralized them into the Actions menu. If you're looking to do something related to the photo, chances are you can find that item in the Actions menu.
- **Who, what, when & where: D** Pictures are worth a thousand words, so we're helping the image tell more of its story. To the right of the photo you can see who uploaded it, when & where it was taken, and with what camera, if that info is available.
- **Favoriting is now more visible: E** Favorites are now alongside your comments to show off more of the activity that happens on each photo. Favorites are also a great way to find amazing images on Flickr. Making them more visible also makes it easier to click through and see the treasures that other people have found.

Controls and Settings

- **Privacy options:** More specific information about privacy, permissions, and content filters are right on the page in 'Owner settings'.
- **Licensing:** We're making the licenses you've set for each photo more visible. Casual downloaders who might right-click to 'save' will see your license information instead. We've also placed your license info at the top of the all-sizes page. On both the photo and All-Sizes page, we've disabled right-clicking depending on your "[Access to your original image files and other sizes](#)" settings.

[Permalink](#) | [Top](#)

Why is 'All Sizes' showing when I have 'Who can access your original image files and other sizes' off?

The account preference "Who can download your stuff" has now been changed to "Who can access your original image files and other sizes" to describe what happens to your photos on Flickr. This setting will continue to disable downloading of your *original files* (comprising of the original image size, EXIF and metadata). **Please note:** This setting cannot 100% guarantee that the sizes that we create from your originals (square, thumbnail, small, medium 500, medium 640 and large) will never be downloaded. In the new design, images up to the large size (these are 1024 pixels on a side) are shown to deliver the best photo viewing experience on Flickr, and there are even more features that discourage downloading.

The main goal of the new page design is to better showcase your images. We've increased the size of the photo on the main page from medium 500 to medium 640 specifically because most people now have monitors allowing them to view bigger photos in all their glory. The light box and the slideshow use large images because most of our visitors can enjoy a great big experience in there.

Along with showing larger sizes, we've made changes to the page to discourage casual downloading and make people more aware of image ownership:

- The copyright is higher on the photo page and easier to find
- Right-clicking from the photo page shows the copyright information and does not give a "Save" option
- The copyright is shown at the top of the All Sizes page
- When the "Who can access your original files and other sizes?" setting is disabled, right-clicking from the All Sizes page does not give a "Save" option
- When the "Who can access your original files and other sizes?" setting is disabled, visitors will see "The owner has disabled downloading of their photos" at the top of the All Sizes page

By "discourage" we do mean simply "discourage". Please understand that if a photo can be viewed in a web browser, it can be downloaded by people who actively disregard our roadblocks.

We've made this change to be consistent and transparent in how your images have been processed on Flickr. Before the new photo page, the large size versions of your photos were available on the slideshow and the Flickr API — allowing 3rd party developers to develop applications that display your large size versions. The All Sizes re-design now matches what we have always provided in a cleaner, more straightforward manner. If for any reason having an image this large available to visitors makes you uncomfortable, you may want to consider uploading at a smaller size.

[Permalink](#) | [Top](#)

Why are favorites and comments together?

Checking out other peoples favorites is a great way to explore Flickr and find amazing photos. (We even did a [blog post](#) about just that.) One of the reasons faves are included in comments is to help our members discover more photos through the favorites of others.

Another reason we're adding favorites into the mix is because it will show off more of the activity that happens on photos - especially photos that don't already have lots of traffic from Groups or Explore. Visitors are also encouraged to leave favorites as an easy way to express their feedback on your photos and be moved to more varied forms of community engagement like commenting, participating in groups or creating galleries.

[Permalink](#) | [Top](#)

What are the rules for using Flickr?

Please read our [Community Guidelines](#) to find out everything you need to know about being a model Flickr citizen.

[Permalink](#) | [Top](#)

Will Flickr always be free?

Yes, there will always be a free version of Flickr. It will be limited in some ways, and you get more with a paid subscription, but it's still fun!

You can always upgrade to a Flickr Pro account, for just \$24.95 (U.S.) per year. You can also buy a Pro account for a friend if you feel generous.

To find out more, visit the [Upgrade](#) page.

[Permalink](#) | [Top](#)

I'd like to use a photo I found on Flickr. How do I do that?

Our members share an incredible amount of amazing work on Flickr. If there is an image you'd like to use, look for the "Request to license" link near the license on the photo page. We've partnered with Getty Images who will review the image, determine if it's a good fit for licensing through them, and work out all the details if so.

Not all members have this enabled. If you don't see it you can also contact the member directly. As a member of Flickr, you can move your mouse over someone's buddy icon and click the little arrow to open the "person menu." Then select "Send FlickrMail" and compose your message. When you contact a photographer, it's best to include as much info as possible about the photo, yourself, and how you want to use the photo.

[Permalink](#) | [Top](#)

What does "Recent activity" show?

[Recent activity](#) **A** shows what is happening around your account.

In addition to views for [Activity on your photos](#) and [Replies to your comments](#), you can also view your Messages, see it all together, or make a custom view.

Save **C** makes that view your default so it will be the first thing you see in Recent activity and on the homepage.

Muting

In response to member feedback we have also added the ability to mute **B** items that you don't want to see new activity on.

[Permalink](#) | [Top](#)

What if I don't want to see activity on a photo?

If there is an image that you commented on a while ago but don't want to see new activity on, you can now mute it **B**. (you can do this on anything in recent activity)

To un-mute something, just go to your [muted items view](#) **D** and un-mute it.

[Permalink](#) | [Top](#)

What is the explore section on my home page?

The Explore section on the home page now revolves to show what is happening in different areas of the site. You might see recent uploads, a glimpse of groups, a piece of places, or something new you wouldn't have stumbled across otherwise. There is a lot happening on Flickr every day and we hope this lets you see a little more of it.

Note: If you don't want to see this Explore content, you can minimize the module at anytime by clicking on the arrow just to the left of the first photo.

[Permalink](#) | [Top](#)

How do I list my meetup on Flickr Blog?

Currently, we use [Upcoming](#) as our event platform, and to add your event to Flickr Blog, simply follow these steps:

If it's your first visit to [Upcoming](#), you will need to create an Upcoming account. To create one, do the following: Simply click on the Sign In link above the Upcoming logo. Then use your Yahoo!, Facebook, or Gmail credentials (can be similar to the ones you use for Flickr) to sign into the Yahoo! Network.

Once you are signed in, visit the [official Flickr group on Upcoming](#) and click the **Join this group** button to become a member.

Click on **+ ADD AN EVENT** in the navigation bar to start filling in your event's details, most importantly:

- **Event name** - This is what will show up on Flickr Blog, so try to make it as informative as possible.
- **Venue** - You can choose from a multitude of already available venues or add your own. Just start typing to give it a try.
- **Date** - This is important, right? You can also choose to add an end date or make it an all day long event.
- **Description** - Yes, please. This will help other community members understand what your event is all about, and whether it's a photowalk through the forest where long pants are highly recommended, or a picnic at the beach where it's best to bring some

bathing togs and a waterproof camera.

If you already have a group discussion running that you would like to link to, it's handy to add an **Event URL** in the **More Options** section.

When you are happy with all the event details, hit the **Preview Event** button. Make sure all the information is accurate and then **Add the Event**. Your event is now live and you will see it's **Upcoming page**.

To add it to Flickr Blog, look for the **Related Groups** section in the right hand column and click the **Send to group?** link. Select "Flickr" from the drop down.

Once listed on the Flickr group on Upcoming, your meetup will automatically be added to the *Community Events* on Flickr Blog. Yay!

[Permalink](#) | [Top](#)

Jump to another FAQ category ▼

Or, [return to the main help page](#).

About Flickr

- [Who we are](#)
- [Take the tour](#)
- [Flickr blog](#)
- [Jobs](#)

Community

- [Community Guidelines](#)
- [Report abuse](#)

Help

- [Need help? Start here!](#)
- [Help forum](#)
- [FAQs](#)
- [About Our Ads](#)

Apps and the API

- [Flickr for mobile](#)
- [App Garden](#)
- [API documentation](#)
- [Developer blog](#)
- [Developer Guide](#)

- [Follow us](#)
- [Like us](#)

[Deutsch](#) | [English](#) | [Español](#) | [Français](#) | [Italiano](#) | [Português](#) | [Tiếng Việt](#) | [Bahasa Indonesia](#)

[Terms of Use](#) | [Your privacy](#) | [Yahoo! Safety](#) | [Copyright/IP Policy](#)
Copyright © 2011 Yahoo! Inc. All rights reserved.



You aren't signed in [Sign In](#) [Help](#)

[Home](#) [The Tour](#) [Sign Up](#) [Explore](#) [Upload](#)

 [Search](#)

Holy smokes! That's cheap!

Just **\$24.95** for a 1 year pro account.

Here's what you get...

- **Unlimited** uploads and storage
- **Unlimited** sets and collections
- **Access** to your original files
- **Stats** on your account
- **Ad-free** browsing and sharing
- **HD playback** for high-definition video uploads

That's about **\$2** a month!

Buy Now!

You can get a **2 year** account for just **\$47.99** and that's even less per month!

Just so you know...

Flickr will always offer a [free account](#), and all prices on the site are quoted in US dollars.

About Flickr
[Who we are](#)
[Take the tour](#)
[Flickr blog](#)
[Jobs](#)

Community
[Community Guidelines](#)
[Report abuse](#)

Help
[Need help? Start here!](#)
[Help forum](#)
[FAQs](#)
[About Our Ads](#)

Apps and the API
[Flickr for mobile](#)
[App Garden](#)
[API documentation](#)
[Developer blog](#)
[Developer Guide](#)

[Follow us](#)
[Like us](#)

[Deutsch](#) | [English](#) | [Español](#) | [Français](#) | [Italiano](#) | [Português](#) | [Tiếng Việt](#) | [Bahasa Indonesia](#)

[Terms of Use](#) | [Your privacy](#) | [Yahoo! Safety](#) | [Copyright/IP Policy](#)
 Copyright © 2011 Yahoo! Inc. All rights reserved.

The New York Times

Bits

Business ■ Innovation ■ Technology ■ Society

JUNE 7, 2011, 7:29 AM

Apple iCloud May Not Be a Threat to Online-Storage Services

By [VERNE G. KOPYTOFF](#)

10:25 a.m. | **Updated** to add a comment from [Drew Houston](#), the chief executive of [Dropbox](#), an online storage service.

A herd of start-ups officially learned they had a new rival when [Apple introduced its iCloud](#) online storage service on Monday.

People will be able to store documents, presentations, photos, digital books and apps in a Web-based locker and get access to them from their Internet-connected devices, Apple said.

Marcio Jose Sanchez/Associated Press
Steven P. Jobs did not discuss whether people would be able to store all kinds of file formats in iCloud during his presentation on Monday.

It is a niche that has plenty of competition, as I discussed in an [article in Monday's paper](#). The rival companies are all trying to help people get into documents saved on their computers, smartphones and tablets without having to e-mail files to themselves.

The question is whether Apple, with its huge base of customers and deep pockets, will squash the companies already in the niche, like [Dropbox](#), [Box.net](#) and [Cx.com](#). The technology giants [Microsoft](#) and [Amazon](#) also have similar online storage services.

Maybe Apple will kill them all. But an early look at iCloud, which is supposed to be available sometime in the fall, makes this unlikely.

[Steven P. Jobs](#), Apple's chief executive, showed off iCloud on Monday at his company's annual developer conference in San Francisco, by talking about how it automatically stores files. He also spoke about how the service makes updated versions of those files available on any device a user owns, a process known as synchronization, or syncing for short.

The audience at the conference – largely Apple fanatics and developers who build on top of its products – applauded Mr. Jobs's presentation. But left untouched was the issue of whether people would be able to store all kinds of file formats in iCloud, and not just Apple's.

Mr. Jobs demonstrated the service using Apple's word processor, Pages, along with Numbers and Presentations. Apple's calendar, contacts and e-mail are also compatible.

There was no mention of Microsoft Office software like Word and Excel or Adobe Acrobat. If it doesn't work with those applications, iCloud would have limited utility for

There was no mention of Microsoft Office software like Word and Excel or Adobe Acrobat. If it doesn't work with those applications, iCloud would have limited utility for many people.

Aaron Levie, chief executive of Box.net, an online storage company, pointed out in [a blog post](#) on Monday after the announcement that the shortcoming would be particularly problematic for business customers, who often use a variety of file formats.

“The first issue with iCloud is that it will be optimized to work with other Apple products,” he wrote. “The de facto difficulty with Apple (speaking as a customer and amateur pundit) is that they are laser-focused on their own ecosystem.”

He continued: “This is fine in the consumer world, where we tend to have considerable flexibility in selecting our own software and hardware. But in the enterprise we're typically using devices, platforms, operating systems, and software that come from an array of vendors — and not always of our choosing.”

Apple said consumers would be able to store documents in iCloud that use Apple's iCloud Storage APIs, or technical specifications. Maybe other major companies will make their products available using those rules, but Apple only announced those specifications on Monday.

Apple also declined to disclose a price for iCloud other than to say that 5 gigabytes of storage would be free. Music, apps and books purchased from Apple will not count against the threshold, nor will photos in Apple's new Photo Stream service.

Competitors to iCloud typically provide a free tier of storage. A paid version with extra storage can cost up to \$20 a month. Dropbox, for example, offers 2 gigabytes of storage free of charge, but charges \$10 a month for 50 gigabytes and \$20 for 100 gigabytes. Amazon's Cloud Drive is free for up to 5 gigabytes and \$20 a year for up to 20 gigabytes, although it has other payment plans. Songs bought through Amazon do not count against the storage limit. Users get 20 gigabytes for free for a year if they buy an MP3 album.

iCloud does not appear to provide a way for users to share documents, at least not at this stage. Sharing files with colleagues or family is a main component of the rival storage services.

For example, users of Dropbox can choose to make a particular file available to others so they can edit the document together.

Drew Houston, chief executive of Dropbox, dismissed the idea that iCloud would upend companies like his. Although he did not mention it, Apple had a previous so-called “cloud” storage service, MobileMe, which shared many of iCloud's capabilities, but never really caught on.

iCloud, Mr. Houston said, is “a step forward for Apple, and the service will evolve, but when iCloud arrives later this year it won't let people access and share all of their files

iCloud, Mr. Houston said, is “a step forward for Apple, and the service will evolve, but when iCloud arrives later this year it won’t let people access and share all of their files everywhere the way tens of millions of people already do with Dropbox today.” He added, “Furthermore, there’s a big world beyond Apple, and our users love that Dropbox works just as well with your Android phone or PC as with your iPad or iPhone.”

[Copyright 2011 The New York Times Company](#) | [Privacy Policy](#) | [NYTimes.com](#) 620 Eighth Avenue New York, NY 10018



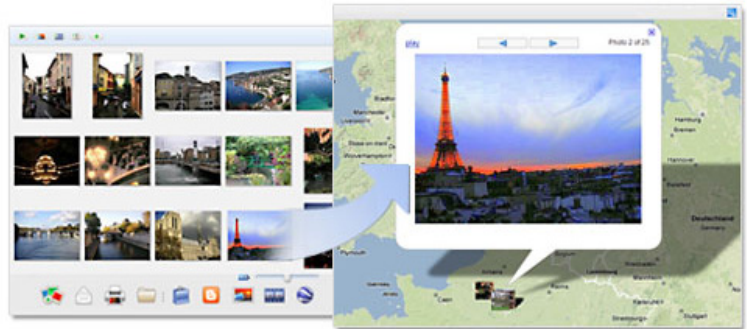
Picasa™ About Picasa 3.8 & Picasa Web Albums

[Help](#)

Fast and easy photo sharing from Google

Together, Picasa & Picasa Web Albums make it easy for you to organize and edit your digital photos, then create online albums to share with friends, family & the world.

Get Started with Picasa Web Albums»



Beautiful web albums


Show your photos at their best. View full-screen slideshows, see your pictures arranged on a global map, enjoy video playback, and more.



Focus on people

People matter in your photos. Our technology helps you automatically organize your photos based on the people in them, and works in Picasa and Picasa Web Albums.

[See name tags in Picasa Web Albums](#)
»



Get organized

Picasa automatically finds all the photos on your PC, wherever they are, and will organize them in seconds.



Be creative

Use Picasa to design and print beautiful photo collages, create fun video slideshows, add photo text or view your favorite photos on your desktop or screensaver.



It's free to use

Picasa is free to download, and Picasa Web Albums provides 1 gigabyte of free storage -- that's enough space for 4,000 wallpaper-size photos.

[Download Picasa 3.8](#)



Sharing made simple

Publish your favorite photos online with one click. Create stunning online photo albums to share with friends and family, or public albums for the world to see. Get notified when your "Favorites" post new photos.



Edit to perfection

Improve almost any picture with Picasa's one-click fixes for common problems like red-eye, color, and lighting. Or, use tuning and effects to make your best photos look even better.



Add places

Easily add geo-tags to your photos so that you can remember exactly where they're from using Google Maps.

[Learn more](#)



Order photos & gifts, or print at home

Picasa makes it easy for you to get the best out of your color printer. And Picasa Web Albums gives your friends the ability to download full-resolution pictures, so they can do the same. When ordering online, choose freely between major retail services.





Picasa and Picasa Web Albums

Search Picasa and Picasa Web Albums help

[Help home](#)[Storage and backups](#)

How it works

Storage and backups

How it works

[Free storage limits](#)[Download from album](#)[Exporting photos](#)[Saving photos](#)[Using an alternate drive](#)[Backup CD or DVD](#)[Rebuilding database](#)[How do I use the Picasa Web Albums Exporter with iPhoto?](#)[Recovering deleted photos](#)[Export to removable media](#)[Back up with Picasa Web Albums](#)[Download album](#)[Download size in Picasa Web Albums](#)[Original photo quality when saving](#)[Viewing edited pictures on the backup CD](#)

Google offers a way to purchase additional storage space shared across Gmail, Google Docs, and Picasa Web Albums (which includes photos uploaded to Blogger). You can [purchase additional storage](#) at any time.

Google Apps users: Google storage is available at the same prices and storage levels for [Google Apps](#) users. However, purchased storage will not apply to your Gmail account. [Learn more](#)

Free storage

Free storage space is specific to each product. Free storage from one product cannot be used by or transferred to another product.

- Gmail provides 7+ GB (and counting) for Gmail messages.
- Docs provides 1 GB for your uploaded files (documents created in Google Docs and converted files don't count towards your storage, but do have some [size limitations](#)).
- Picasa Web provides 1 GB for photos and videos. Files under certain sizes [don't count towards this limit](#).
- Google+ provides unlimited storage for photos, which are automatically resized to 2048 pixels. Videos up to 15 minutes in length are also free.

Paid storage

If you run out of free storage for Gmail, Google Docs, or Picasa Web, you can purchase additional storage that is shared across these products. Choose from the following storage plans:

- 20 GB - \$5/yr
- 80 GB - \$20/yr
- 200 GB - \$50/yr
- 400 GB - \$100/yr
- 1 TB - \$256/yr

Related

[Sign up](#)[Get Started](#) › [Picasa Web Albums](#)[Add name tags in Picasa Organize](#) › [Text tags and searches](#)[Album visibility](#)[Share](#) › [Album visibility and access](#)[How do I select multiple photos in the Library?](#)[Organize](#) › [Folders and albums](#)[500 internal server error](#)[Fix a Problem](#) › [Errors](#)

- 2 TB - \$512/yr
- 4 TB - \$1024/yr
- 8 TB - \$2048/yr
- 16 TB - \$4096/yr

Note that 1 TB equals 1024 GB, and these prices don't include any applicable taxes or fees associated with your country of residence.

Google storage purchases are annual subscriptions that renew after one year. You can cancel your subscription at any time, but storage purchases are non-refundable. You can also upgrade to a larger storage plan for the pro-rated difference at any time. Paid storage for one Google Account cannot be transferred to a different account. [Learn more about Google storage refunds, renewals and cancellations](#)

It's normal to experience a delay of up to 24 hours before purchased storage is added to your account. [See our storage troubleshooting steps](#) for more information.

updated 10/04/2011

Picasa and Picasa Web Albums - Contacting Us - Help with other Google products - Change language: English (US) ▼

©2011 Google - Google Home - Privacy Policy - Terms of Service

Find Open Source Software

[\(/\)](#)

[Browse \(/directory/\)](#) [Blog \(/blog/\)](#) [Support \(/support\)](#) [Jobs \(http://jobs.sourceforge.net/\)](#) [Newsletters \(https://sourceforge.net/user/registration\)](#)
[Resources \(http://geek.net/events/\)](#)

[Register \(https://sourceforge.net/user/registration\)](#) [Log In \(https://sourceforge.net/account/login.php\)](#)

About

SourceForge (<http://sourceforge.net/>) is dedicated to making [open source \(http://opensource.org/\)](http://opensource.org/) projects successful.

We thrive on community collaboration to help us create the leading resource for open source software development and distribution. With [the tools we provide \(/create\)](#), 2.7 million developers create powerful software in over [260,000 projects \(/directory/\)](#). Our [popular directory \(/directory/\)](#) connects more than 46 million consumers with these open source projects and serves more than 2,000,000 downloads a day. SourceForge is where open source happens.

SourceForge.net is owned and operated by [Geeknet, Inc. \(http://geek.net/\)](http://geek.net/), a publicly traded US-based company.

Connect with us:

[@sourceforge on twitter \(http://twitter.com/sourceforge\)](http://twitter.com/sourceforge)

[#sourceforge on irc.freenode.net \(http://webchat.freenode.net/?randomnick=0&channels=sourceforge\)](http://webchat.freenode.net/?randomnick=0&channels=sourceforge)

[sfnet_ops@geek.net \(mailto:sfnet_ops@geek.net\)](mailto:sfnet_ops@geek.net)

[Status \(http://twitter.com/sfnet_ops\)](http://twitter.com/sfnet_ops) [Terms \(/apps/trac/sitelegal/wiki/Terms_of_Use\)](#) [Privacy \(http://geek.net/privacy-statement\)](http://geek.net/privacy-statement) [Advertise \(http://geek.net/advertise/\)](http://geek.net/advertise/) [About \(/about\)](#)

© 2011 [Geeknet, Inc. \(http://geek.net\)](http://geek.net/)

Already a member? Log in here. [Forgot password?](#)[Home](#)[Advertisers](#)[How it Works](#)[Blog](#)[Videos](#)[Help](#)

Earn money
when you deliver
quality traffic

Apply Now to the eBay Partner Network

Drive quality traffic to one of our partners and get paid for it.

Apply Now!

How it works

You can earn a lot of money in the eBay Partner Network, by driving high quality traffic to eBay or one of our partners.

How to generate revenue through your website

Advertise

Getting Paid

Optimize

[Register](#)[Tools](#)[Tracking](#)[Quality Click Pricing](#)[Payment](#)[Reports](#)[Optimize](#)

Blog

[ePN TV: Carrie Smith Walks Through the Daily Deals Holiday Calendar](#)

11/07/2011 07:34:44 AM

[See You at BlogWorld Los Angeles!](#)

11/01/2011 08:55:05 AM

[ePN TV: Julia Barrett Talks Holiday Promotions](#)

10/28/2011 09:05:28 AM

[View the blog](#)

Quality Click Pricing

Find out more about Quality Click Pricing.

Tools, Widgets & Links

Find the right tools for advertising on your site.

Developers Program

Find out how to tap into eBay's marketplace with our API.

Learn more about Quality Click Pricing

[How do I improve my traffic's quality?](#)

[FAQs](#)

Copyright © 2011 eBay Inc. All Rights Reserved. Designated trademarks and brands are the property of their respective owners.

Use of this website constitutes acceptance of the [eBay Partner Network Agreement](#) and [eBay Partner Network Privacy Policy](#)

Language: [English](#) [Français](#) [Deutsch](#) [Español](#) [Italiano](#)



What happened to gatekeeper.dec.com?

- » **HP Labs**
- » Research
- » Technical reports
- » Downloads

History

The FTP archive on gatekeeper.dec.com, later to become gatekeeper.research.compaq.com, first went on-line in the late 1980's. It was assembled from spare parts and brought to life by Paul Vixie. Located at [Digital Equipment Corporation's](#) Western Research Laboratory (DECWRL) in Palo Alto, California, it was maintained by a handful of volunteers from Digital Corporate Research. The spare parts were eventually replaced with production hardware.

In its heyday, gatekeeper was a prominent Internet FTP site. Just about any public domain software package you wanted or needed could be found on gatekeeper. One of gatekeepers' primary functions was to give Digital software developers, living on the DECNET based internal corporate network, access to public domain software available from the Internet. (The Digital internal DECNET host, DECPA:, mounted the gatekeeper FTP archive via NFS.) Gatekeeper was also used by Digital product groups to provide software updates and patches to Internet customers.

Today

Over the years, acquisitions, mergers, and business needs have changed the role of gatekeeper. Early in 2008, gatekeeper.dec.com and gatekeeper.research.compaq.com were redirected to apotheca.hpl.hp.com, a shiny new server at HP Labs in Palo Alto, California. ftp.hpl.hp.com is also directed to apotheca.

Apotheca doesn't mirror any of the public domain software repositories that were mirrored on gatekeeper. Nowadays, there's plenty of Internet sites that do a fine job of serving those bits. This is why you'll find ~ftp/pub rather sparse, compared to what you use to find on the old gatekeeper.

The legacy content from the Digital and Compaq research labs is available in [~ftp/gatekeeper/](#).

Richard Schedler

Gatekeeper Archive Administrator 1993-2007

[Privacy statement](#)

[Using this site means you accept its terms](#)

[Feedback to HP Labs](#)

© 2007 Hewlett-Packard Development Company, L.P.

[Online Video News](#)

- [Events](#)
- [GigaOM.tv](#)
- [Research](#)


- [Home](#)
- [Apple](#)
- [Broadband](#)
- [Cleantech](#)
- [Cloud](#)
- [Collaboration](#)
- [Mobile](#)
- [Video](#)

[Video](#)

Fingerprint Tech Waits for a Grand Entrance

By [Liz Gannes](#) Oct. 9, 2007, 12:19pm PT [2 Comments](#)

-  Tweet  0
-
-
-

It turns out the forces of change can't be held back by crappy anti-piracy technology after all. [Unfulfilled promises](#) by Google (GOOG) to launch video fingerprinting technology on YouTube in order to prevent uploads of copyrighted content have had little effect on the rest of the industry. In the meantime, big media is adopting the tools of the trade to put more and more content online. And along the way, it's learning that distributing programming encourages fans — while closing up programming encourages piracy.

But down the road, there's still an opportunity for video fingerprinting technology to make content owners genuinely comfortable with online distribution. While [we've found](#) early implementations of audio fingerprinting have problems detecting piracy, the MPAA claims that recent tests have been quite promising.

We've seen a couple of developments on that front just this week: Santa Clara, Calif.-based [Vobile](#), maker of the "VideoDNA" video fingerprinting and tracking product, [said](#) it had scored Gideon Yu — who is known for being chief financial officer of YouTube immediately prior to its sale to Google, and is currently CFO of Facebook — as a board member. Elsewhere, Paris-based Dailymotion said it was employing Ina, a

digital image bank, to detect copyrighted videos on the French version of its site.

In a recent test of video fingerprinting technologies, multiple vendors were effective at identifying copyrighted files they had been trained to recognize, [according to the Los Angeles Times](#).

The MPAA and MovieLabs, its R&D arm, spent months testing a dozen software programs designed to identify copyrighted videos from brief samples of their data. Each “fingerprinting” program was fed about 1,000 test files of wide-ranging format and quality, including a number of camcorder movies and other bootlegs downloading from the Net. Of the technologies submitted (11 by firms and one by a Scottish university), the MPAA said, three successfully identified more than 90 percent of the files with no false positives.

Reportedly, though the MPAA and MovieLabs wouldn’t disclose it, Vobile was the overall winner of those tests.

Video sites such as Microsoft’s (MSFT) Soapbox and MySpace already employ fingerprinting technology from provider Audible Magic, and Veoh has said it will implement fingerprinting this month. A Dailymotion spokesperson clarified that the site will continue to use Audible Magic for audio fingerprinting, but is adding in Ina for video fingerprinting. He emphasized that the initial Ina rollout on Dailymotion is only in France, though the company expects it to be implemented across all versions of its sites by 2008.

Other fingerprinting providers we’ve covered include [Activated Content](#), [Autonomy](#), [MotionDSP](#), [iPharro](#), and Gracenote.

Related research and analysis from GigaOM Pro:

Subscriber content. [Sign up for a free trial](#).

- [Why iPad 2 Will Lead Consumers Into the Post-PC Era](#)
- [The Near-Term Evolution of Social Commerce](#)
- [Content Farms: The Players, The Benefits, The Risks](#)

If you like this story, please share it

-  Tweet 0
-
-
-



- [Comments](#)
2

Solutions

Overview

▶ Application Integration

▶ Process Automation

▶ Content Migration

▼ Web Intelligence

Web Data Extraction

Social Media Monitoring

News Syndication

Online Account Aggregation

Background Checks

[AntiPiracy](#)

OSINT

▶ Mobile Enablement

▶ Analytics & Business Intelligence

Web Intelligence Solution: Antipiracy

[Home](#) > [Solutions](#) > [Web Intelligence](#) > [AntiPiracy](#)

Fight media piracy and protect against lost revenue with automated web intelligence

Media piracy on the web accounts for hundreds of billions of dollars in lost revenue every year for media companies and artists. The ability to automatically monitor and collect evidence from pirated downloads and BitTorrent sites has become a critical weapon in fighting this problem.

Featured Customer Success Stories in AntiPiracy:



In the fight against software and movie piracy, Kapow delivers a major advance in intellectual property and revenue protection.

▶ Close ▶ Det

Challenge

- Identify instances of copyright infringement, where the company's products have been illegally copied and shared online.
- Implement a cost effective mechanism for searching thousands of sites daily for infringing content.
- Address day-to-day requirements as well as peak demand surrounding major product releases.

Solution

- Kapow solution searches daily for illegally copied materials from the company's back catalog.
- Scales to meet peak demand surrounding new product releases using flexible scheduling rules and automatic load-balancing.

- Provides a central resource for the company's global anti-piracy initiative.

Results

- Enables the company to identify and remove infringing content more rapidly than ever before.
- Provides a centralized anti-piracy infrastructure available to all departments.
- Enables targeting of specific new product releases for the first time, resulting in the detection of millions of copyright infringements for a single product
- Reduces labor costs by automating web intelligence and aggregation, enabling staff to be redeployed.



Entertainment

Kapow is a powerful weapon in the fight against web piracy.

▶ Close ▶ Deti

Challenge

- Search the web for illegally copied material.
- Extract structured data from thousands of sites to find copyright infringements.
- Gather evidence for use in copyright infringement prosecution cases.

Solution

- Using current artist/title catalog , the agency uses Kapow robots to search hundreds of forums, blogs and peer-to-peer websites every day to find illegally posted content.
- Copyright infringements are reported to the copyright owner and takedown notices are automatically issued to the hosting website.

Results

- Enabled hugely scalable copyright infringement detection system.
- Allowed redeployment of content analysts, resulting in higher morale due to more rewarding work.
- Reduced average detection time to minutes/hours rather than days (representing potential cost savings for copyright owners).

- Enabled successful prosecution of serial copyright infringing websites.

Top Links

Kapow Software New

FEATURED SOLUTIONS

- Application Integration
- Cloud / SaaS Integration
- Process Automation
- API Enabling
- Analytics and Business Intelligence
- Competitive Price Monitoring
- Mobile Enablement
- Web Intelligence
- Social Media Monitoring
- Content Migration
- View all Solutions ▶

DISCOVER KAPOW SOFTWARE

- Overview
- Data Delivery Challenge
- Kapow Software Revolutionary Approach
- Delivery Options
- Resource Center

PRODUCTS

- Kapow Katalyst Platform
- Kapow Mobile Katalyst
- Request a Trial

CUSTOMERS AND PARTNERS

- Customer Successes
- Partner Directory
- Kapow Partner Program

SERVICES

- Overview
- Consulting Services
- Education Services
- Technical Support Services

Kapow Software Intro Katalyst 8.2

Kapow Software Rep With 64 Percent Growth

Kapow Software Named Vice President of Marketing

Kapow Software Achieves Validated Integration with WebCenter

More News ▶

1-800-805-0828



Request Sales Call

© 2010 Kapow Software Terms of Use

You are here: LAT Home > Blogs > Bit Player

BIT PLAYER

HOLLYWOOD'S LOVE-HATE RELATIONSHIP WITH TECHNOLOGY

« Online media round-up | Main | Content recognition, part 2 »

The content-recognition bakeoff



There was one obvious winner in the MPAA's **test of content-recognition technologies**, and that would be Hollywood. The association disclosed a brief summary of the results at a **conference** on online entertainment held Thursday in cooperation with the University of California. Paradoxically, though, the results might prove to be a boon for user-generated content sites and the people who love them.

Here's the background. The MPAA and **MovieLabs**, its R&D arm, spent months testing a dozen software programs designed to identify copyrighted videos from brief samples of their data. Each "fingerprinting" program was fed about 1,000 test files of wide-ranging format and quality, including a number of camcorder movies and other bootlegs downloading from the Net. Of the technologies submitted (11 by firms and one by a Scottish university), the MPAA said, three successfully identified more than 90 percent of the files with no false positives. In other words, when they erred, it was in not spotting a copyrighted movie within a file, rather than identifying a movie that wasn't there.

Steve Weinstein, ceo of MovieLabs, summarized the results as follows: "This stuff works." Of course, that's just in a laboratory setting. The next step for MovieLabs and the MPAA, in addition to testing more iterations of the technology, is to try out the technologies in the field, where commercially important issues such as scaling (how well they handle a large volume of ID requests) can be measured.

Weinstein and other officials declined to say which firms had scored highest in the test. That honor was claimed by Santa Clara-based startup **Vobile**, although executives from two more established competitors, Gracenote and Audible Magic, also said their technology performed well. For the studios, though, the important thing is that several firms actually seem to be able to do what they claim. That could raise the pressure on networks with a large amount of unauthorized copying, such as user-generated video sites and college campuses, to deploy fingerprinting technologies. In fact, Joshua Metzger, a senior vice president at the UGC site Veoh.com, said at the conference that his company would add fingerprinting technology within 30 days.

A lot of details remain unsettled, and the courts have **yet to weigh in** on whether copyright law actually obliges anyone to install fingerprinting technologies on their system. But a fingerprinting technology that works well enough to satisfy the studios could be the key to advertiser-supported sites such as Veoh and YouTube striking major content deals and ramping up their ad sales. After all, one of the things holding back revenues at UGC sites is uncertainty over liability. If fingerprinting settles that issue, that's a powerful motivator for companies to deploy it. On the other hand, if it chokes off so much content that users run away, then it's a lose-lose situation.

Posted by Jon Healey on September 21, 2007 | Permalink

Comments

[View the entire comment thread.](#)

Post a comment

If you are under 13 years of age you may read this message board, but you may not participate.

Here are the full legal terms you agree to by using this comment form.

Comments are moderated, and will not appear until they've been approved.

This weblog only allows comments from registered users. To comment, please **Sign In**.

News/Opinion

- California | Local
- National
- World
- Business
- Sports
- Washington
- Science
- Environment
- Opinion

Arts/Entertainment

- Entertainment
- The Guide
- Company Town
- Arts & Culture
- Calendar
- The Envelope
- TV Listings
- Movie Showtimes

Living

- Travel
- Health
- Autos
- Home & Garden
- Food
- Image
- Books
- Brand X

Data Desk

- Video
- Photography
- Obituaries
- Crosswords/Sudoku
- Your Scene

Blogs

- Columnists
- Print Edition
- Readers Rep
- Corrections
- All Sections

Buy, Sell & More

- Jobs
- Cars
- Real Estate
- Foreclosure Sale
- Rentals
- Personals
- Local Values
- Coupons
- Newspaper Ads

Place an Ad

- In the Newspaper
- Online

Settings/Services

- Sign In

Our Blogger

Times editorial writer **Jon Healey** pens opinion pieces about a variety of business issues, and blogs about technologies that are changing the entertainment industry's business model.

Recent Comments

- **Aryol Prefabrik** on **Nero burns a new mission**
- **Mathew** on **Consumer groups blast MPAA proposal (a dog bites man story)**
- **Roy Anderson** on **Should local radio pay royalties?**
- **Jon Healey** on **Gizmodo on journalism**
- **Socrates** on **Gizmodo on journalism**
- **Chris K.** on **A new day for WebTV?**

Search this blog

- Technorati search
- all blogs this blog
-

Archives

- August 2008
- July 2008
- June 2008
- May 2008
- April 2008
- March 2008
- February 2008
- January 2008
- December 2007
- November 2007
- October 2007
- September 2007
- August 2007
- July 2007
- June 2007
- May 2007
- April 2007
- March 2007
- February 2007
- January 2007
- December 2006
- November 2006
- October 2006
- September 2006
- August 2006
- July 2006
- June 2006

Subscribe to this Blog - What is RSS?

Recent Posts

- **New RSS feed**
- **Moving to a new blog**
- **The Cablevision DVR ruling**
- **More on the FCC, Comcast and BitTorrent**
- **The FCC's Comcast decision**

Now Playing

- **Brazilian Girls**
Mano De Dios

Categories

- Apple
- Audio
- Business models
- Computers
- and other gadgets
- Copyrights
- DRM
- Games
- Geek speak
- Mobile
- Music**
- Net

[Reprint Requests](#)
[Work for Us](#)

Home Delivery

[Customer Support](#)
[Subscribe](#)

[neutrality](#) [Odds and ends](#) [Regulation](#)
Software TV, film and video
[Video](#)

Where I've Been Lately

- **Web Scout**
Web content highlights
 - **TechCrunch**
Well reported tech-news site
 - **Goodyblog**
A smart take on TV and pop culture
 - **PaidContent**
Nice roundup of news and conferences
 - **Good Morning Silicon Valley**
The view from San Jose, snarkily
 - **Techmeme**
A window into the tech blogosphere
 - **Engadget**
Allgadget
 - **Coolfer**
On music and the music biz
 - **Freedom to Tinker**
Professor Ed Felten's musings about devices and DRM
 - **Progress and Freedom Foundation IP blog**
Where regulation is the problem, not the solution
-

More on LATimes.com

[California/Local](#) | [National](#) | [World](#) | [Sports](#)
[Opinion](#) | [Entertainment](#) | [Travel](#) | [Health](#)
[L.A. Wheels](#) | [Real Estate](#)

Partners

[Hoy](#) | [KTLA](#) | [Boodle.com](#) | [ShopLocal](#)
[Shopping](#) | [Grocery Coupons](#)

Classifieds

[Career Builder](#) | [Cars.com](#) | [Apartments.com](#)
[FSBO \(For Sale By Owner\)](#) | [Open Houses](#)

Los Angeles Times

Copyright 2011 Los Angeles Times

[Privacy Policy](#) | [Terms of Service](#) | [Advertise](#) | [Home Delivery](#) | [Reprint Requests](#) | [Help & Services](#) | [Contact](#) | [Site Map](#)

BIT PLAYER

HOLLYWOOD'S LOVE-HATE RELATIONSHIP WITH TECHNOLOGY

« [The content-recognition bakeoff](#) | [Main](#) | [Art Tatum live, kinda sorta](#) »

Content recognition, part 2



As I mentioned in [my last post](#), the MPAA and MovieLabs' **tests** of 12 content-recognition technologies found several that performed quite well in the lab. According to information released at a joint MPAA-University of California **event** last week, seven of the technologies correctly identified at least 80% of the test files, and three got better than 90% right with no false positives. But several speakers at the event predicted that the technologies would prove much more useful for distributing content than for blocking it.

The event was designed to explore how colleges and the entertainment industry could work together to address piracy. Again and again, speakers opined that making content easier to get legitimately, rather than trying to make it harder to obtain illegitimately, was the best way to combat illegal downloading. Fingerprinting can be used for either purpose, but it's most effective when used to track usage online and determine what royalties are owed. For example, panelists from AT&T Labs and USC talked about the prohibitive expense and/or intrusiveness of examining every bit of traffic flowing across large networks, which is what a network operator would have to do if it was using fingerprints to deter unauthorized file-sharing. "You don't look inside what people are doing unless there's a legitimate reason to do so," said Marty Loman, vice president of content protection at AT&T Labs. "You have to narrow down who you're going to look at."

Fingerprinting and DRM techniques will drive students away if they make it harder to consume what they want on the devices they want to use, several speakers warned. Ashwin Navin, president of [BitTorrent Inc.](#), put it this way: "If demand exists for content in a certain model, it will be fulfilled. The question for rights holders is, do we want that demand fulfilled legally or do we want that fulfilled illegitimately?" Added UCLA student Colin Iberti, "To put things on your iPod is very important to the college demographic. Not any other MP3 player, but iPods specifically." The authorized content that colleges offer typically uses a DRM not compatible with iPods, and "it really alienates the 98% of us who want to put it on their iPods."

In sum, the MPAA and MovieLabs' comments about fingerprinting technology had to be encouraging to copyright owners that are looking to boost online revenue. But if the speakers at the Universal City Hilton were right, fingerprinting won't be a silver bullet for copyright holders eager to kill the online bootlegging beast.

Posted by Jon Healey on September 24, 2007 in [DRM](#), [Music](#), [TV, film and video](#) | [Permalink](#)

Comments

[View the entire comment thread.](#)

Post a comment

If you are under 13 years of age you may read this message board, but you may not participate.

News/Opinion

[California](#) | [Local](#)
[National](#)
[World](#)
[Business](#)
[Sports](#)
[Washington](#)
[Science](#)
[Environment](#)
[Opinion](#)

Arts/Entertainment


[Entertainment](#)
[The Guide](#)
[Company Town](#)
[Arts & Culture](#)
[Calendar](#)
[The Envelope](#)
[TV Listings](#)
[Movie Showtimes](#)

- [Real Estate](#)
- [Foreclosure Sale](#)
- [Rentals](#)
- [Personals](#)
- [Local Values](#)
- [Coupons](#)
- [Newspaper Ads](#)

Place an Ad

- [In the Newspaper](#)
- [Online](#)

Settings/Services

- [Sign In](#)
- [Register](#)
- [Personalized News](#)
- [E-Mail Newsletters](#)
-  [RSS Feeds](#)
- [Help](#)
- [Contact Us](#)
- [L.A. Times Archives](#)
- [Reprint Requests](#)
- [Work for Us](#)

Home Delivery

- [Customer Support](#)
- [Subscribe](#)

Here are the full legal terms you agree to by using this comment form.

Comments are moderated, and will not appear until they've been approved.

This weblog only allows comments from registered users. To comment, please **Sign In**.

Our Bl

Times e
about a
technol
industry

Recen

- [Aryol](#)
- [Mathe](#)
- [propos](#)
- [Roy A](#)
- [royaltie](#)
- [Jon H](#)
- [Socra](#)
- [Chris](#)

Search

•

Archiv

- [August](#)
- [July 20](#)
- [June 20](#)
- [May 20](#)
- [April 20](#)
- [March ;](#)
- [Februa](#)
- [Januar](#)
- [Decemi](#)
- [Novem](#)
- [Octobe](#)
- [Septem](#)
- [August](#)
- [July 20](#)
- [June 20](#)
- [May 20](#)
- [April 20](#)
- [March ;](#)
- [Februa](#)
- [Januar](#)
- [Decemi](#)
- [Novem](#)
- [Octobe](#)
- [Septem](#)
- [August](#)
- [July 20](#)
- [June 20](#)

Subsc

Recen

Now P

Categ

Where

-
-
-
-
-
-
-
-
-
-

More on LATimes.com

California/Local | National | World | Sports

Partners

Hoy | KTLA | Boodle.com | ShopLocal

Classifieds

Career Buil

[Opinion](#) | [Entertainment](#) | [Travel](#) | [Health](#)
[L.A. Wheels](#) | [Real Estate](#)


[Shopping](#) | [Grocery Coupons](#)

[FSBO \(For](#)

Los Angeles Times

Copyright 2011 Los Angeles Times

[Privacy Policy](#) | [Terms of Service](#) | [Advertise](#) | [Home Delivery](#) | [Reprir](#)



Refer a Friend Returns!

Subject: Refer a Friend Returns!
From: Ooma <newsletter@ooma.com>
Date: Thu, 20 Oct 2011 07:09:36 -0400
To: Andrew Cromarty <and[REDACTED]@[REDACTED].com>



Refer a Friend Returns!

Dear Andrew,

Ooma brings you huge savings over your old phone company, while offering all of the features you love plus more. And now with our Refer a Friend program, you can earn rewards for sharing your secret to savings with your loved ones.



Your friends and family will get a great price of \$199 for a new Ooma Telo, plus free shipping. Even better, when someone uses your referral code to purchase a new system, you will receive a \$25 Amazon gift card. This means you can earn up to \$75 in total!

Here's how to get started:

1. Log in to the [Refer a Friend](#) page in My Ooma to generate referral codes you can share with your friends and family. Everyone has up to three one-time use referral codes to send out.
2. When the referral link is clicked, your friends and family will be taken to a page where they can learn more about Ooma and purchase directly on the site.
3. You'll earn a \$25 Amazon gift card for each purchase completed with your referral codes. Check back in My Ooma to track how many gift cards you've earned!

If you're ready to share the savings, then act now because this offer ends on November 15th. Start [referring your friends](#) today!

Until later,
Team Ooma

This referral program expires on 11/15/2011. Referral codes must be redeemed by this date to be eligible for reward. Purchases must be made using the referral code and link provided to be eligible for promotional offer and referral reward. Ooma will not be held responsible for any lost or stolen gift cards. This offer cannot be combined with any other program. Ooma reserves the right to cancel or modify the terms of this program at any time.

Refer a Friend Returns!

New Voicemail-to-Text

Read your voicemail instead
of listening to it

[LEARN MORE >](#)

New International Bundle

Call 70 countries for less
than 1c per minute

[CHECK IT OUT >](#)

Step up to Premier 

Fire up the full power of
Ooma with Premier!

[SUBSCRIBE NOW >](#)

This email message was sent from:
Ooma Inc, 1840 Embarcadero Road, Palo Alto, CA
Need assistance? Visit www.ooma.com/support

You received this message because you are currently subscribed to receive news from Ooma.
[Click here](#) to unsubscribe from future tips and newsletter e-mails from Ooma.

Network.http.sendRefererHeader

The title given to this article is incorrect due to [technical limitations](#). The correct title is **network.http.sendRefererHeader**.

This article describes the preference **network.http.sendRefererHeader**. To add, delete, or modify this preference, you will need to [edit your configuration](#) — do *not* edit this article.

Contents [\[hide\]](#)

- [1 Background](#)
- [2 Possible values and their effects](#)
 - [2.1 0](#)
 - [2.2 1](#)
 - [2.3 2](#)
- [3 Caveats](#)
- [4 Recommended settings](#)
- [5 First checked in](#)
- [6 Has an effect in](#)
- [7 Related bugs](#)
- [8 Related preferences](#)

Background

[\[edit\]](#)

[HTTP](#) is the application-layer protocol with which most web pages are transferred. As part of HTTP, requests can include a ["Referer" \(sic\) header](#) that tells the server which page the user was on that initiated the request. Servers use this information to track users' paths through the site and possibly provide additional features.

Additionally, in JavaScript, the current page's referrer is exposed in the DOM through [document.referrer](#). Scripts running on the page can consult this property to see the same information that was sent in the Referer header.

This preference controls when to send the Referer header and set `document.referrer`.

[\[edit\]](#)

Possible values and their effects

0

Never send the Referer header or set `document.referrer`.

[\[edit\]](#)

1

[\[edit\]](#)

Main Page

This Page

- [Article](#)
- [Discussion](#)
- [Edit](#)
- [History](#)
- [What links here](#)
- [Related changes](#)
- [Printable version](#)

Your Account

- [Log in/Create Account](#)

Tools

- [Recent Changes](#)
- [Upload file](#)
- [Special pages](#)

Search

[Old Search](#)

[submit news](#)

- [forums](#)
- [weblogs \(feedHouse\)](#)
- [knowledge base](#)

- [chat](#)
- [members](#)
- [store](#)
- [about](#)

[fr](#) [ja](#) [de](#) [ko](#) [es](#) [hu](#)

Send the Referer header when clicking on a link, and set `document.referrer` for the following page.

2

[\[edit\]](#)

Send the Referer header when clicking on a link or loading an image, and set `document.referrer` for the following page. (Default)

Caveats

[\[edit\]](#)

- Disabling Referer headers may cause some functionality on some sites to no longer work.

Recommended settings

[\[edit\]](#)

Those concerned with privacy can set this to 0, realizing that this may adversely affect some sites. Those wanting to ensure compatibility should leave it at the default.

First checked in

[\[edit\]](#)

[2001-05-11 by Darin Fisher](#)

Has an effect in

[\[edit\]](#)

- Netscape (all versions since 6.1)
- Mozilla Suite (all versions since 0.9)
- Mozilla Phoenix (all versions)
- Mozilla Firebird (all versions)
- Mozilla Firefox (all versions)
- SeaMonkey (all versions)
- Camino (all versions)
- Minimo (all versions)

Related bugs

[\[edit\]](#)

- [Bug 1582 - \[NECKO\]\[webshell\] Send HTTP Referer field to server, "network.sendRefererHeader"](#)
- [Bug 76866 - http spews many "private" events before any real data events](#)

Related preferences

[\[edit\]](#)

- [browser.send_pings](#)
- [network.http.sendSecureXSiteReferrer](#)

Categories: [Preferences](#) | [Security and privacy-related preferences](#)

This page was last modified 21:30, 10 June 2007. This page has been accessed 143,077 times. [About MozillaZine Knowledge Base](#) - [Disclaimers](#)

MozillaZine and the MozillaZine Logo Copyright © 1998-2007 MozillaZine. All Rights Reserved - [Privacy Policy](#)

Network Working Group
Request for Comments: 2616
Obsoletes: 2068
Category: Standards Track

R. Fielding
UC Irvine
J. Gettys
Compaq/W3C
J. C. Mogul
Compaq
H. Frystyk
W3C/MIT
L. Masinter
Xerox
P. Leach
Microsoft
T. Berners-Lee
W3C/MIT
June, 1999

Hypertext Transfer Protocol -- HTTP/1.1

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the “Internet Official Protocol Standards” (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice


Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

The Hypertext Transfer Protocol (HTTP) is an application-level protocol for distributed, collaborative, hypermedia information systems. It is a generic, stateless, protocol which can be used for many tasks beyond its use for hypertext, such as name servers and distributed object management systems, through extension of its request methods, error codes and headers [47]. A feature of HTTP is the typing and negotiation of data representation, allowing systems to be built independently of the data being transferred.

HTTP has been in use by the World-Wide Web global information initiative since 1990. This specification defines the protocol referred to as “HTTP/1.1”, and is an update to RFC 2068 [33].

Exhibit Excerpted
(112 pages removed)



51 Reasons to Invest in the National Information Infrastructure

Early in 1993, the Federation of American Research Networks teamed with Interop, Inc. and the Coalition for Networked Information to collect, organize, and disseminate stories which describe the transformational potential of high performance networks such as the Internet. The goal was to document the positive affects the Internet was having upon the provision of services in education, research, manufacturing or other areas. Fifty-one of these so-called "network success stories" were published as a booklet and sent to all the members of the current U.S. Congress as well as to the Clinton Administration. All of the contributed submissions are compiled here.

[Call for Submission of Stories \(March, 1993\)](#)

[Form for Submission of Stories](#)


[Geographical Index of Stories](#)

[Categorical Index of Stories](#)

 info@farnet.org



FARNET: 51 Network Success Stories



FARNET Stories Project

51 Reasons to Invest in the National Information Infrastructure

story149.NM

Submitted by:

Marianne Granoff
Director of Operations
New Mexico Technet
4100 Osuna NE, Suite 103
Albuquerque, NM 87109 USA

v: (505) 345-6555

f: (505) 345-6559

e: granoff@technet.nm.org

Categories:

Education, higher; Education, continuing or distance ; Research, academic; Research, government; Research, commercial; Library; Other

Keywords:

Innovative or improved ways of doing things; More equitable access to technology or electronic information; Creation of new ideas, products, or services; Technology transfer; Local commitment to network-based activities; Leverage of public funding; Volunteer contributions of time and energy; Partnerships between public and private sector

Supporting Documentation (contact author for more information):

Software; Documentation; CD-ROM; Other

Story Site (if other than location listed above):

White Sands Missile Range
New Mexico

The Story:

THE STORY OF SIMTEL20

SIMTEL20 is well known among the Internet community as the world's largest on-line repository of freely accessible software and documentation on the Internet. What is not well known today, is how it came into being.

In 1979, CP/M (Control Program for Microcomputers) was the primary operating system on smaller (micro) computers. One summer evening back then, Frank Wancho, an employee of WSMR (White Sands Missile Range) in New Mexico, and at the time, the volunteer sponsor of the MIT (Massachusetts Institute of Technology) INFO-CPM mailing list, made a phone call from his home in El Paso, Texas to a newly published phone number for a computer bulletin board system outside Detroit, Michigan. The 2:00 am phone call was answered by the BBS operator, Keith Petersen, rather than by the computer, as expected.

Keith was, at the time, employed by the local Detroit CBS affiliate as a technical engineer, but was also an avid computer enthusiast. As a result of his computing interests, Keith had started collecting CP/M utilities, files, software, etc. from other enthusiasts and keeping them on his system. He was informally encouraging others to send him new software to add to his collection. The unusual part of this activity was that he provided a way for other users to SECURELY do this using the actual remote CP/M commands, by using multiple directories for uploading and downloading.

As a result of that phone call, Keith began to be an active contributor to the MIT INFO-CPM mailing list. This was at a time when the PDP-10 at MIT had no file transfer protocol that was compatible with CP/M. Keith used "blind uploads" via modem to send new public domain CP/M software, patches, and documentation to the MIT machine on a daily basis. The files themselves were sent directly to the MIT mailing list as messages (i.e. each person on the list received a separate copy of the file).

It soon became clear that all this traffic was straining the MIT mail delivery system. MIT agreed to set aside disk space to hold the CP/M collection, and only the announcements of the new software were sent to the mailing list. Several volunteers at MIT, notably Gail Zacharias, soon managed to write mainframe versions of the MODEM2 file transfer protocol and other CP/M utility programs for the MIT machine. These significantly improved the accuracy of the file contents that Keith uploaded.

In 1983, as availability and access to the MIT computer was removed, Frank Wancho (the System Administrator for the WSMR computer), arranged for the contents of the MIT CP/M

and the newly formed MS-DOS collections to be moved to a DECSYSTEM 20 that had excess CPU and Disk capacity at WSMR. This computer was also on the ARPANET. Because of the existing network requirements at that time, ARPANET hosts had to have 8 character names. SIMTEL20 was named for the SIMulation and TELEprocessing organization's DECSYSTEM 20 machine by a vote of the local employees at WSMR.

Along with the original CP/M and MS-DOS collections, the CP/M User's Group, SIG/M(icros), PC/Blue, Ada Software Repository, and the Unix/C collections have been added to SIMTEL20 over time. The MS-DOS collection alone contains over 9,000 ZIPped files in over 200 subject subdirectories. A full range of topics is covered, including learning aides for children, handicap aids, and the latest anti-virus checking programs. SIMTEL20 is the primary distribution center for the INFO-CPM, INFO-MICRO, INFO-IBMPC, INFO-ADA, and other mailing list digests, and hosts the archives for those lists and 20 others.

The availability of freely distributable software programs from SIMTEL20 has had a significant impact on the advancement of computer technology both within the United States, and in many foreign countries, especially those who have recently experienced freedom. Many new software authors (even some Shareware author cooperatives) have emerged from these countries - adding their own freely distributable programs to the SIMTEL20 collection.


The management and staff of SIMTEL20 (and Frank Wancho, who was the system administrator for many years), with the help of Keith Petersen and numerous supporters around the world, are proud to be active, contributing members of the Global Internet Community.

(SIMTEL20 is still owned and operated by the Unites States Army at WSMR. It is connected to the MILNET, and to the Internet through the facilities of New Mexico Technet, Inc., a private non-profit corporation, and the WESTNET Regional Network. Access to its archives is available through Technet and via anonymous FTP.)



info@farnet.org





I LIKE HOW WE'VE HAD THE INTERNET FOR DECADES,
YET "SENDING FILES" IS SOMETHING EARLY
ADOPTERS ARE STILL FIGURING OUT HOW TO DO.

- [|<](#)
- [< Prev](#)
- [Random](#)
- [Next >](#)
- [>|](#)

Permanent link to this comic: <http://xkcd.com/949/>

Image URL (for hotlinking/embedding): http://imgs.xkcd.com/comics/file_transfer.png



US008009861B2

(12) **United States Patent**
Lu et al.

(10) **Patent No.:** **US 8,009,861 B2**
(45) **Date of Patent:** **Aug. 30, 2011**

(54) **METHOD AND SYSTEM FOR FINGERPRINTING DIGITAL VIDEO OBJECT BASED ON MULTIREOLUTION, MULTIRATE SPATIAL AND TEMPORAL SIGNATURES**

2003/0202660 A1* 10/2003 Zhou et al. 380/210
2005/0175224 A1 8/2005 Venkatesan et al.
2006/0165179 A1* 7/2006 Feuer et al. 375/240.18
2007/0253594 A1* 11/2007 Lu et al. 382/100
2008/0273741 A1* 11/2008 Fujii et al. 382/100

FOREIGN PATENT DOCUMENTS

(75) Inventors: **Jian Lu**, Cupertino, CA (US); **Yangbin Wang**, Milpitas, CA (US)

WO WO 2005/079072 * 8/2005
WO WO 2007/127590 A2 11/2007

(73) Assignee: **Vobile, Inc.**, Santa Clara, CA (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1109 days.

Mohatny, S. P., "Digital Watermarking: A Tutorial Review," from Department of Computer Science, University of South Florida, prepared in 1999 at the Indian Institute of Technology, Bangalore, India.
Su et al., "Statistical invisibility for collusion-resistant digital video watermarking," *IEEE Transactions on Multimedia*, Publication Date: Feb. 2005, vol. 7, Issue (1), p. 43-51.

(21) Appl. No.: **11/681,567**

* cited by examiner

(22) Filed: **Mar. 2, 2007**

Primary Examiner — Jingge Wu

(65) **Prior Publication Data**

Assistant Examiner — Tahmina Ansari

US 2007/0253594 A1 Nov. 1, 2007

(74) *Attorney, Agent, or Firm* — Richard T. Ogawa; Ogawa P.C.

Related U.S. Application Data

(57) **ABSTRACT**

(60) Provisional application No. 60/795,786, filed on Apr. 28, 2006.

A method and system for generating a fingerprint for a video object. The method includes obtaining a plurality of frames associated with a video object. Additionally, the method includes, for each of the plurality of frames, processing information associated with the plurality of frames, determining a plurality of spatial signatures for the each of the plurality of frames based on at least information associated with the each of the plurality of frames, and determining a plurality of temporal signatures for the each of the plurality of frames based on at least information associated with the plurality of frames. The plurality of spatial signatures corresponds to a plurality of resolutions respectively, and the plurality of temporal signatures corresponding to a plurality of frame rates respectively.

(51) **Int. Cl.**
G06K 9/00 (2006.01)

(52) **U.S. Cl.** **382/100**


(58) **Field of Classification Search** None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,792,226 A 12/1988 Fishbine et al.
6,957,350 B1* 10/2005 Demos 380/203
2002/0150164 A1* 10/2002 Felts et al. 375/240.19
2003/0185417 A1 10/2003 Alattar et al.

31 Claims, 6 Drawing Sheets






Figure 1: Process of fingerprinting a Video Object




Figure 2: Computing Spatial Signatures at multiple resolutions. (a) a frame is divided into 2x2 blocks; (b) a frame is divided into 4x4 blocks.




Figure 3:: Computing the Base Spatial Signature (BSS) over 2x2 blocks. (a) the mean pixel value of each block; (b) the ordinal rank of each block; (c) the BSS vector.






Figure 4: Positioning a sliding window in computing Temporal Signatures.



(a)



(b)

Figure 5: Computing the Temporal Signature over a downsampled group of frames. (a) the sum of absolute pixel difference between two consecutive frames; (b) the ordinal rank of each frame; the TS of the current frame is its ordinal rank.

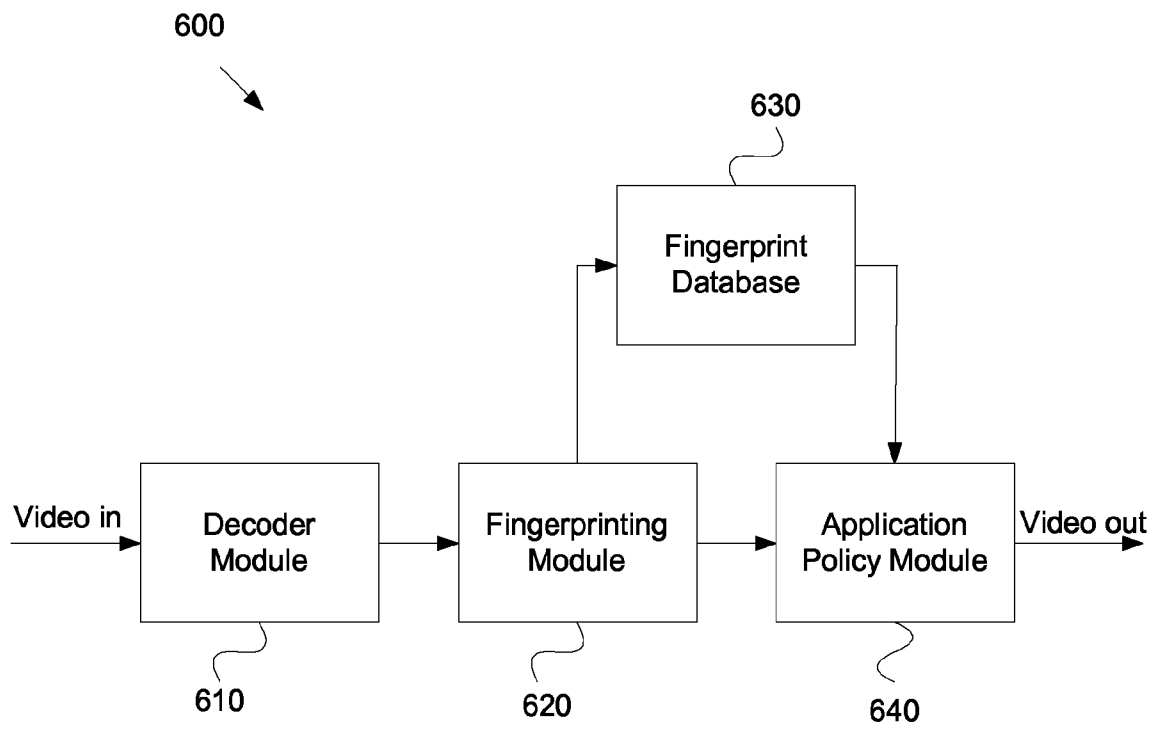


Figure 6

**METHOD AND SYSTEM FOR
FINGERPRINTING DIGITAL VIDEO OBJECT
BASED ON MULTIREOLUTION,
MULTIRATE SPATIAL AND TEMPORAL
SIGNATURES**

CROSS-REFERENCES TO RELATED
APPLICATIONS

This application claims priority to U.S. Provisional Application No. 60/795,786, filed Apr. 28, 2006, which is incorporated by reference herein.

STATEMENT AS TO RIGHTS TO INVENTIONS
MADE UNDER FEDERALLY SPONSORED
RESEARCH OR DEVELOPMENT

Not Applicable

REFERENCE TO A "SEQUENCE LISTING," A
TABLE, OR A COMPUTER PROGRAM LISTING
APPENDIX SUBMITTED ON A COMPACT DISK

Not Applicable

BACKGROUND OF THE INVENTION

This invention relates to techniques for characterizing and fingerprinting digital video object. In particular, this invention relates to method and system for generating a unique and robust identifier for a digital video object based multiresolution, multirate spatial and temporal signatures.

Digital video has become very popular in the last decade. There are many sources from which digital video is created, recorded and distributed, such as DV camcorders, DVD, DVR, and video download and streaming over the Internet. A piece of digital video is called a digital video object or simply video object in this document. It may be a file that is saved on a storage media such as a hard disk drive, or a bitstream that is transmitted over a broadcast channel or over the Internet. The constantly increasing number of digital video objects and proliferation of digital video entertainment and services demand effective and efficient methods and systems for indexing and identifying digital video objects.

A common method for uniquely identifying a digital object is to pass it through a hash function that produces a fixed-length output known as hash sum or message digest. A popular hash function is MD5 that is specified by RFC 1321. While a hash sum as an identifier is useful for certain purposes such as data integrity check, it is often inadequate for content identifications. For example, a digital video object may be encoded in various formats such as MPEG4 and Windows Media, and at various bitrates such as 2 Mbps for broadcast and 700 Kbps for Internet download. The hash sum value will be different for each of these formats though the content is the same. A fingerprint of a digital video object is different from a hash sum in that the former is a unique identifier for the video content while the latter is a unique identifier for the file. Having a unique and robust fingerprint for each and every video object enables many applications, such as video content indexing, search and retrieval, content filtering, broadcast monitoring, and metadata services.

BRIEF SUMMARY OF THE INVENTION

The present invention relates in general to video signal processing. More particularly, the invention provides a

method and system for characterizing a digital video object. Merely by way of example, the invention is described as it applies to obtaining spatial signatures for multiple resolutions, temporal signatures for multiple frame rates, and/or spatial-temporal signatures, but it should be recognized that the invention has a broader range of applicability.

According to an embodiment of the present invention, a method for generating a fingerprint for a video object includes obtaining a plurality of frames associated with a video object. Additionally, the method includes, for each of the plurality of frames, processing information associated with the plurality of frames, determining a plurality of spatial signatures for the each of the plurality of frames based on at least information associated with the each of the plurality of frames, and determining a plurality of temporal signatures for the each of the plurality of frames based on at least information associated with the plurality of frames. The plurality of spatial signatures corresponds to a plurality of resolutions respectively, and the plurality of temporal signatures corresponding to a plurality of frame rates respectively. Moreover, the method includes, for each of the plurality of frames, processing information associated with the plurality of spatial signatures and the plurality of temporal signatures, and determining a frame fingerprint for the each of the plurality of frames, the frame fingerprint including the plurality of spatial signatures corresponding to the plurality of resolutions respectively and the plurality of temporal signatures corresponding to the plurality of frame rates respectively. Also, the method includes processing a plurality of frame fingerprints for the plurality of frames respectively, the plurality of frame fingerprints including the frame fingerprint for the each of the plurality of frames. Additionally, the method includes determining a video fingerprint for the video object, the video fingerprint including the plurality of frame fingerprints.

According to another embodiment of the present invention, a method for generating a spatial signature for a frame of a video object includes obtaining a frame associated with a video object, and dividing the frame into a plurality of blocks, the plurality of blocks corresponding to a plurality of locations respectively. Each of the plurality of blocks includes a plurality of pixels, and the plurality of pixels corresponds to a plurality of pixel values respectively. Additionally, the method includes determining a plurality of average pixel values for the plurality of blocks respectively, processing the plurality of average pixel values, and determining a plurality of ranks for the plurality of blocks respectively based on at least information associated with the plurality of average pixel values. Each of the plurality of ranks corresponds to a block. Moreover, the method includes processing information associated with the plurality of ranks, and determining a sequence of ranks based on at least information associated with the plurality of ranks and the plurality of locations. A spatial signature for the frame includes information associated with the sequence of ranks.

According to yet another embodiment of the present invention, a method for generating a temporal signature for a frame of a video object includes obtaining a first plurality of frames associated with a video object. The first plurality of frames includes at least a frame, each of the first plurality of frames includes a first plurality of pixels and corresponds to an adjacent frame, and the adjacent frame includes a second plurality of pixels. Additionally, the method includes processing information associated with the first plurality of frames, and determining a plurality of difference values for the first plurality of frames respectively. Each of the plurality of difference values corresponds to the each of the first plurality of frames and the adjacent frame. Moreover, the method includes processing

information associated with the plurality of difference values, and determining a plurality of ranks corresponding to the first plurality of frames respectively based on at least information associated with the plurality of difference values. The plurality of ranks includes a rank corresponding to the frame, and the rank is a temporal signature for the frame.

According to yet another embodiment of the present invention, a method for generating a spatial-temporal signature for a frame of a video object includes obtaining a first plurality of frames associated with a video object. The first plurality of frames includes at least a frame, and each of the first plurality of frames corresponds to an adjacent frame. Additionally, the method includes dividing each of the first plurality of frames. The divided frame includes a first plurality of blocks corresponding to a plurality of locations respectively, each of the first plurality of blocks corresponds to a second plurality of blocks including the each of the first plurality of blocks, and the second plurality of blocks is associated with blocks on the first plurality of frames respectively. Moreover, the method includes processing information associated with the first plurality of frames. Also, the method includes, for each of the first plurality of blocks, determining a plurality of difference values for the second plurality of blocks respectively. Each of the plurality of difference values is associated with the each of the second plurality of blocks and a corresponding block on the adjacent frame. Additionally, the method includes, for each of the first plurality of blocks, processing information associated with the plurality of difference values, determining a first plurality of ranks corresponding to the second plurality of blocks respectively based on at least information associated with the plurality of difference values, processing information associated with the first plurality of ranks, and determining a rank for the each of the first plurality of blocks based on at least information associated with the first plurality of ranks. Moreover, the method includes processing information associated with a second plurality of ranks corresponding to the first plurality of blocks respectively, and determining a sequence of ranks based on at least information associated with the second plurality of ranks and the plurality of locations for the first plurality of blocks. The second plurality of ranks includes the rank, and a spatial-temporal signature for the frame includes information associated with the sequence of ranks.

According to yet another embodiment of the present invention, a computer program product includes a computer readable medium including instructions for generating a fingerprint for a video object. The computer readable medium includes one or more instructions for obtaining a plurality of frames associated with a video object. Additionally, the computer readable medium includes one or more instructions for, for each of the plurality of frames, processing information associated with the plurality of frames, determining a plurality of spatial signatures for the each of the plurality of frames based on at least information associated with the each of the plurality of frames, and determining a plurality of temporal signatures for the each of the plurality of frames based on at least information associated with the plurality of frames. The plurality of spatial signatures corresponds to a plurality of resolutions respectively, and the plurality of temporal signatures corresponds to a plurality of frame rates respectively. Moreover, the one or more instructions are for, for each of the plurality of frames, processing information associated with the plurality of spatial signatures and the plurality of temporal signatures, and determining a frame fingerprint for the each of the plurality of frames. The frame fingerprint includes the plurality of spatial signatures corresponding to the plurality of resolutions respectively and the plurality of temporal sig-

natures corresponding to the plurality of frame rates respectively. Also, the computer readable medium includes one or more instructions for processing a plurality of frame fingerprints for the plurality of frames respectively, and one or more instructions for determining a video fingerprint for the video object. The plurality of frame fingerprints includes the frame fingerprint for the each of the plurality of frames, and the video fingerprint includes the plurality of frame fingerprints.

According to yet another embodiment of the present invention, a computer program product includes a computer readable medium including instructions for generating a spatial signature for a frame of a video object. The computer readable medium includes one or more instructions for obtaining a frame associated with a video object, and one or more instructions for dividing the frame into a plurality of blocks. The plurality of blocks corresponds to a plurality of locations respectively, each of the plurality of blocks includes a plurality of pixels, and the plurality of pixels corresponds to a plurality of pixel values respectively. Additionally, the computer readable medium includes one or more instructions for determining a plurality of average pixel values for the plurality of blocks respectively, one or more instructions for processing the plurality of average pixel values, and one or more instructions for determining a plurality of ranks for the plurality of blocks respectively based on at least information associated with the plurality of average pixel values, each of the plurality of ranks corresponding to a block. Moreover, the computer readable medium includes one or more instructions for processing information associated with the plurality of ranks, and one or more instructions for determining a sequence of ranks based on at least information associated with the plurality of ranks and the plurality of locations. A spatial signature for the frame includes information associated with the sequence of ranks.

According to yet another embodiment of the present invention, a computer program product includes a computer readable medium including instructions for generating a temporal signature for a frame of a video object. The computer readable medium includes one or more instructions for obtaining a first plurality of frames associated with a video object. The first plurality of frames includes at least a frame, each of the first plurality of frames includes a first plurality of pixels and corresponds to an adjacent frame, and the adjacent frame includes a second plurality of pixels. Additionally, the computer readable medium includes one or more instructions for processing information associated with the first plurality of frames, and one or more instructions for determining a plurality of difference values for the first plurality of frames respectively. Each of the plurality of difference values corresponding to the each of the first plurality of frames and the adjacent frame. Moreover, the computer readable medium includes one or more instructions for processing information associated with the plurality of difference values, and one or more instructions for determining a plurality of ranks corresponding to the first plurality of frames respectively based on at least information associated with the plurality of difference values. The plurality of ranks includes a rank corresponding to the frame, and the rank is a temporal signature for the frame.

According to yet another embodiment of the present invention, a computer program product includes a computer readable medium including instructions for generating a spatial-temporal signature for a frame of a video object. The computer readable medium includes one or more instructions for obtaining a first plurality of frames associated with a video object. The first plurality of frames includes at least a frame, and each of the first plurality of frames corresponds to an

adjacent frame. Additionally, the computer readable medium includes one or more instructions for dividing each of the first plurality of frames. The divided frame includes a first plurality of blocks corresponding to a plurality of locations respectively, each of the first plurality of blocks corresponds to a second plurality of blocks including the each of the first plurality of blocks, and the second plurality of blocks is associated with blocks on the first plurality of frames respectively. Moreover, the computer readable medium includes one or more instructions for processing information associated with the first plurality of frames. Also, the computer readable medium includes one or more instructions for, for each of the first plurality of blocks, determining a plurality of difference values for the second plurality of blocks respectively, processing information associated with the plurality of difference values, determining a first plurality of ranks corresponding to the second plurality of blocks respectively based on at least information associated with the plurality of difference values, processing information associated with the first plurality of ranks, and determining a rank for the each of the first plurality of blocks based on at least information associated with the first plurality of ranks. Each of the plurality of difference values is associated with the each of the second plurality of blocks and a corresponding block on the adjacent frame. Additionally, the computer readable medium includes one or more instructions for processing information associated with a second plurality of ranks corresponding to the first plurality of blocks respectively, and one or more instructions for determining a sequence of ranks based on at least information associated with the second plurality of ranks and the plurality of locations for the first plurality of blocks. The second plurality of ranks includes the rank, and a spatial-temporal signature for the frame includes information associated with the sequence of ranks.

Many benefits are achieved by way of the present invention over conventional techniques. Certain embodiments of the present invention can generate a robust fingerprint which either does not change or change only slightly with different formats, bitrates, or resolutions, and/or with certain alterations and/or distortions for the same video object. Some embodiments of the present invention can generate a fingerprint that is highly discriminating so that two video objects containing different video contents would yield significantly different fingerprints. Certain embodiments of the present invention can generate a fingerprint that is compact for storage, and can be stored in a form for efficient search and matching.

Depending upon embodiment, one or more of these benefits may be achieved. These benefits and various additional objects, features and advantages of the present invention can be fully appreciated with reference to the detailed description and accompanying drawings that follow.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a flowchart of the algorithm for characterizing and fingerprinting digital video object according to the present invention.

FIG. 2 illustrates the ways of subdividing a frame for computing the Spatial Signatures at multiple resolutions according to the present invention.

FIG. 3 illustrates the process of computing the Base Spatial Signature over 2x2 blocks in a specific embodiment according to present invention.

FIG. 4 illustrates the positioning of a sliding window for computing Temporal Signatures according to present invention.

FIG. 5 illustrates the process of computing the Temporal Signature in a downsampled group of frames in a specific embodiment according to present invention.

FIG. 6 is a simplified system implementing the method for characterizing and fingerprinting a digital video object according to an embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

The present invention covers method and system for characterizing video content by its intrinsic features and transforming these features into a compact signature or fingerprint. Because the same video content may be encoded in different formats, bitrates, or resolutions, and the video content may be cut, edited, or subject to various degree of distortion, it is important that the fingerprint characterizing the video object is robust, to the degree that it is invariant or varying only slightly under these circumstances. On the other hand, the fingerprint must be highly discriminating in the sense that two video objects containing different video content should yield very different fingerprints. Furthermore, the data representation of the video fingerprint must be sufficiently compact for storage, and can be stored in a form for efficient search and matching.

Fingerprinting Video Object

The process of fingerprinting a video object according to the present invention is shown with the block diagram in FIG. 1. Each frame is processed in display order, producing a spatial signature block (SSB), a temporal signature block (TSB), and an optional spatial-temporal signature block (STSB). The SSB consists of one or multiple spatial signatures (SS) in multiple resolutions; the TSB consists of one or multiple temporal signatures (TS) in multiple framerates; the STSB consists of spatial-temporal signatures (STS) in multiple resolutions and multiple framerates. For each frame in a video object, the SSB, TSB, and STSB form a fingerprint block of the corresponding frame. The sequence of all fingerprint blocks in frame order form the fingerprint for the video object. In the following the method for computing SSB, TSB, and STSB is described. In this specific embodiment, all signatures and the fingerprints are computed in Luma component only. In more generalized cases, the signatures and fingerprints can be computed in both Luma and Chroma components, or in any or all components in RGB or other color spaces.

Spatial Signatures

Spatial Signatures (SS) for a video frame can be computed at multiple resolutions. In one implementation according to the present invention, a frame is divided evenly into 2x2 or 4 blocks of equal size, as shown in FIG. 2(a). This is the lowest resolution for frame subdivision in computing the SS; the resulting SS is called Base Spatial Signature, or BSS. Going up one level in resolution, the frame can be divided into 4x4 or 16 blocks, as shown in FIG. 2(b). Going up further in resolution by finer frame subdivision, such as 8x8 or 64 blocks, can produce SS in finer resolutions. It is possible to compute the BSS using a different frame subdivision, such as 3x3 or 9 blocks, and extend to higher resolution by successively doubling its dimensions to, e.g., 6x6=36 blocks, and 12x12=144 blocks. But we find that computing BSS using 2x2=4 blocks is more robust against certain image transformations such as an aspect ratio change.

Using the 2x2 block pattern in FIG. 2(a), the BSS is computed as follows:

1. For each block, compute of mean of pixel value as follows:

$$B_i = (\sum \times (k)) / N_i, k=1, 2, 3, \dots, N_i$$

where B_i is the mean pixel value of the i -th block, $x(k)$ is a pixel value inside the i -th block, and N_i is the number of pixels in the i -th block. See an example in FIG. 3(a).

2. Compare and rank the value of B_i among the blocks and assign the ordinal rank to each block. See an example in FIG. 3(b).
3. Generate the BSS by collecting the ordinal rank of each block in raster order and forming a BSS vector. See an example in FIG. 3(c).

The SS at higher resolution can be computed following the same steps as above. The only difference is that the SS vector at a higher resolution has higher dimensions. For example, the BSS vector computed using 2×2 blocks is of dimension 4, while the SS vector computed using 4×4 blocks is of dimension 16.

Since each ordinal rank in 2×2 blocks can be represented with $\log_2(4)=2$ bits, it's easy to see that the BSS can be represented with $4 \times \log_2(4)=8$ bits. Similarly, it's easy to see that the SS over 4×4 blocks can be represented with $16 \times \log_2(16)=64$ bits. Required bits for representing the SS at even higher resolution can be calculated in a similar way.

The SSB for a frame is formed by stacking the BSS and the SS at all resolutions that are available. In the specific embodiment that is described above, the SSB can be represented with $8+64=72$ bits.

Temporal Signatures

The Temporal Signatures (TS) are computed for each frame over a sliding time window.

For each frame, the window is positioned such that the current frame is on the right edge of the window (See FIG. 4). The position of the window is moved one frame at a time along the temporal axis. The size of the window is a parameter that can be adjusted. In one embodiment according to the present invention, the size of the sliding window is defined to be 1 second in time. This means there will be a number of frames falling in the sliding window at any position of the window, and that number depends on the framerate of the video object. For example, if the framerate of a video object is 30 fps, there will be 30 frames in the sliding window at any position of the window. If the framerate of a video object is less than 1 fps, there will be no frame inside the sliding window for any position of the window.

In order to make the TS comparable for video objects of different framerates, the framerate of each video object is downsampled to a set of common framerates. The set of common framerates for a video object to be downsampled to can be adjusted. Generally speaking, they are designed to be representative of the TS at multiple framerates. In one embodiment according to the present invention, the specific set of common framerates that are used to compute the TS is {6 fps, 12 fps, 24 fps}.

Framerate downsampling to the set of common framerates produces multiple groups of frames. The TS is computed over each group of frames, resulting multirate TS. For clarity in this document, the TS computed from a particularly group will be labeled by the downsampled framerate of that group. For example, TS6 indicates that it is the TS computed from the group of frames of 6 fps.

Using the downsample set of {6 fps, 12 fps, 24 fps} and their associated downsampled groups of frames, the TS is computed in the following steps:

1. For each group of frames, compute the sum of absolute difference of corresponding pixels between two consecutive frames in the group, that is,

$$D_i = \sum |x_i(k) - x_{i-1}(k)|, k=1,2,3 \dots N.$$

where i is the index for the i -th frame in the group, and $k=1, 2, 3, \dots, N$ is the pixel index in a frame. See FIG. 5(a).

2. Compare and rank the value of D_i among the frames in the group and assign the ordinal rank to each frame. See FIG. 5(b).

3. Record the ordinal rank of the current frame (i.e. the frame on the right edge of the sliding window). This is the TS of the current frame in the group. See FIG. 5(b).

Since the ordinal rank is no greater than the number of frames in the group, it is easy to see TS6 can be represented with $\gamma \log_2(6)/3$ bits; TS12 can be represented with $\gamma \log_2(12)/4$ bits; TS24 can be represented with $\gamma \log_2(24)/5$ bits. Here the operator γ denotes a mathematical ceiling function. For example, $\gamma 4.2=5$.

The TSB for a frame is formed by stacking the TS at all framerates that are available. In the specific embodiment that is described above, the TSB can be represented with $3+4+5=12$ bits.

Spatial-Temporal Signatures

Spatial-Temporal Signatures (STS) are computed for each frame over a sliding window of subdivided frames. Frames may be subdivided in various ways as described previously for computing the SS at multiple resolutions. In one specific embodiment, the frame is subdivided into 2×2 blocks. The TS is computed for each block in a frame in the same way as described previously, resulting 4 TS per frame for each downsampled group. The STS is formed by collecting the resulting TS in raster order and put them in a vector. The STSB is formed by stacking the STS at all framerates that are available. In this specific embodiment, the STSB can be represented with $4 \times (3+4+5)=48$ bits.

Fingerprint from Spatial and Temporal Signatures

The collection of SSB, TSB, and optionally STSB for a frame form the Fingerprint Block (FB) for the corresponding frame. A sequence of FBs corresponding to each frame in a video object is defined to be the fingerprint for the video object. In the specific embodiment that is described in this document, a FB can be represented with 84 bits without STSB, or 132 bits with STSB. The fingerprint so defined will have a data rate of 2,520 bits/s (3,930 bits/s if STSB included) for video objects with framerate of 30 fps.

As discussed above, FIG. 1 is a simplified method for characterizing and fingerprinting a digital video object according to an embodiment of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The method 100 includes a process 110 for receiving video object, a process 120 for obtaining next frame, a process 130 for determining one or more spatial signatures, a process 140 for determining one or more temporal signatures, a process 150 for determining one or more spatial-temporal signatures, a process 160 for storing fingerprint block, and a process 170 for determining whether all frames have been processed. Although the above has been shown using a selected group of processes for the method, there can be many alternatives, modifications, and variations. For example, some of the processes may be expanded and/or combined. Other processes may be inserted to those noted above. Depending upon the embodiment, the sequence of processes may be interchanged with others replaced. As an example, some or all processes of the method are performed by a computer or a processor directed by a code. In another example, some or all processes of the method are performed according to instructions included by a computer-readable medium in a computer pro-

gram product. Further details of these processes are found throughout the present specification.

After a video object is received at the process 110, the video object is processed. For example, the video object is a piece of video, such as a piece of digital video. In another example, the video object includes one or more frames, which are obtained and processed according to FIG. 1. In one embodiment, the frames are processed in their display order. In another embodiment, for each frame, a spatial signature block (SSB), a temporal signature block (TSB), and/or a spatial-temporal signature block (STSB) are determined. For example, the SSB includes one or more spatial signatures (SS) in one or more resolutions; the TSB includes one or more temporal signatures (TS) in one or more framerates; and/or the STSB includes one or more spatial-temporal signatures (STS) in one or more resolutions and one or more framerates. In yet another embodiment, the SSB, the TSB, and/or the STSB of the same frame form at least part of a fingerprint block of this frame. For example, a fingerprint block is a frame fingerprint. According to FIG. 1, after the fingerprint block is stored, it is determined whether all frames have been processed at the process 170. If not all frames have been processed, next frame is obtained and processed. The sequence of all fingerprint blocks in frame order form at least part of fingerprint for the video object according to an embodiment of the present invention.

Further emphasized here, FIG. 1 is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, the process 170 for determining whether all frames have been processed can be replaced by a process for determining whether additional frames need to be processed. In another example, one or more of the process 130 for determining one or more spatial signatures, the process 140 for determining one or more temporal signatures, and the process 150 for determining one or more spatial-temporal signatures are skipped.

As discussed above, FIG. 2 illustrates examples of subdividing a frame for determining Spatial Signatures at multiple resolutions according to an embodiment of the present invention. FIG. 2 includes FIGS. 2(a) and (b). These diagrams are merely examples, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications.

As shown in FIG. 2(a), a frame is divided into m_b by n_b blocks according to an embodiment. For example, each of m_b and n_b is a positive integer. In another example, m_b and n_b are the same or different in value. In yet another example, all of the m_b by n_b blocks have the same size. As shown in FIG. 2(b), the frame is divided into m_h by n_h blocks according to an embodiment. For example, each of m_h and n_h is a positive integer. In another example, m_h and n_h are the same or different in value. In yet another example, all of the m_h by n_h blocks have the same size.

In one embodiment, the spatial signature (SS) based on m_b by n_b blocks is called Base Spatial Signature (BSS), and the spatial signature (SS) based on m_h by n_h blocks is called Spatial Signature (SS) at higher resolution. For example, m_h is larger than m_b , and/or n_h is larger than n_b . In another example, both m_b and n_b are equal to 2, and both m_h by n_h are equal to 2 multiplied by 2^n , where n is a positive integer. In yet another example, both m_b and n_b are equal to 3, and both m_h by n_h are equal to 3 multiplied by 2^n .

As discussed above, FIG. 3 illustrates a process of determining the Base Spatial Signature over 2×2 blocks according to an embodiment of the present invention. FIG. 3 includes FIGS. 3(a), (b), and (c). These diagrams are merely examples,

which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, the BSS based on m_b by n_b blocks can be calculated in substantially the same manner, even if either m_b or n_b or both are not equal to 2. In another example, the SS based on m_h by n_h blocks can be calculated in substantially the same manner.

As shown in FIG. 3, the method for determining spatial signature based on m_b by n_b blocks or m_h by n_h blocks includes the following three processes:

1. For each block, determining average of pixel values as follows:

$$B_i = \sum_{k=1}^{N_i} x(k) / N_i$$

where B_i is the average pixel value of the i -th block, $x(k)$ is the pixel value for the k th pixel inside the i -th block, and N_i is the number of pixels in the i -th block. In one embodiment, $i=1, 2, \dots, m_b \times n_b$. In another embodiment, $i=1, 2, \dots, m_h \times n_h$.

2. Determining a ranking number for each block. For example, the value of B_i is compared among the blocks, and the ordinal ranking number for the B_i is assigned to the corresponding block. In one embodiment, the total number of blocks equals $m_b \times n_b$, so the ranking number ranges from 1 to $m_b \times n_b$. In another embodiment, the total number of blocks equals $m_h \times n_h$, so the ranking number ranges from 1 to $m_h \times n_h$.

3. Determining spatial signature based on ranking numbers. For example, the spatial signature is BSS, or SS at higher resolution. In another example, the spatial signature includes a vector. Within the vector, the ranking numbers for the blocks are arranged based on the physical locations of these blocks within the frame. In one embodiment, the ranking numbers are collected in raster order and thus a spatial signature is generated.

After spatial signatures at multiple resolutions are determined for a frame, the SSB is determined by stacking the BSS and the SS at one or more higher resolutions according to an embodiment of the present invention.

FIG. 4 illustrates the positioning of a sliding window for computing Temporal Signatures according to an embodiment of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications.

As shown in FIG. 4, for each frame, the sliding window is positioned such that the current frame is on the right edge of the window. The position of the window is moved one frame at a time along the temporal axis. The width of the window is a parameter that can be adjusted. In one embodiment, the width of the window is 1 second in time. For example, a number of frames fall within the sliding window, and the number of such frames depends on the framerate of the video object. According to one embodiment, if the framerate used for determining temporal signature is the original framerate of the video object, the frame or frames are the original frame or frames of the video object. According to another embodiment, if the framerate used for determining temporal signature is different from the original framerate of the video object, the frame or frames are the sampled frame or frames of the video object. For example, the framerate used for determining temporal signature is lower than the original framerate of the video object, so the video object is temporally downsampled.

For a particular sampled framerate used, the method for determining temporal signature includes the following three processes:

1. Determining differences between a frame and its adjacent frame. For example, the sum of absolute differences of corresponding pixels between two consecutive frames is calculated as follows:

$$D_i = \sum_{k=1}^N |x_i(k) - x_{i-1}(k)|$$

where i is the index for a frame. For example, D_i is calculated for each frame within the sliding window. Additionally, k represents the k th pixel and N represents the total number of pixels in the frame.

2. Determining a ranking number for the current frame. For example, the value of D_i are compared among all the frames within the sliding window, and the ordinal ranking number for the D_i is assigned to the corresponding frame. In one embodiment, the ordinal ranking number is assigned to the current frame, which resides at the right edge of the sliding window.

3. Recording the ordinal ranking number of the current frame as the temporal signature of the current frame.

After temporal signatures based on multiple framerates are determined for a frame, the TSB is determined by stacking the TS at multiple framerates according to an embodiment of the present invention.

As discussed above, Spatial-Temporal Signatures (STS) are computed for each frame over a sliding window of subdivided frames according to an embodiment of the present invention. For example, a frame is divided into m by n blocks, where each of m and n is a positive integer. In another example, m and n are the same or different in value. In yet another example, all of the m by n blocks have the same size.

In one embodiment, the TS is computed for each block in the frame. For a particular sampled framerate used, the method for determining temporal signature for each block includes the following two processes:

1. Determining differences between a block on a frame and its corresponding block on an adjacent frame. According to one embodiment, if the framerate used for determining temporal signature is the original framerate of the video object, the frames are the original frames of the video object. According to another embodiment, if the framerate used for determining temporal signature is different from the original framerate of the video object, the frames are the sampled frames of the video object. For example, the framerate used for determining temporal signature is lower than the original framerate of the video object, so the video object is temporally downsampled.

For example, the sum of absolute differences of corresponding pixels between two corresponding blocks on two consecutive frames is calculated as follows:

$$D_i = \sum_{k=1}^N |x_i(k) - x_{i-1}(k)|$$

where i is the index for a frame. For example, D_i is calculated for the corresponding block of each frame within the sliding window. Additionally, k represents the k th pixel within the corresponding block, and N represents the total number of pixels in the block.

2. Determining a ranking number for the block on the current frame. For example, the value of D_i are compared among all the corresponding blocks on all the frames within the sliding window, and the ordinal ranking number for the D_i is assigned to the corresponding block. In one embodiment, the ordinal ranking number is assigned to the block on the current frame, which resides at the right edge of the sliding window.

These two processes are repeated to determine the ordinal ranking numbers for all m -by- n blocks on the current frame according to an embodiment. The spatial-temporal signature is then determined based on ranking numbers of blocks on the current frame. For example, the spatial-temporal signature includes a vector. Within the vector, the ranking numbers for the blocks are arranged based on the physical locations of these blocks within the current frame. In one embodiment, the ranking numbers are collected in raster order and thus a spatial-temporal signature is generated.

As discussed, the spatial-temporal signature is determined based on the framerate used for determining the TS for each block of the current frame. In one embodiment, spatial-temporal signatures are determined based on multiple framerates for the current frame, and the STSB is determined by stacking the STS at multiple framerates according to an embodiment of the present invention.

For each frame, one or more of SSB, TSB, and STSB, with or without any other information, can form the Fingerprint Block (FB) for the corresponding frame according to an embodiment of the present invention. For example, a Fingerprint Block (FB) is a frame fingerprint, which includes one or more spatial signatures, one or more temporal signatures, and/or one or more spatial-temporal signatures. For a sequence of frames of a video object, the corresponding sequence of FBs can be used as the fingerprint for the video object according to another embodiment of the present invention. For example, if the framerate used for determining temporal signature is the original framerate of the video object, the frames are the original frames of the video object. According to another embodiment, if the framerate used for determining temporal signature is different from the original framerate of the video object, the frames are the sampled frames of the video object. The fingerprint of the video object is stored in a database according to an embodiment of the present invention. For example, the database includes one or more fingerprints of one or more corresponding video objects.

FIG. 6 is a simplified system implementing the method 100 for characterizing and fingerprinting a digital video object according to an embodiment of the present invention. This diagram is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. The system 600 includes a decoder module 610, a fingerprinting module 620, a fingerprint database 630, and an application policy module 640. Although the above has been shown using a selected group of components for the system, there can be many alternatives, modifications, and variations. For example, some of the modules may be expanded and/or combined. Other modules may be inserted to those noted above. Depending upon the embodiment, the arrangement of modules may be interchanged with others replaced. As an example, some or all processes of the method are performed by a computer or a processor directed by a code. In another example, some or all processes of the method are performed according to instructions included by a computer-readable medium in a computer program product. Further details of these processes are found throughout the present specification.

As shown in FIG. 6, an input video is decoded by the decoder module 610 and fed to the fingerprinting module 620. The fingerprinting module 620 performs the method 100 according to an embodiment of the present invention. For example, the fingerprinting module 620 is implemented according to FIGS. 1, 2, 3, 4, and/or 5. The resulting video fingerprint is compared to the ones stored in the fingerprint database for identification, and the identification result is returned to the application along with associated metadata (e.g., title and ownership of the video content). Based on the identification result, the application applies certain policy at the application policy module 640. For example, if the video is identified to be a pirated version or copy, the application applies filtering.

As discussed above and further emphasized here, FIG. 6 is merely an example, which should not unduly limit the scope of the claims. One of ordinary skill in the art would recognize many variations, alternatives, and modifications. For example, anyone of the modules 610, 620, 630, and 640 can be either hardware or software, or a combination of hardware and software. In another example, the fingerprint database 630 can be embedded in an application or resided outside the application on a local hard drive or a remote server.

According to another embodiment of the present invention, a method for generating a fingerprint for a video object includes obtaining a plurality of frames associated with a video object. Additionally, the method includes, for each of the plurality of frames, processing information associated with the plurality of frames, determining a plurality of spatial signatures for the each of the plurality of frames based on at least information associated with the each of the plurality of frames, and determining a plurality of temporal signatures for the each of the plurality of frames based on at least information associated with the plurality of frames. The plurality of spatial signatures corresponds to a plurality of resolutions respectively, and the plurality of temporal signatures corresponding to a plurality of frame rates respectively. Moreover, the method includes, for each of the plurality of frames, processing information associated with the plurality of spatial signatures and the plurality of temporal signatures, and determining a frame fingerprint for the each of the plurality of frames, the frame fingerprint including the plurality of spatial signatures corresponding to the plurality of resolutions respectively and the plurality of temporal signatures corresponding to the plurality of frame rates respectively. Also, the method includes processing a plurality of frame fingerprints for the plurality of frames respectively, the plurality of frame fingerprints including the frame fingerprint for the each of the plurality of frames. Additionally, the method includes determining a video fingerprint for the video object, the video fingerprint including the plurality of frame fingerprints. For example, the method is implemented according to FIGS. 1, 2, 3, 4, and/or 5.

According to yet another embodiment of the present invention, a method for generating a spatial signature for a frame of a video object includes obtaining a frame associated with a video object, and dividing the frame into a plurality of blocks, the plurality of blocks corresponding to a plurality of locations respectively. Each of the plurality of blocks includes a plurality of pixels, and the plurality of pixels corresponds to a plurality of pixel values respectively. Additionally, the method includes determining a plurality of average pixel values for the plurality of blocks respectively, processing the plurality of average pixel values, and determining a plurality of ranks for the plurality of blocks respectively based on at least information associated with the plurality of average pixel values. Each of the plurality of ranks corresponds to a

block. Moreover, the method includes processing information associated with the plurality of ranks, and determining a sequence of ranks based on at least information associated with the plurality of ranks and the plurality of locations. A spatial signature for the frame includes information associated with the sequence of ranks. For example, the method is implemented according to FIGS. 1, 2, and/or 3.

According to yet another embodiment of the present invention, a method for generating a temporal signature for a frame of a video object includes obtaining a first plurality of frames associated with a video object. The first plurality of frames includes at least a frame, each of the first plurality of frames includes a first plurality of pixels and corresponds to an adjacent frame, and the adjacent frame includes a second plurality of pixels. Additionally, the method includes processing information associated with the first plurality of frames, and determining a plurality of difference values for the first plurality of frames respectively. Each of the plurality of difference values corresponds to the each of the first plurality of frames and the adjacent frame. Moreover, the method includes processing information associated with the plurality of difference values, and determining a plurality of ranks corresponding to the first plurality of frames respectively based on at least information associated with the plurality of difference values. The plurality of ranks includes a rank corresponding to the frame, and the rank is a temporal signature for the frame. For example, the method is implemented according to FIGS. 1, 4, and/or 5.

According to yet another embodiment of the present invention, a method for generating a spatial-temporal signature for a frame of a video object includes obtaining a first plurality of frames associated with a video object. The first plurality of frames includes at least a frame, and each of the first plurality of frames corresponds to an adjacent frame. Additionally, the method includes dividing each of the first plurality of frames. The divided frame includes a first plurality of blocks corresponding to a plurality of locations respectively, each of the first plurality of blocks corresponds to a second plurality of blocks including the each of the first plurality of blocks, and the second plurality of blocks is associated with blocks on the first plurality of frames respectively. Moreover, the method includes processing information associated with the first plurality of frames. Also, the method includes, for each of the first plurality of blocks, determining a plurality of difference values for the second plurality of blocks respectively. Each of the plurality of difference values is associated with the each of the second plurality of blocks and a corresponding block on the adjacent frame. Additionally, the method includes, for each of the first plurality of blocks, processing information associated with the plurality of difference values, determining a first plurality of ranks corresponding to the second plurality of blocks respectively based on at least information associated with the plurality of difference values, processing information associated with the first plurality of ranks, and determining a rank for the each of the first plurality of blocks based on at least information associated with the first plurality of ranks. Moreover, the method includes processing information associated with a second plurality of ranks corresponding to the first plurality of blocks respectively, and determining a sequence of ranks based on at least information associated with the second plurality of ranks and the plurality of locations for the first plurality of blocks. The second plurality of ranks includes the rank, and a spatial-temporal signature for the frame includes information associated with the sequence of ranks. For example, the method is implemented according to FIGS. 1, 2, 3, 4, and/or 5.

According to yet another embodiment of the present invention, a computer program product includes a computer read-

able medium including instructions for generating a fingerprint for a video object. The computer readable medium includes one or more instructions for obtaining a plurality of frames associated with a video object. Additionally, the computer readable medium includes one or more instructions for, for each of the plurality of frames, processing information associated with the plurality of frames, determining a plurality of spatial signatures for the each of the plurality of frames based on at least information associated with the each of the plurality of frames, and determining a plurality of temporal signatures for the each of the plurality of frames based on at least information associated with the plurality of frames. The plurality of spatial signatures corresponds to a plurality of resolutions respectively, and the plurality of temporal signatures corresponds to a plurality of frame rates respectively. Moreover, the one or more instructions are for, for each of the plurality of frames, processing information associated with the plurality of spatial signatures and the plurality of temporal signatures, and determining a frame fingerprint for the each of the plurality of frames. The frame fingerprint includes the plurality of spatial signatures corresponding to the plurality of resolutions respectively and the plurality of temporal signatures corresponding to the plurality of frame rates respectively. Also, the computer readable medium includes one or more instructions for processing a plurality of frame fingerprints for the plurality of frames respectively, and one or more instructions for determining a video fingerprint for the video object. The plurality of frame fingerprints includes the frame fingerprint for the each of the plurality of frames, and the video fingerprint includes the plurality of frame fingerprints. For example, the computer program product is implemented according to FIGS. 1, 2, 3, 4, 5 and/or 6.

According to yet another embodiment of the present invention, a computer program product includes a computer readable medium including instructions for generating a spatial signature for a frame of a video object. The computer readable medium includes one or more instructions for obtaining a frame associated with a video object, and one or more instructions for dividing the frame into a plurality of blocks. The plurality of blocks corresponds to a plurality of locations respectively, each of the plurality of blocks includes a plurality of pixels, and the plurality of pixels corresponds to a plurality of pixel values respectively. Additionally, the computer readable medium includes one or more instructions for determining a plurality of average pixel values for the plurality of blocks respectively, one or more instructions for processing the plurality of average pixel values, and one or more instructions for determining a plurality of ranks for the plurality of blocks respectively based on at least information associated with the plurality of average pixel values, each of the plurality of ranks corresponding to a block. Moreover, the computer readable medium includes one or more instructions for processing information associated with the plurality of ranks, and one or more instructions for determining a sequence of ranks based on at least information associated with the plurality of ranks and the plurality of locations. A spatial signature for the frame includes information associated with the sequence of ranks. For example, the computer program product is implemented according to FIGS. 1, 2, 3, and/or 6.

According to yet another embodiment of the present invention, a computer program product includes a computer readable medium including instructions for generating a temporal signature for a frame of a video object. The computer readable medium includes one or more instructions for obtaining a first plurality of frames associated with a video object. The first plurality of frames includes at least a frame, each of the first

plurality of frames includes a first plurality of pixels and corresponds to an adjacent frame, and the adjacent frame includes a second plurality of pixels. Additionally, the computer readable medium includes one or more instructions for processing information associated with the first plurality of frames, and one or more instructions for determining a plurality of difference values for the first plurality of frames respectively. Each of the plurality of difference values corresponding to the each of the first plurality of frames and the adjacent frame. Moreover, the computer readable medium includes one or more instructions for processing information associated with the plurality of difference values, and one or more instructions for determining a plurality of ranks corresponding to the first plurality of frames respectively based on at least information associated with the plurality of difference values. The plurality of ranks includes a rank corresponding to the frame, and the rank is a temporal signature for the frame. For example, the computer program product is implemented according to FIGS. 1, 4, 5 and/or 6.

According to yet another embodiment of the present invention, a computer program product includes a computer readable medium including instructions for generating a spatial-temporal signature for a frame of a video object. The computer readable medium includes one or more instructions for obtaining a first plurality of frames associated with a video object. The first plurality of frames includes at least a frame, and each of the first plurality of frames corresponds to an adjacent frame. Additionally, the computer readable medium includes one or more instructions for dividing each of the first plurality of frames. The divided frame includes a first plurality of blocks corresponding to a plurality of locations respectively, each of the first plurality of blocks corresponds to a second plurality of blocks including the each of the first plurality of blocks, and the second plurality of blocks is associated with blocks on the first plurality of frames respectively. Moreover, the computer readable medium includes one or more instructions for processing information associated with the first plurality of frames. Also, the computer readable medium includes one or more instructions for, for each of the first plurality of blocks, determining a plurality of difference values for the second plurality of blocks respectively, processing information associated with the plurality of difference values, determining a first plurality of ranks corresponding to the second plurality of blocks respectively based on at least information associated with the plurality of difference values, processing information associated with the first plurality of ranks, and determining a rank for the each of the first plurality of blocks based on at least information associated with the first plurality of ranks. Each of the plurality of difference values is associated with the each of the second plurality of blocks and a corresponding block on the adjacent frame. Additionally, the computer readable medium includes one or more instructions for processing information associated with a second plurality of ranks corresponding to the first plurality of blocks respectively, and one or more instructions for determining a sequence of ranks based on at least information associated with the second plurality of ranks and the plurality of locations for the first plurality of blocks. The second plurality of ranks includes the rank, and a spatial-temporal signature for the frame includes information associated with the sequence of ranks. For example, the computer program product is implemented according to FIGS. 1, 2, 3, 4, 5 and/or 6.

The present invention has various advantages. Certain embodiments of the present invention can generate a robust fingerprint which either does not change or change only slightly with different formats, bitrates, or resolutions, and/or

with certain alterations and/or distortions for the same video object. Some embodiments of the present invention can generate a fingerprint that is highly discriminating so that two video objects containing different video contents would yield significantly different fingerprints. Certain embodiments of the present invention can generate a fingerprint that is compact for storage, and can be stored in a form for efficient search and matching.

Although specific embodiments of the present invention have been described, it will be understood by those of skill in the art that there are other embodiments that are equivalent to the described embodiments. Accordingly, it is to be understood that the invention is not to be limited by the specific illustrated embodiments, but only by the scope of the appended claims.

What is claimed is:

1. A computer implemented method for generating a fingerprint for a video object that is performed by a computer system programmed to perform the method, comprising:

obtaining, by the computer system, a plurality of frames associated with a video object;

for each of the plurality of frames,

processing, by the computer system, information associated with the plurality of frames;

determining, by the computer system, a plurality of spatial signatures for the each of the plurality of frames based on at least information associated with the each of the plurality of frames, the plurality of spatial signatures corresponding to a plurality of resolutions respectively;

determining, by the computer system, a plurality of temporal signatures for the each of the plurality of frames based on at least information associated with the plurality of frames, the plurality of temporal signatures corresponding to a plurality of frame rates respectively;

processing, by the computer system, information associated with the plurality of spatial signatures and the plurality of temporal signatures; and

determining a frame fingerprint for the each of the plurality of frames, the frame fingerprint including the plurality of spatial signatures corresponding to the plurality of resolutions respectively and the plurality of temporal signatures corresponding to the plurality of frame rates respectively;

processing, by the computer system, a plurality of frame fingerprints for the plurality of frames respectively, the plurality of frame fingerprints including the frame fingerprint for the each of the plurality of frames;

determining, by the computer system, a video fingerprint for the video object, the video fingerprint including the plurality of frame fingerprints; and

storing, by the computer system, the video fingerprint for the video object in a data storage separate from the video object, wherein the video object is not modified by the video fingerprint.

2. The method of claim 1, and further comprising:

for the each of the plurality of frames,

determining, by the computer system, a plurality of spatial-temporal signatures for the each of the plurality of frames based on at least information associated with the plurality of frames, the plurality of spatial-temporal signatures corresponding to a plurality of combinations of resolutions and frame rates respectively;

processing, by the computer system, information associated with the plurality of spatial-temporal signatures;

wherein the frame fingerprint further including the plurality of spatial-temporal signatures corresponding to the plurality of combinations of resolutions and frame rates respectively.

3. The method of claim 1, and further comprising:

for the each of the plurality of frames, storing, by the computer system the frame fingerprint for the each of the plurality of frames in the data storage.

4. The method of claim 1, and further comprising storing, by the computer system, the video fingerprint for the video object in a video fingerprint database.

5. The method of claim 1 wherein the determining a plurality of spatial signatures comprises for each of the plurality of resolutions, dividing, by the computer system, the each of the plurality of frames into a number of columns and a number of rows, the each of the plurality of resolutions begin represented by, at least, the number of columns and the number of rows.

6. The method of claim 1 wherein the determining a plurality of temporal signatures comprises for each of the plurality of frame rates, determining, by the computer system, a difference value corresponding to the each of the plurality of frames and an adjacent frame associated with the each of the plurality of frames, the each of the plurality of frame rates being inversely proportional to an time interval between the each of the plurality of frames and the adjacent frame.

7. A computer implemented method for generating a spatial signature for a frame of a video object with a computer system programmed to perform the method, comprising:

obtaining, by the computer system, a frame associated with a video object;

dividing, by the computer system, the frame into a plurality of blocks, the plurality of blocks corresponding to a plurality of locations respectively, each of the plurality of blocks including a plurality of pixels, the plurality of pixels corresponding to a plurality of pixel values respectively;

determining, by the computer system, a plurality of average pixel values for the plurality of blocks respectively; processing, by the computer system, the plurality of average pixel values;

determining, by the computer system, a plurality of ranks for the plurality of blocks respectively based on at least information associated with the plurality of average pixel values, each of the plurality of ranks corresponding to a block;

processing, by the computer system, information associated with the plurality of ranks;

determining, by the computer system, a sequence of ranks based on at least information associated with the plurality of ranks and the plurality of locations; and

storing a spatial signature for the frame in a data storage separate from the video object, wherein the spatial signature for the frame includes information associated with the sequence of ranks, and wherein the frame of the video object is not modified by the spatial signature.

8. The method of claim 7 wherein the determining a sequence of ranks comprises:

scanning, by the computer system, the plurality of blocks based on a predetermined pattern;

determining, by the computer system, a sequence of blocks as a result of the scanning process; and

arranging, by the computer system, the each of the plurality of ranks based on a location of the corresponding block in the sequence of blocks.

9. The method of claim 8 wherein the predetermined pattern is associated with a raster scan.

19

10. The method of claim 7 wherein the plurality of ranks is a plurality of ordinal ranks.

11. The method of claim 7 wherein:

the dividing the frame into a plurality of blocks includes dividing, by the computer system, the frame into a number of columns and a number of rows;

wherein the spatial signature for the frame corresponds to a resolution represented by, at least, the number of columns and the number of rows.

12. The method of claim 11 wherein the resolution is a base resolution, and the spatial signature is a based spatial signature.

13. A computer implemented method for generating a temporal signature for a frame of a video object that is performed by a computer system programmed to perform the method, comprising:

obtaining, by the computer system, a first plurality of frames associated with a video object, the first plurality of frames including at least a frame, each of the first plurality of frames including a first plurality of pixels and corresponding to an adjacent frame, the adjacent frame including a second plurality of pixels;

processing, by the computer system, information associated with the first plurality of frames;

determining, by the computer system, a plurality of difference values for the first plurality of frames respectively, each of the plurality of difference values corresponding to the each of the first plurality of frames and the adjacent frame;

processing, by the computer system, information associated with the plurality of difference values;

determining, by the computer system, a plurality of ranks corresponding to the first plurality of frames respectively based on at least information associated with the plurality of difference values, the plurality of ranks including a rank corresponding to the frame, the rank being a temporal signature for the frame; and

storing, by the computer system, the temporal signature for the frame in a data store separate from the video object, wherein the first plurality of frames associated with the video object are not modified by the temporal signature.

14. The method of claim 13, and further comprising:

obtaining, by the computer system, a second plurality of frames;

processing, by the computer system, information associated with the second plurality of frames; and

determining, by the computer system, the first plurality of frames based on at least information associated with the second plurality of frames.

15. The method of claim 14 wherein the determining the first plurality of frames comprises:

receiving, by the computer system, information associated with a time period including a start time and an end time;

placing, by the computer system, the frame at the start time of the time period; and

selecting, by the computer system, at least some frames from the second plurality of frames within the time period, the first plurality of frames including the selected frames.

16. The method of claim 13 wherein the determining a plurality of difference values comprises:

for the each of the first plurality of frames,

determining, by the computer system, magnitudes of differences between first values for the first plurality of pixels and second values for the second plurality of pixels respectively;

20

summing, by the computer system, the magnitudes of differences to determine one of the plurality of difference values.

17. The method of claim 13 wherein the plurality of ranks is a plurality of ordinal ranks.

18. The method of claim 13 wherein:

the first plurality of frames corresponds to a first frame rate, the rank being the temporal signature for the frame corresponding to the first frame rate.

19. The method of claim 18, wherein

the video object corresponds to a second frame rate; wherein if the first frame rate is lower than the second frame rate, the first plurality of frames are down-sampled frames for the video object.

20. A computer implemented method for generating a spatial-temporal signature for a frame of a video object with a computer system programmed to perform the method, comprising:

obtaining, by the computer system, a first plurality of frames associated with a video object, the first plurality of frames including at least a frame, each of the first plurality of frames corresponding to an adjacent frame;

dividing, by the computer system, each of the first plurality of frames, wherein the divided frame includes a first plurality of blocks corresponding to a plurality of locations respectively, each of the first plurality of blocks corresponding to a second plurality of blocks including the each of the first plurality of blocks, the second plurality of blocks being associated with blocks on the first plurality of frames respectively;

processing, by the computer system, information associated with the first plurality of frames;

for each of the first plurality of blocks,

determining, by the computer system, a plurality of difference values for the second plurality of blocks respectively, each of the plurality of difference values being associated with the each of the second plurality of blocks and a corresponding block on the adjacent frame;

processing, by the computer system, information associated with the plurality of difference values;

determining, by the computer system, a first plurality of ranks corresponding to the second plurality of blocks respectively based on at least information associated with the plurality of difference values;

processing, by the computer system, information associated with the first plurality of ranks;

determining, by the computer system, a rank for the each of the first plurality of blocks based on at least information associated with the first plurality of ranks;

processing, by the computer system, information associated with a second plurality of ranks corresponding to the first plurality of blocks respectively, the second plurality of ranks including the rank;

determining, by the computer system, a sequence of ranks based on at least information associated with the second plurality of ranks and the plurality of locations for the first plurality of blocks; and

storing, by the computer system, a spatial-temporal signature for the frame independent of the video object, wherein the spatial-temporal signature for the frame includes information associated with the sequence of ranks, and wherein the first plurality of frames associated with the video object are not modified by the spatial-temporal signature.

21. The method of claim 20 wherein the first plurality of blocks are on a same frame at a same instant in time.

21

22. The method of claim 20 wherein the second plurality of blocks are at a same location on different frames respectively along a temporal axis.

23. The method of claim 20 wherein the determining a sequence of ranks comprises:

scanning, by the computer system, the first plurality of blocks based on a predetermined pattern;

determining, by the computer system, a sequence of blocks on the frame as a result of the scanning process; and arranging, by the computer system, the each of the second plurality of ranks based on a location of the corresponding block in the sequence of blocks.

24. The method of claim 20 wherein:

the first plurality of frames corresponds to a first frame rate, the spatial-temporal signature for the frame corresponding to the first frame rate.

25. The method of claim 24, wherein

the video object corresponds to a second frame rate; if the first frame rate is lower than the second frame rate, the first plurality of frames are down-sampled frames for the video object.

26. The method of claim 20, and further comprising:

obtaining, by the computer system, a second plurality of frames;

processing, by the computer system, information associated with the second plurality of frames;

determining, by the computer system, the first plurality of frames based on at least information associated with the second plurality of frames.

27. The method of claim 26 wherein the determining the first plurality of frames comprises:

receiving, by the computer system, information associated with a time period including a start time and an end time;

placing, by the computer system, the frame at the start time of the time period; and

selecting, by the computer system, at least some frames from the second plurality of frames within the time period, the first plurality of frames including the selected frames.

28. A non-transitory computer readable medium including instructions that programs a processor of a computer system to generate a fingerprint for a video object, the non-transitory computer readable medium comprising:

one or more instructions that programs the processor to obtain a plurality of frames associated with a video object;

one or more instructions for each of the plurality of frames, that programs the processor to process information associated with the plurality of frames;

that programs the processor to determine a plurality of spatial signatures for the each of the plurality of frames based on at least information associated with the each of the plurality of frames, the plurality of spatial signatures corresponding to a plurality of resolutions respectively;

that programs the processor to determine a plurality of temporal signatures for the each of the plurality of frames based on at least information associated with the plurality of frames, the plurality of temporal signatures corresponding to a plurality of frame rates respectively;

that programs the processor to process information associated with the plurality of spatial signatures and the plurality of temporal signatures;

that programs the processor to determine a frame fingerprint for the each of the plurality of frames, the frame fingerprint including the plurality of spatial signa-

22

tures corresponding to the plurality of resolutions respectively and the plurality of temporal signatures corresponding to the plurality of frame rates respectively;

one or more instructions that programs the processor to process a plurality of frame fingerprints for the plurality of frames respectively, the plurality of frame fingerprints including the frame fingerprint for the each of the plurality of frames;

one or more instructions that programs the processor to determine a video fingerprint for the video object, the video fingerprint including the plurality of frame fingerprints; and

one or more instructions that programs to processor to store the video fingerprint for the video object in a memory storage separately from the video object, wherein the plurality of frames associated with the video object are not modified.

29. A non-transitory computer readable medium including instructions that programs a processor of a computer system to generate a spatial signature for a frame of a video object, the non-transitory computer readable medium comprising:

one or more instructions that programs the processor to obtain a frame associated with a video object;

one or more instructions that programs the processor to divide the frame into a plurality of blocks, the plurality of blocks corresponding to a plurality of locations respectively, each of the plurality of blocks including a plurality of pixels, the plurality of pixels corresponding to a plurality of pixel values respectively;

one or more instructions that programs the processor to determine a plurality of average pixel values for the plurality of blocks respectively;

one or more instructions that programs the processor to process the plurality of average pixel values;

one or more instructions that programs the processor to determine a plurality of ranks for the plurality of blocks respectively based on at least information associated with the plurality of average pixel values, each of the plurality of ranks corresponding to a block;

one or more instructions that programs the processor to process information associated with the plurality of ranks;

one or more instructions that programs the processor to determine a sequence of ranks based on at least information associated with the plurality of ranks and the plurality of locations; and

one or more instructions that programs the processor to store a spatial signature for the frame in a memory separately from the video object, wherein the spatial signature for the frame includes information associated with the sequence of ranks, wherein the frame associated with the video object is not modified by the spatial signature.

30. A non-transitory computer readable medium including instructions that programs a processor of a computer system to generate a temporal signature for a frame of a video object, the non-transitory computer readable medium comprising:

one or more instructions that programs the processor to obtain a first plurality of frames associated with a video object, the first plurality of frames including at least a frame, each of the first plurality of frames including a first plurality of pixels and corresponding to an adjacent frame, the adjacent frame including a second plurality of pixels;

one or more instructions that programs the processor to process information associated with the first plurality of frames;

23

one or more instructions that programs the processor to determine a plurality of difference values for the first plurality of frames respectively, each of the plurality of difference values corresponding to the each of the first plurality of frames and the adjacent frame;

one or more instructions that programs the processor to process information associated with the plurality of difference values;

one or more instructions that programs the processor to determine a plurality of ranks corresponding to the first plurality of frames respectively based on at least information associated with the plurality of difference values, the plurality of ranks including a rank corresponding to the frame, the rank being a temporal signature for the frame; and

one or more instructions that programs the processor to store the temporal signature for the frame in a fingerprint database separate from the video object, wherein the video object is not modified by the temporal signature.

31. A non-transitory computer readable medium including instructions that programs a processor a computer system to generate a spatial-temporal signature for a frame of a video object, the non-transitory computer readable medium comprising:

one or more instructions that programs the processor to obtain a first plurality of frames associated with a video object, the first plurality of frames including at least a frame, each of the first plurality of frames corresponding to an adjacent frame;

one or more instructions that programs the processor to dividing each of the first plurality of frames, wherein the divided frame includes a first plurality of blocks corresponding to a plurality of locations respectively, each of the first plurality of blocks corresponding to a second plurality of blocks including the each of the first plurality of blocks, the second plurality of blocks being associated with blocks on the first plurality of frames respectively;

24

one or more instructions that programs the processor to process information associated with the first plurality of frames;

one or more instructions, for each of the first plurality of blocks,

that programs the processor to determine a plurality of difference values for the second plurality of blocks respectively, each of the plurality of difference values being associated with the each of the second plurality of blocks and a corresponding block on the adjacent frame; that programs the processor to process information associated with the plurality of difference values;

that programs the processor to determine a first plurality of ranks corresponding to the second plurality of blocks respectively based on at least information associated with the plurality of difference values;

that programs the processor to process information associated with the first plurality of ranks;

that programs the processor to determine a rank for the each of the first plurality of blocks based on at least information associated with the first plurality of ranks;

one or more instructions for processing information associated with a second plurality of ranks corresponding to the first plurality of blocks respectively, the second plurality of ranks including the rank;

one or more instructions that programs the processor to determine a sequence of ranks based on at least information associated with the second plurality of ranks and the plurality of locations for the first plurality of blocks; and

one or more instructions that programs the processor to store the spatial-temporal signature for the frame in a fingerprint database separate from the video object, wherein the spatial-temporal signature for the frame includes information associated with the sequence of ranks, and wherein the video object is not modified by the spatial-temporal signature.

* * * * *

2010 11 16

To: Mr. Patrick Owens
Senior Analyst
Costing and Competitive Services
Canadian Radio-television and
Telecommunications Commission
Ottawa, Ontario
K1A 0N2

Subject: **Complaint regarding Bell Canada's Internet traffic management practices**

Dear Mr. Owens,

s.19(1)

1. Bell Canada (or the Company) is in receipt of a Commission staff letter dated 27 October 2010, indicating that Commission staff received a complaint regarding Internet traffic management practices (ITMPs) applied by the Company. The complaint by [REDACTED] attached to the Commission staff letter alleges that the Company has "decided to throttle (i.e. lower the bandwidth) on websites beyond only P2P (peer to peer) traffic as stated on the Company's "Network Management" web page." More particularly, Mr. Singh states that downloads from the website www.hotfile.com are being "throttled" and requests that the Company be required to specify which websites it "throttles" pursuant to its ITMPs. The following constitutes the Company's response.

2. As the Company has consistently stated, it does not intentionally shape any non-P2P file-sharing traffic as part of its network management practices and, as such, there are no websites that the Company "targets" for traffic shaping. However, as the Company has stated in the past, the deployment of any network technology such as technical ITMPs may have some unintended consequences and the Company encourages its customers to bring such unintended consequences to its attention in order to resolve them.¹

¹ See Bell Canada(CRTC)15May08-6 CAIP Part VII and The Companies(CRTC)4Dec08-15 PN 2008-19.

Bell Canada
Suzanne Morin
Floor 19
160 Elgin Street
Ottawa, Ontario K2P 2C4

Telephone: (613) 785-6282
Facsimile: (613) 560-0472
bell.regulatory@bell.ca

3. As a result of the complaint letter, the Company has performed various tests in an attempt to recreate the reported problem. As accurately indicated in the complaint letter, www.hotfile.com is a storage type of site that leverages http (web) direct downloading mechanisms for the sharing of files. The site offers two types of access, a free access and a premium (paid for subscription) access. Whereas the free access has limitations in terms of the number of concurrent downloads as well as throughput/speed that can be reduced if the site www.hotfile.com is congested, the premium service is not limited in concurrent downloads or throughput/speed as a result of congestion at the site itself.

4. The Company's testing has revealed that in some instances, and for both free and premium hotfile accounts, certain non-P2P www.hotfile.com traffic was indeed being identified by the Company's deep packet inspection (DPI) devices as being P2P file-sharing protocols. Further investigation allowed the Company to determine that www.hotfile.com traffic was being classified by one of the signatures used by the Company's DPI devices as a P2P file-sharing application and as a result traffic from both free and premium accounts were being shaped as P2P file sharing traffic. The Company notes that the inaccurate identification of www.hotfile.com traffic is not the result of any changes to the signatures used by the Company.

5. In order to address the issue of inadvertent shaping of www.hotfile.com traffic, there was a requirement to update the signatures used in the various DPI platforms used by the Company. A solution was implemented for the vast majority of the Company's DPI devices on 2 November 2010 and a second solution for the remaining DPI devices will be implemented by 30 November 2010. As such, www.hotfile.com traffic should no longer be shaped during peak hours after the 30 November 2010.

6. Based on the above, the Company considers that it was and remains in compliance with the ITMP framework established in Telecom Regulatory Policy 2009-657. The Company would like to thank the Commission and [REDACTED] for bringing this issue to the Company's attention and continues to encourage its customers to report any problems in order to allow the Company to investigate and resolve any unintended consequences of its network management practices.

Yours truly,

s.19(1)

[Original signed by Suzanne Morin]

Suzanne Morin

Assistant General Counsel - Regulatory Law & Policy

c.c.: Lynne Fancy (CRTC)
[REDACTED]

PG/lp

*** End of Document ***

A FILE TRANSFER PROTOCOL

I. Introduction

Computer network usage may be divided into two broad categories -- direct and indirect. Direct usage implies that you, the network user, are "logged" into a remote host and use it as a local user. You interact with the remote system via a terminal (teletypewriter, graphics console) or a computer. Differences in terminal characteristics are handled by host system programs, in accordance with standard protocols (such as TELNET (RFC 97) for teletypewriter communications, NETRJS (RFC 88) for remote job entry). You, however, have to know the different conventions of remote systems, in order to use them.

Indirect usage, by contrast, does not require that you explicitly log into a remote system or even know how to "use" the remote system. An intermediate process makes most of the differences in commands and conventions invisible to you. For example, you need only know a standard set of network file transfer commands for your local system in order to utilize remote file system. This assumes the existence of a network file transfer process at each host cooperating via a common protocol.

Indirect use is not limited to file transfers. It may include execution of programs in remote hosts and the transfer of core images. The extended file transfer protocol would facilitate the exchange of programs and data between computers, the use of storage and file handling capabilities of other computers (possibly including the trillion-bit store data computer), and have programs in remote hosts operate on your input and return an output.

The protocol described herein has been developed for immediate implementation on two hosts at MIT, the GE645/Multics and the PDP-10/DM/CG-ITS (and possibly Harvard's PDP-10). An interim version with limited capabilities is currently in the debugging stage. [1] Since our implementation involves two dissimilar systems (Multics is a "service" system, ITS is not) with different file systems (Multics provides elaborate access controls, ITS provides none), we feel that the file transfer mechanisms proposed are generalizable. In addition, our specification reflects a consideration of other file systems on the network. We conducted a survey [2] of network host

systems to determine the requirements and capabilities. This paper is a "first cut" at a protocol that will allow users at any host on the network to use the file system of every cooperating host.

II. Discussion

A few definitions are in order before the discussion of the protocol. A file is an ordered set consisting of computer instructions and/or data. A file can be of arbitrary length [3]. A named file is uniquely identified in a system by its file name and directory name. The directory name may be the name of a physical directory or it may be the name of a physical device. An example of physical directory name is owner's project-programmer number and an example of physical device name is tape number.

A file may or may not have access controls associated with it. The access controls designate the users' access privileges. In the absence of access controls, the files cannot be protected from accidental or unauthorized usage.

A principal objective of the protocol is to promote the indirect use of computers on the network. Therefore, the user or his program should have a simple and uniform interface to the file systems on the network and be shielded from the variations in file and storage systems of different host computers. This is achieved by the existence of a standard protocol in each host.

Criteria by which a user-level protocol may be judged were described by Mealy in RFC 91, as involving the notion of logical records, ability to access files without program modifications, and implementability. I would add to these efficiency, extendibility, adaptability, and provision of error-recovery mechanisms.

The attempt in this specification has been to enable the reliable transfer of network ASCII (7-bit ASCII in 8-bit field with leftmost bit zero) as well as "binary" data files with relative ease. The use of other character codes, such as EBCDIC, and variously formatted data (decimal, octal, ASCII characters packed differently) is facilitated by inclusion of data type in descriptor headings. An alternative mechanism for defining data is also available in the form of attributes in file headings. The format control characters reserved for the syntax of this protocol have identical code representation in ASCII and EBCDIC. (These character are SOH, STX, ETX, DC1, DC2, DC3, US, RS, GS, and FS.)

The notion of messages (the physical blocks of data communicated between NCP's) is suppressed herein and that of "logical" records and transactions is emphasized. The data passed by the NCP is parsed into logical blocks by use of simple descriptors (code and count mechanisms) as described in Section III. The alternative to count is fixed length blocks or standard end-of-file characters (scan data stream). Both seem less desirable than count.

The cooperating processes may be "daemon" processes which "listen" to agreed-upon sockets, and follow the initial connection protocol much in the same way as a "logger" does. We recommend using a single full-duplex connection for the exchange of both data and control information [4], and using CLS to achieve synchronization when necessary (a CLS is not transmitted until a RFNM is received).

The user may be identified by having the using process send at the start of the connection the user's name information (either passed on by user or known to the using system) [5]. This user name information (a sequence of standard ASCII characters), along with the host number (known to the NCP), positively identifies the user to the serving process.

At present, more elaborate access control mechanisms, such as passwords, are not suggested. The user, however, will have the security and protection provided by the serving system. The serving host, if it has access controls, can prevent unprivileged access by users from other host sites. It is up to the using host to prevent its own users from violating access rules.

The files in a file system are identified by a pathname, similar to the labels described in RFC 76 (Bouknight, Madden, and Grossman). The pathname contains the essential information regarding the storage and retrieval of data.

In order to facilitate use, default options should be provided. For example, the main file directory on disk would be the default on the PDP-10/ITS, and a pool directory would be the default on Multics.

The file to be transferred may be a complete file or may consist of smaller records. It may or may not have a heading. A heading should contain ASCII or EBCDIC characters defining file attributes. The file attributes could be some simple agreed-upon types or they could be described in a data reconfiguration or interpretation language similar to that described in RFC 83 (Anderson, Haslern, and Heffner), or a combination.

The protocol does not restrict the nature of data in the file. For example, a file could contain ASCII text, binary core image, graphics data or any other type of data. The protocol includes an "execute" request for files that are programs. This is intended to facilitate the execution of programs and subroutines in remote host computers [6].


III. SPECIFICATIONS

1. Transactions

1A. The protocol is transaction-oriented. A transaction is defined to be an entity of information communicated between cooperating processes.

1B. Syntax

A transaction has three fields, a 72-bit descriptor field and variable length (including zero) data and filler fields, as shown below. The total length of a transaction is (72 + data + filler) bits.



1C. Semantics

The code field has three 8-bit bytes. The first byte is interpreted as transaction type, the second byte as data type and the third byte as extension of data type.

The filler count is a binary count of bits used as "filler" (i.e., not information) at the end of a transaction [7]. As the length of the filler count field is 8-bits, the number of bits of filler shall not exceed 255 bits.

The data count is a binary count of the number of data (i.e., information) bits in the data field, not including filler bits. The number of data bits is limited to $(2^{24}-1)$, as there are 24 bits in the data count field.

The NUL bytes are inserted primarily as fillers in the descriptor field and allow the count information to appear at convenient word boundaries for different word length machines [8].

2. Transaction Types

- 2A. A transaction may be of the following four basic types: request, response, transfer and terminate. Although large number of request and transfer types are defined, implementation of a subset is specifically permitted. Host computers, on which a particular transaction type is not implemented, may refuse to accept that transaction by responding with an unsuccessful terminate.

The following transaction type codes are tentatively defined:

Transaction Type	Transaction Type Code		
	ASCII	Octal	Hexidecimal
Request			
Identify	I	111	49
Retrieve	R	122	52
Store	S	123	53
Append	A	101	41
Delete	D	104	44
Rename	N	116	4E
addname (Plus)	P	120	50
deletename (Minus)	M	115	4D
Lookup	L	114	4C
Open	O	117	4F
Close	C	103	43
Execute [9]	E	105	45
Response			
ready-to-receive (rr)	<	074	3C
ready-to-send (rs)	>	076	3E
Transfer			
complete_file	*	052	
heading	#	043	23
part_of_file	'	054	2C
last_part	.	056	2E
Terminate			
successful (pos.)	+	053	2B
unsuccessful (neg.)	-	055	2D

2B. Syntax

In the following discussion US, RS, GS, FS, DC1, DC2, and DC3 are the ASCII characters, unit separator (octal 037), record separator (octal 036), group separator (octal 035), file separator (octal 034), device control 1 (octal 021), device control 2 (octal 022), and device control 3 (octal 023), respectively. These have an identical interpretation in EBCDIC.

2B.1 Requests

Identify, retrieve, store, append, delete, open, lookup and execute requests have the following data field:

<path name>

Rename request has the data field:

<path name> GS <name>

Addname and deletename requests have the data field:

<path name> GS <filenames>

where pathname [10], name and filenames have the following syntax (expressed in BNF, the metalanguage of the ALGOL 60 report):

<pathname> ::= <device name> | <name> | <pathname> US <name>
<device name> ::= DC1 <name>

<name> ::= <char> | <name> <char>
<char> ::= All 8-bit ASCII or EBCDIC characters except

US, RS, GS, FS, DC1, DC2, AND DC3.

<filenames> ::= <name> | <filenames> RS <name>

The data type for the request transaction shall be either A (octal 101 for ASCII, or E (octal 105) for EBCDIC [11].

Some examples of pathname are:

DC1 MT08
DC1 DSK 1.2 US Net<3> US J.Doe US Foo
udd US proj. US h,n/x US user US file
filename 1 filename 2

2B.2 Responses

The response transactions shall normally have an empty data field.

2B.3 Transfers

The data types defined in section 4 will govern the syntax of the data field in transfer transactions. No other syntactical restrictions exist.

2B.4 Terminates

The successful terminate shall normally have an empty data field. The unsuccessful terminate may have a data field defined by the data types A (octal 101) for ASCII, E (octal 105) for EBCDIC, or S (octal 123) for status.

A data type code of 'S' would imply byte oriented error return status codes in the data field. The following error return status codes are defined tentatively:

Error Code Meaning	Error Code		
	ASCII	Octal	Hexadecimal
Undefined error	U	125	55
Transaction type error	T	124	54
Syntax error	S	123	53
File search failed	F	106	46
Data type error	D	104	44
Access denied	A	101	41
Improper transaction sequence	I	111	49
Time-out error	O	117	4F
Error condition by system	E	105	45

2C. Semantics

2C.1 Requests

Requests are always sent by using host. In absence of a device name or complete pathname, default options should be provided for all types of requests.

Identify request identifies the user as indicated by <pathname> from serving to using host.

Retrieve request achieves the transfer of file specified in <pathname> from serving to using host.

Store request achieves the transfer of file specified in <pathname> from using to serving host.

Append request causes data to be added to file specified in pathname.

Rename request causes name of file specified in <pathname> to be replaced by name specified in <name>.

Delete request causes file specified in <pathname> to be deleted. If an extra level of protection for delete is desired (such as the query 'Do you wish to delete file x?'), it is to be a local implementation option.

Addname and _deletename_ requests cause names in <filenames> to be added or deleted to existing names of file specified in <pathname>. These requests are useful in systems such as Multics which allow multiple names to be associated with a file.

Lookup request achieves the transfer of attributes (such as date last modified, access list, etc) of file specified in <pathname>, instead of the file itself.

Open request does not cause a data transfer, instead file specified in <pathname> is "opened" for retrieve (read) or store (write). Subsequent requests are then treated as requests pertaining to the file that is opened till such a time that a close request is received.

Execute request achieves the execution of file specified in <pathname>, which must be an executable program. Upon receipt of rr response, using host will transmit the necessary input data (parameters, arguments, etc). Upon completion of execution serving host will send the results to using host and terminate [12].

2C.2 Response

Responses are always sent by serving host. The rr response indicates that serving host is ready to receive the file indicated in the preceding request. The rs response indicates that the next transaction from serving host will be the transfer of file indicated in the preceding request.

2C.3 Transfers

Transfers may be sent by either host. Transfer transactions indicate the transfer of file indicated by a request. Files can be transferred either as `complete_file` transactions or as `part_of_file` transactions followed by `last_part` transactions. The file may also have a heading transaction in the beginning. The syntax of a file, therefore, may be defined as:

```
<file> ::= <text> | <heading> <text>
<text> ::= <complete_file> | <parts> <last_part>
<parts> ::= <part_of_file> | <parts> <part_of_file>
```

Headings may be used to communicate the attributes of files. The form of headings is not formally specified but is discussed in Section IV as possible extension to this protocol.

2C.4 Terminates

The successful terminate is always sent by serving host. It indicates to using host that serving host has been successful in serving the request and has gone to an initial state. Using host will then inform user that his request is successfully served, and go to an initial state.

The unsuccessful terminate may be sent by either host. It indicates that sender of the terminate is unable to (or does not wish to) go through with the request. Both hosts will then go to their initial states. The using host will inform the user that his request was aborted. If any reasons for the unsuccessful terminate (either as text or as error return status codes) are received, these shall be communicated to the user.

3. Transaction Sequence

3A. The transaction sequence may be defined as an instance of file transfer, initiated by a request and ended by a terminate [13]. The exact sequence in which transactions occur depends on the type of request. A transaction sequence may be aborted anytime by either host, as explained in Section 3C.

3B. Examples

The identify request doesn't require a response or terminate and constitutes a transaction sequence by itself.

Rename, delete, addname, deletename and open requests involve no data transfer but require terminates. The user sends the request and the server sends a successful or an unsuccessful terminate depending on whether or not he is successful in complying with the request.

Retrieve and Lookup requests involve data transfer from the server to the user. The user sends the request, the server responds with a rs, and transfers the data specified by the request. Upon completion of the data transfer, the server terminates the transaction sequence with a successful terminate if all goes well, or with an unsuccessful terminate if errors were detected.

Store and Append requests involve data transfer from the user to server. The user sends the request and the server responds with a rr. The user then transfers the data. Upon receiving the data, the server terminates the sequence.

Execute request involves transfer of inputs from user to server, and transfer of outputs from server to user. The user sends the request to which the server responds with rr. The user then transfers the necessary inputs. The server "executes" the program or subroutine and transfers the outputs to the user. Upon completion of the output transfer, the server terminates the transaction sequence.

3C. Aborts

Either host may abort the transaction sequence at any time by sending an unsuccessful terminate, or by closing the connection (NCP to transmit a CLS for the connection). The CLS is a more drastic type of abort and shall be used when there is a catastrophic failure or when an abort is desired in the middle of a long file transfer. The abort indicates to the receiving host that the other host wishes to terminate the transaction sequence and is now in the initial state. When CLS is used to abort, the using host will reopen the connection.

4. Data Types

- 4A. The data type code together with the extension code defines the manner in which the data field is to be parsed and interpreted [14]. Although a large number of data types are defined, specific implementations may handle only a limited subset of data types. It is recommended that all host sites accept the

"network ASCII" and "binary" data types. Host computers which do not "recognize" particular data types may abort the transaction sequence and return a data type error status code.

- 4B. The following data types are tentatively defined. The code in the type and extension field is represented by its ASCII equivalent with 8th bit as zero.

Data Type	Byte Size	Code Type	Extension
ASCII character, bit8=0 (network)	8	A	NUL
ASCII characters, bit8=1	8	A	1
ASCII characters, bit8=even parity	8	A	E
ASCII characters, bit8=odd parity	8	A	O
ASCII characters, 8th bit info.	8	A	8
ASCII characters, 7 bits	7	A	7
ASCII characters, in 9-bit field	9	A	9
ASCII formatted files (with SOH, STX, ETX, etc.)	8	A	F
DEC-packed ASCII (5 7-bit char., 36th bit 1 or 0)	36	A	D
EBCDIC characters	8	E	NUL
SIXBIT characters	6	S	NUL
Binary data	1	B	NUL
Binary bytes (size is binary ext.)	1-255	B	(any)
Decimal numbers, net ASCII	8	D	A
Decimal numbers, EBCDIC	8	D	E
Decimal numbers, sixbit	6	D	S
Decimal numbers, BCD (binary coded)	4	D	B
Octal numbers, net. ASCII	8	O	A
Octal numbers, EBCDIC	8	O	E
Octal numbers, SIXBIT	6	O	S
Hexadecimal numbers, net. ASCII	8	H	A
Hexadecimal numbers, EBCDIC	8	H	E
Hexadecimal numbers, SIXBIT	6	H	S
Unsigned integers, binary (ext. field is byte size)	1-225	U	(any)
Sign magnitude integers (field is binary size)	1-255	I	(any)
2's complement integers (ext. field is byte size)	1-255	2	(any)
1's complement integers (ext. field is byte size)	1-255	1	(any)
Floating point (IBM360)	32	F	I
Floating point (PDP-10)	36	F	D
Status codes	8	S	NUL

- 4C. The data type information is intended to be interpretive. If a host accepts a data type, it can interpret it to a form suited to its internal representation of characters or numbers [15]. Specifically when no conversion is to be performed, the data type used will be binary. The implicit or explicit byte size is useful as it facilitates storing of data. For example, if a PDP-10 receives data types A, A1, AE, or A7, it can store the ASCII characters five to a word (DEC-packed ASCII). If the datatype is A8 or A9, it would store the characters four to a word. Sixbit characters would be stored six to a word. If conversion routines are available on a system, the use of system program could convert the data from one form to another (such as EBCDIC to ASCII, IBM floating point to DEC floating point, Decimal ASCII to integers, etc.).

5. Initial Connection, CLS, and Identifying Users

- 5A. There will be a prearranged socket number [16] for the cooperating process on the serving host. The connection establishment will be in accordance with the initial connection protocol of RFC 66 as modified by RFC 80. The NCP dialog would be:

user to server: RTS<us><3><p>

if accepted, server to user: STR<3><us><CLS><3><us>
server to user on link p: <ss>
server to user: STR<ss+1><us>RTS<ss><us+1><q>
user to server: STR<us><ss+1>RTS<us+1><ss><r>

This sets up a full-duplex connection between user and server processes, with server receiving through local socket ss from remote socket us+1 via link q, and sending to remote socket us through local socket ss+1 via link r.

- 5B. The connection will be broken by trading a CLS between the NCP'S for each of the two connections. Normally the user will initiate the CLS.

CLS may also be used by either the user or the server to abort a data transmission in the middle. If a CLS is received in the middle of a transaction sequence, the whole transaction sequence will be aborted. The using host will then reopen the connection.

- 5C. The first transaction from the user to server will be the identify transaction. The users will be identified by the pathname in data field of the transaction which should be a

form acceptable to the server. The server is at liberty to truncate pathnames for its own use. Since the identify transaction does not require a response or terminate, the user can proceed directly with other requests.

IV. Extensions to Protocol

The protocol specified above has been designed to be extendable. The obvious extensions would be in the area of transaction types (new types of requests), error return status words, and data types. Some of the non-obvious extensions, that I can visualize are provisions of access control mechanisms, developing a uniform way of specifying file attributes in headings of files, increasing the scope of the execute command to include subroutine mediation, and the provision of transaction sequence identification numbers to facilitate handling of multiple requests over the same connection pair.

Users of protected file systems should be able to have a reasonable degree of confidence in the ability of the serving process to identify remote users correctly. In the absence of such confidence, some users would not be willing to give access to the serving process (especially write access). Inclusion of access control mechanisms such as passwords, is likely to enhance the indirect use of network by users who are concerned about privacy and security. A simple extension to the protocol would be to have the serving host sent a transaction type "password?" after it receives user name. Upon receipt of "password?" the using host will transmit the password, which when successfully acknowledged, would indicate to the user that requests may proceed.

There are a number of file attributes which properly belong in the heading of a file rather than the file itself or the data type in descriptors of transactions. Such attributes include access control lists, date file was last modified, information about the nature of file, and description of its contents in a data description or data reconfiguration language. Some uniformity in the way file attributes are specified would be useful. Until then, the interpretation of the heading would be up to the user or the using process. For example, the heading of files which are input to a data reconfiguration (form) machine may be the desired transformations expressed in the reconfiguration language.

The "execute" command which achieves the execution of programs resident in remote hosts is a vital part of indirect use of remote hosts. The present scope of the execute command, as outlined in the specifications, is somewhat limited. It assumes that the user or

using process is aware of the manner in which the arguments and results should be exchanged. One could broaden the scope of the execute command by introducing a program mediation protocol [17].

The present specification of the protocol does not allow the simultaneous transfer and processing of multiple requests over the same pair of connections. If such a capability is desired, there is an easy way to implement it which only involves a minor change. A transaction sequence identification number (TSid) could replace a NUL field in the descriptor of transactions. The TSid would facilitate the coordination of transactions, related to a particular transaction sequence. The 256 code combinations permitted by the TSid would be used in a round-robin manner (I can't see more than 256 outstanding requests between two user-processes in any practical implementation). An alternate way of simultaneous processing of requests is to open new pairs of connection. I am not sure as to how useful simultaneous processing of requests is, and which of the two is a more reasonable approach.

V. Conclusions

I tried to present a user-level protocol that will permit users and using programs to make indirect use of remote host computers. The protocol facilitates not only file system operations but also program execution in remote hosts. This is achieved by defining requests which are handled by cooperating processes. The transaction sequence orientation provides greater assurance and would facilitate error control. The notion of data types is introduced to facilitate the interpretation, reconfiguration and storage of simple and limited forms of data at individual host sites. The protocol is readily extendible.

Endnotes

[1] The interim version of the protocol, limited to transfer of ASCII files, was developed by Chander Ramchandani and Howard Brodie of Project MAC. The ideas of transactions, descriptors, error recovery, aborts, file headings and attributes, execution of programs, and use of data types, pathnames, and default mechanisms are new here. Howard Brodie and Neal Ryan have coded the interim protocol in the PDP-10 and the 645, respectively.

[2] The network system survey was conducted last fall by Howard Brodie of Project MAC, primarily by telephone.

[3] PDP-10 Reference Handbook, page 306.

[4] We considered using two full-duplex links, one for control information, the other for data. The use of a separate control link between the cooperating processes would simplify aborts, error recoveries and synchronization. The synchronization function may alternatively be performed by closing the connection (in the middle of a transaction sequence) and reopening it with an abort message. (The use of INR and INS transmitted via the NCP control link has problems as mentioned by Kalin in RFC 103.) We prefer the latter approach.

[5] Identifying users through use of socket numbers is not practical, as unique user identification numbers have not been implemented, and file systems identify users by name, not number.

[6] This subject is considered in detail by Bob Metcalfe in a forthcoming paper.

[7] Filler bits may be necessary as particular implementations of NCP's may not allow the free communication of bits. Instead the NCP's may only accept bytes, as suggested in RFC 102. The filler count is needed to determine the boundary between transactions.

[8] 72-bits in descriptor field are convenient as 72 is the least common multiple of 6, 8, 9, 18, 24 and 30, the commonly encountered byte sizes on the ARPA network host computers.

[9] The execute request is intended to facilitate the indirect execution of programs and subroutines. However, this request in its present form may have only limited use. A subroutine or program mediation protocol would be required for broader use of the execute feature. Metcalfe considers this problem in a forthcoming paper.

[10] The pathname idea used in Multics is similar to that of labels in RFC 76 by Bouknight, Madden and Grossman.

[11] We, however, urge the use of standard network ASCII.

[12] The exact manner in which the input and output are transmitted would depend on specific mediation conventions. Names of input and output files may be transmitted instead of data itself.

[13] The transactions (including terminate) are not "echoed", as echoing does not solve any "hung" conditions. Instead time-out mechanisms are recommended for avoiding hang-ups.

[14] The data type mechanism suggested here does not replace data reconfiguration service suggested by Harslem and Heafner in RFC 83 and NIC5772. In fact, it complements the reconfiguration. For

example, data reconfiguration language can be expressed in EBCDIC, Network ASCII or any other code that form machine may "recognize". Subsequent data may be transmitted binary, and the form machine would reconfigure it to the required form. I have included in data types, a large number suggested by Harslem and Heafner, as I do not wish to preclude interpretation, reconfiguration and storage of simple forms of data at individual host sites.

[15] The internal character representation in the hosts may be different even in ASCII. For example PDP-10 stores 7-bit characters, five per word with 36th bit as don't care, while Multics stores them four per word, right-justified in 9-bit fields.

[16] It seems that socket 1 has been assigned to logger and socket 5 to NETRJS. Socket 3 seems a reasonable choice for the file transfer process.

[17] The term program mediation was suggested by Bob Metcalfe who is intending to write a paper on this subject.

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Ryan Kato 6/01]

D. Exhibit X: Additional Materials Relied Upon



Hello, Andrew Cromarty. We have recommendations for you. (Not Andrew?)

The All-New Kindle Family: from \$79

Andrew's Amazon.com | Today's Deals | Gifts & Wish Lists | Gift Cards

Your Digital Items | Your Account | Help

Shop All Departments

Search Movies & TV Dragon Tattoo

GO Cart Your Lists

Movies & TV

New Releases Best Sellers Today's Deals Blu-ray Browse Genres TV Shows Amazon Instant Video Prime Instant Videos Your Video Library

Format

- Clear
- Amazon Instant Video
- Blu-ray (9)
- DVD (22)
- VHS (1)

Department

- < Any Department
- Movies & TV**
- Movies (6)

Amazon Prime

- Clear
- Prime Eligible

Purchase Type

- Purchase (3)
- Rental (3)

Genre

- Action & Adventure (3)

Movies & TV > Amazon Instant Video > Prime Eligible > "Dragon Tattoo"

Showing 6 Results

Sort by Relevance



1. The Girl With the Dragon Tattoo (English dubbed) Starring Michael Nyqvist, Noomi Rapace, Lena Endre and Peter Haber

★★★★☆ (56)

amazon instant video

Prime members: \$0.00 Prime (unlimited streaming) Buy: \$9.99

Watch instantly on your PC, Mac, compatible TV or device.



2. The Girl With the Dragon Tattoo (English subtitled) Starring Michael Nyqvist, Noomi Rapace, Lena Endre and Peter Haber

★★★★☆ (507)

amazon instant video

Prime members: \$0.00 Prime (unlimited streaming) Rent: \$2.99



6 March 2009

Internet Service Provider Copyright Code of Practice – TCF Consultation Draft

Google is grateful for the opportunity to present its comments in response to the Telecommunications Carriers' Forum's (TCF) Consultation Draft of the *Internet Service Provider Copyright Code of Practice* (the **Draft Code**). Google wishes to also take this opportunity to express its concerns regarding the recent amendments to the Copyright Act 1994 (the **Act**), in particular section 92A. In Part 1 of this submission, Google provides its views on copyright and its relevant international experience with copyright in an online environment.

In Part 2 of this submission, Google provides preliminary comments on a number of issues raised by the Draft Code. In a number of places in Part 2, Google provides recommendations to the TCF, as well as suggestions for alternative drafting of the Draft Code noted as underlined text. Google would also be able to provide a marked-up version of the Draft Code if this would assist the TCF in its consultation process on the Draft Code.

Google would be pleased to provide further information on either of these Parts.

Summary of Google's submission

In Part 1 of this submission, Google addresses the following key points:

- Google is strongly committed to preserving the Internet's fundamental openness and protecting the benefits to technology innovation and content creation that are gained from this openness.
- Google supports a flexible and adaptable legal and regulatory framework that aims to provide an effective balance between the interests of rights holders' copyright protection and users' self-expression, in order to promote an environment favourable to innovation,



creativity and diversity.

- Google has a number of concerns around the new section 92A and the impact the section 92A obligation is likely to have on the balance of interests served by copyright law:
 - » Section 92A undermines the incredible social and economic benefits of the open and universally accessible Internet, by providing for a remedy of account termination or disconnection that is disproportionate to the harm of copyright infringement online.
 - » Section 92A puts users' procedural and fundamental rights at risk, by threatening to terminate users' Internet access based on mere allegations and reverse the burden of proof onto a user to establish there was no infringement. In Google's experience, there are serious issues regarding the improper use and inaccuracy of copyright notices by rights holders.
 - » Section 92A could impose significant burdens on ISPs, as it threatens to require enforcement of policies based simply on rights holders' allegations of infringement.

In Part 2 of this submission, Google addresses the following key points:

- As the Draft Code is a voluntary industry code of practice, Google submits that the Draft Code should expressly acknowledge that becoming a Party is not necessary or relevant for compliance with section 92A.
- There are numerous types of ISPs falling within the scope of the Act, therefore a 'one size fits all' approach to section 92A compliance is impractical and could impose excessive cost and procedural requirements on ISPs. Google submits that the Draft Code should clarify and confirm that there may be a variety of ISP termination policies that meet the requirements of section 92A.



- User education is stated as a “primary purpose” of the Draft Code, though there is little reference to education about users’ rights, including limitations and exceptions enabling lawful use of copyright protected works, in addition to their obligations. Google submits that the Draft Code clarify the concept of ‘user education’ to include this information.
- The Draft Code provides for notification and termination based solely on allegations of infringement, without an independent evaluation of evidence and determination of infringement. Google submits that the Draft Code should provide for greater independence in the process, as well as including acknowledgement of the possible defences and exceptions to copyright infringement.
- Whether a Party is in compliance with the Draft Code does not indicate whether that Party is meeting its obligations under section 92A. Google submits that the Draft Code should expressly state this, and in this context should also remove the particularly severe and serious enforcement measures for non-compliance with the Draft Code.
- The Draft Code provides for ‘Processing Fees’, but these are not further detailed in the Draft Code itself. Google submits that the Draft Code should include additional details of any fees and the basis for those fees.



PART 1: Copyright on the Internet

Introduction to Google

Google initially became familiar to most Internet users as the provider of the Google search engine and subsequently as the provider of email, instant messaging and specialised search and information services, including Google News, Google Maps and Google Finance. Google is also the provider of the well known YouTube service, a platform for people to watch and share original videos through a web experience, which includes video content from New Zealand users featured on youtube.co.nz.

Google's breakthrough technology and continued innovation serve its mission of 'organising the world's information and making it universally accessible and useful to the public'. Google's business model has focused on what is known as the 'long tail' of the Internet – the millions of individuals and small businesses that cater to niche interests and markets. Google's services endeavour to democratise the means of accessing, creating and communicating information across local, national and global boundaries.

Open Internet

Google is strongly committed to preserving the Internet's fundamental openness and protecting the benefits to technology innovation and content creation that are gained from this openness. As the Internet becomes increasingly essential as the fundamental communications infrastructure, promoting and nourishing the Internet's openness is paramount. Google believes that the opportunities for freedom of expression, creativity and access to information are the greatest benefits that the Internet offers to consumers and, more broadly, to society.

Innovation and free expression have thrived in an online environment because the Internet's architecture enables any and all users to generate new ideas, content and technologies. The Internet has dramatically lowered the barriers for any individual to develop transformative



technologies and has created unparalleled avenues for social discourse across national and international boundaries. It is gaining a crucial role in the daily life of users and progressively leading to a major shift in users' habits, from passive media consumption to active content selection and creation. Along with facilitating access to information, entertainment content and public services, the Internet constitutes a central means of communication and self-expression in both personal and professional life (email, VoIP, blogs, instant messaging etc). In this way, the Internet is creating opportunities for a more participative society, and providing enormous potential for creativity and availability of new content.¹

Copyright legislation

Protection of copyright on the Internet

Copyright's core goal is to benefit the public by providing adequate incentives for the creation and dissemination of creative works, and it plays a number of extremely important roles: part cultural policy, part innovation policy, part consumer protection policy and part competition policy. In order for copyright to fully achieve its purpose, it must balance the interests of rights holders' copyright protection and users' self-expression, promoting an environment favourable to innovation, creativity and diversity. Google has always supported legislation and regulation that aims to provide an effective balance.

While inadequate copyright protection can reduce incentives to create, excessive copyright protection can stifle creativity, choke innovation, impoverish culture and block free and fair competition. As both an intermediary and an innovator in online technologies, Google supports a flexible and adaptable legal framework that provides those who create and invest in new technologies the freedom to innovate without fear that their efforts will be hindered by an overly restrictive approach to copyright. Copyright must have sufficient flexibility so that new, legitimate and socially desirable uses, enabled by new technologies, can flourish.

¹ Note also that the UN has recognized broadband as an essential utility for individuals, in the *United Nations Conference on Trade and Development Information Economy Report*, UNCTAD/SDTE/ECB/2006/1, Nov 2006.



Google considers that a focus of any legislative or regulatory framework should be to ensure that all stakeholders can realise the benefits of the Internet, and maximise the scope and potential for innovative content services to emerge out of New Zealand. Depriving users from access to the Internet will sharply limit efficiency gains and productivity improvements in the economy, encumber innovation, and undermine the social benefits of the Internet.

ISP liability provisions in the Copyright Act 1994

Google supports the Government's efforts to bring New Zealand's copyright framework up to date with recent developments and innovations in online technology. The recent amendments to the Act, which include protection for ISPs against liability for copyright infringement on the Internet, are issues that Google has experience with in other jurisdictions.

Google's relevant experience with copyright online

Google strongly supports the rights of both content creators and users. Google considers that the most effective way to improve access to online content and to address rights holders' interests in the online environment is through the development of innovative content services and tools.

Google is proactive in informing its users about copyright and infringing content online, while also working to assist rights holders in protecting and exploiting their copyright online. Google supports the promotion of consumer understanding of copyright, as evidenced by Google's well-developed policies and guidelines on copyright and copyright infringement, as well as the tips and articles for users in complying with copyright while using Google's products.²

Google works to notify users of their requirements to comply with copyright when uploading content. As stated in its copyright policy, Google requires users, when uploading content online,

² For example, see: YouTube 'Copyright Infringement Policy': <http://www.google.com/support/youtube/bin/answer.py?answer=55772&topic=13656>; and YouTube 'Copyright



to confirm that they either own the copyright to the content or are authorised to deal with the content. Google respects the rights of copyright holders and provides them with the opportunity to notify Google of an instance of copyright infringement. Google makes it clear to users that infringing content will be removed by Google once it becomes aware of any unauthorised use of copyright protected material.

Google has also created and continues to develop powerful tools that enable artists and creators to manage and control their copyright online. For example, Google has recently developed the YouTube Video Identification Tool, which identifies matches between user uploads and copyright protected material and gives rights holders the ability to maximise the use of their content according to their individual preferences.

Below, we provide an overview of some of the tools developed by Google for the YouTube service, which protect rights holders and inform users about the use of copyright protected works:

- YouTube offers an automated notification and take-down tool which enables rights holders to easily search for and identify YouTube videos that contain their content, and promptly remove them with the click of a mouse.
- YouTube also uses technology that creates unique identifiers of files that are removed from YouTube for reasons of copyright infringement, and prevents identical files from being uploaded to the YouTube website.
- YouTube informs its users about copyright and strongly discourages infringement (in a 'Copyright Tips' section of the YouTube website). There are clear and prominent messages concerning copyright ownership, notified to users at the point that users upload user-created content, as well as in YouTube's terms of use and via links to YouTube's copyright policies on every page of the YouTube website.

Tips': http://www.youtube.com/t/howto_copyright.



- YouTube also provides a feature called 'AudioSwap' which enables users to illustrate their original videos with music that YouTube licenses from music publishers and record labels. The purpose of this is to give YouTube users easily accessible options for being creative and using licensed music in their user-created content.
- YouTube Content Management Tools: YouTube provides technology to enable rights holders to produce unique identifiers of their copyright protected audiovisual works. If such works are then identified, then rights holders are empowered by being able to decide what should be done with this content, including by managing and monetising it.

New innovative tools such as the ones developed by Google for YouTube are offering a way for users to be creative, by gaining access to wide range of content for free, and for rights holders to protect their content, by identifying, managing access and monetising it. These types of tools and services can only be implemented if the regulatory framework allows for the necessary flexibility for new methods of cooperation between Internet intermediaries and rights holders to develop.

Section 92A of the Act

Google understands that the Government intends to bring section 92A into force on March 27, extended from February 28, and applauds the Government's recognition of the need for further discussion and consideration around the serious issues that section 92A raises. Google has a number of concerns around section 92A and the impact such an obligation is likely to have on innovation, free expression and the overall balance among the interests served by copyright law.

Google's concerns with section 92A

As discussed above, Google supports the promotion of universal access to the Internet, viewing the Internet as a vital part of individuals' and businesses' daily lives. It is critical to ensure that the legal and regulatory environment preserve the essential openness of the Internet.

Google believes that section 92A conflicts with the underlying principles of an open and



universally accessible Internet, by providing for a remedy that is disproportionate to the harm of copyright infringement online. The remedy of account termination or disconnection is especially serious and disproportionate in this context. This is particularly true when considering termination of a user's entire Internet access.

Google further believes that section 92A raises concerns regarding the protection of users' rights. Mere allegations of copyright infringement should not trump users' rights. Copyright law is often complex and context sensitive, and only a court is qualified to adjudicate allegations of copyright infringement. Indeed, in Google's experience, there are serious issues regarding the improper use and inaccuracy of copyright notices by rights holders.³ In this context, the responsibility should not fall to ISPs to determine cases of infringement.

Google believes that section 92A threatens to effectively reverse the burden of proof onto a user to establish that there was no infringement of copyright, which may put user's fundamental rights at risk. It also threatens to put substantial burdens on ISPs based on enforcement of such policies. This outcome may negatively affect innovation and deter users from creative expression on the Internet.⁴

³ A recent study undertaken in the United States reported on findings from takedown notices issued to Google under the Digital Millennium Copyright Act 1998 (US), concluding that over half (57%) of notices sent to Google for removal of material were sent by business targeting competitors and over one third (37%) of notices were not valid copyright claims. See J Urban & L Quilter, 'Efficient Process or "Chilling Effects"? Takedown Notices Under Section 512 of the Digital Millennium Copyright Act', http://mylaw.usc.edu/documents/512Rep-ExecSum_out.pdf.

⁴ These are a few examples of either mistaken or over-reaching notifications alleging copyright infringement online: "No Downtime For Free Speech Campaign", <http://www.eff.org/issues/ip-and-free-speech>; Letter to Google and YouTube from McCain/Palin 2008 campaign regarding takedown of political speech on YouTube, <http://www.eff.org/files/McCain%20YouTube%20copyright%20letter%2010.13.08.pdf>, October 13, 2008; "Viacom: Fair use is what we say it is", *Wired Blog Network*, <http://blog.wired.com/monkeybites/2007/08/viacom-fair-use.html>, August 31, 2007; "Boy dupes YouTube to delete videos", *The Sydney Morning Herald*, April 14, 2007.



PART 2: TCF's Consultation Draft of the ISP Copyright Code of Practice

Introduction

While Google has serious concerns with section 92A, we support effective stakeholder cooperation and the development of industry self-regulation and co-regulation to improve the respect of copyright in the online environment. In this context, Google supports the TCF's efforts, in light of the anticipated section 92A amendment coming into force, to strike a balance between copyright holders' and users' interests in developing the Draft Code.

Google's involvement in the TCF process

The TCF has developed the Draft Code in order to assist ISPs in complying with their obligations under the new section 92A. As Google is likely to fall within the definition of an 'Internet service provider' under the Act, Google will have the option of becoming a party to the Draft Code and therefore takes this opportunity to submit its comments on the Draft Code.

Timing for the development of the Draft Code

Google notes the Government's recent decision to delay bringing section 92A into force until March 27, in place of the original date of February 28. Google supports this important acknowledgement of the serious implications of section 92A and of the need for further discussion and consideration of these issues by the industry and the Government.

This extended time period therefore provides the TCF and industry stakeholders with a significant opportunity to play an even larger role in determining the details and process around section 92A through the Draft Code. It is imperative that the Draft Code strikes the appropriate balance between the interests represented in copyright law.

Purpose and focus of the Draft Code

The purpose of the Draft Code is to assist ISPs in section 92A compliance; however, it appears to



be directed towards only those ISPs acting as ‘Internet *access* providers’, rather than the all-encompassing definition of ‘Internet service provider’ set out in the Act. This interpretation is evidenced both through the purpose and principles sections of the Draft Code (Parts A and C) and the detailed processes set out in the Draft Code. For example, in clause 4.5.1, the drafting refers to “the ISP that directly *provides and controls* the Internet Account of a User”.

As a content host and provider of email, search and information services, Google does not fall neatly within the subset of ISPs that the Draft Code appears to be focused on. Instead of providing access to the Internet itself, Google acts as an intermediary to facilitate access for its users to the vast amount of content and information available on the Internet.

Voluntary nature of the Draft Code

The consequences of a narrower focus in the Draft Code may lead some ISPs to choose to not become a party to the Draft Code. Google notes that the Draft Code is a voluntary industry code of practice. The Draft Code itself acknowledges this to a limited extent in Part D (for example, in clause 5).

However, there remains potential for the Draft Code to form an ‘industry standard’ for compliance with section 92A. Although industry-wide self-regulation or co-regulation is often preferred, it is important to acknowledge that in the development of a code of practice for compliance with section 92A, not all ISPs can be dealt with under the same overall framework. Further, a number of ISPs may already be meeting their section 92A obligations through existing policies and Internet terms of use.

Recommendation: The Draft Code should acknowledge that complying with the Draft Code is not necessary or relevant for compliance with 92A. In addition, some ISPs may have already developed a ‘termination policy’ and may therefore elect not to sign up to the voluntary Draft Code.



Multiple options for a 'termination policy' under section 92A

Google submits that there is no single 'termination policy' which is appropriate for the numerous different types of ISPs falling within the definition in the Act. A 'one size fits all' approach to section 92A compliance may simply be impractical, imposing additional, perhaps excessive, cost and procedural requirements on ISPs and negatively impacting on their relationships with users and with rights holders.

A single, narrow approach would also encumber many industry participants in a way that discourages further innovation and stifles growth, as ISPs would be diverted from innovation and development towards seeking to comply with strictly defined provisions of an industry code. The lack of flexibility in such a narrow approach may also hamper the development of innovative business models that benefit rights holders, users and ISPs alike.

Part 1 explained Google's ongoing efforts in developing online tools, as well as implementing policies and terms of use, as evidence of Google's existing approach to copyright and infringement in an online environment. It is possible that, in addition to Google, a number of ISPs are also already in compliance with section 92A. This situation reinforces Google's view that there is no single 'termination policy' that is appropriate for all industry players in complying with section 92A.

Recommendation: Because it has the potential to be viewed as a default industry standard, the Draft Code should clarify and confirm that there may be a variety of ISP termination policies that meet the requirements of section 92A.

Education of users

The Draft Code includes recognition of the need for continued user education with respect to copyright. This is stated clearly both in the purpose (clause 1.3) and the principles (clause 4.5) sections, yet receives little attention elsewhere in the Draft Code. This is an area on which Google places great importance and dedicates significant resources towards the development of



online tools to educate and assist users, as described in Part 1. Google believes that the Draft Code must clarify that users must be fully informed about their obligations, but also about their rights under copyright, including limitations and exceptions that enable lawful use of copyright protected works.

Recommendation: Clarify the concept of ‘user education’ to include information about users’ rights, as well as limitations and exceptions that enable lawful use of copyrighted materials.

Suggested wording for the Draft Code

Google submits that the Draft Code should expressly acknowledge the varied ways in which it may be appropriate for different ISPs to comply with their section 92A obligations. In this case Google recommends that the TCF recognise, and reflect in the wording of the Draft Code, that there may be a diverse range of termination policies – the voluntary Draft Code is only one of these many options – for compliance with section 92A.

In particular, Google submits that:

(i) clause 1 be amended to the following:

- 1.1 “provide a reasonable policy and process for Parties, as one option for compliance with section 92A...”
- 1.3 “assist Copyright Holders to educate internet Users and Downstream ISPs as to their rights and obligations with respect to copyright, including any limitations and exceptions that allow lawful uses of copyright protected works.”
- 1.4 “provide a policy and process to enable Parties to operate a fair system...”

(ii) that clause 1.5 be deleted, as the Draft Code should only seek to address the obligations and procedural responsibilities of those ISPs that are Parties to the Draft Code.



(iii) clause 4 be amended to the following:

4.3 “All ISPs are required to comply with the Act (whether or not they are Parties to this Code). The fact that an ISP is or is not a Party to this Code does not, either directly or indirectly, indicate whether an ISP is complying with the Act. When a Party passes a Copyright Holder Notice to a Downstream ISP, it therefore does so in reliance on such compliance and in particular in reliance on that Downstream ISP having and implementing a termination policy complying with section 92A of the Act, whether that termination policy is in the form of the policy and process set out in this Code or is an alternative policy to this Code.”

4.5 “A primary purpose of this Code is to assist in the education of Users and Downstream ISPs as to their rights and obligations with respect to copyright, including any limitations and exceptions that allow lawful uses of copyright protected works.”

(iv) clause 5 be amended to the following:

5 “The Code is a voluntary code of practice that is applicable to those ISPs that have agreed in writing to be bound by it.”

(v) the definition of ‘Downstream ISP’ be confined only to those downstream ISPs who are Parties to the Draft Code and that clause 24 be amended accordingly.

(vi) in a number of clauses, in particular in Part C (‘Principles’), the term ‘ISP’ is sometimes used in place of ‘Party’, which is the appropriate term for the Draft Code (except in the case of clause 4.3).

(vii) clause 60 be amended to read:

60 “In order to promote understanding of users’ rights and obligations with



respect to copyright, including limitations and exceptions, and the effective operation of this Code...”

Lack of independence in the policy and process

Policy and process based on allegations alone

The Draft Code provides for a notification and termination process based solely on allegations of infringement of copyright made by rights holders. The process in the Draft Code following the notification stage allows for both the evaluation of evidence and the determination of infringement disputes by the Parties and Copyright Holders. These questions are more appropriate for an independent court to consider and finally determine and enforce.

Google is concerned that the Draft Code allows for termination without the exercise of an independent process for the consideration of allegations. This is particularly important as, in Google’s experience, there are serious issues regarding the improper use and inaccuracy of copyright notices by rights holders, which raise serious concerns with regard to the right to freedom of expression.⁵

Recommendation: Provide for greater independence in the evaluation and determination of allegations of copyright infringement, particularly in relation to evaluating any evidence of infringement, and establish greater safeguards against inaccurate and erroneous infringement notices.

Independent determination of copyright infringement

In any determination of copyright infringement, there are detailed and complex questions of fact and law to be determined by a qualified court. These questions include:

- evaluation of the admissibility of evidence;

⁵ See notes 3 and 4 above.



- consideration of any defences and exceptions to infringement provided for under the Act (Part 3), in accordance with the appropriate legal standard of proof;
- the extent of the protection of the right to freedom of expression; and
- determination of the reasonable and proportionate remedy or sanction where infringement is proven.

In contrast to these important roles of an independent court, the Draft Code generally sets out a process which:

- allows parties to an infringement dispute to determine evidential questions themselves, or requires a third party ISP to determine such evidential questions, with reference in the Draft Code to “acceptable” or “sufficient” evidence but without any provision for independent review;
- makes no reference to the possible existence of defences to infringement or exceptions for the use of copyright works, effectively placing the burden of proof on users to establish there was no infringement of copyright;
- makes no reference to the need to safeguard users’ fundamental rights, such as freedom of expression; and
- requires Parties to impose a remedy of ‘account termination’ which, in many circumstances, may be disproportionate and likely to go beyond what a court would order as reasonable.

***Recommendation:* Remove references in the Draft Code to areas that are more appropriately the responsibility of an independent court and include acknowledgement of the possible existence of defences and exceptions to copyright infringement.**

Suggested wording for the Draft Code

Google submits that the Draft Code should include greater acknowledgement of users’ rights,



including reference to possible defences or exceptions to copyright infringement. Google further submits that the TCF provide for greater independence and safeguards for users throughout the notification and termination process, reflecting the roles that are appropriately held by a court in evaluating and determining questions of copyright infringement.

Enforcement measures and cost provisions in the Draft Code

Non-compliance with the Act and the Draft Code

Every ISP's primary responsibility is to comply with its obligations under section 92A; the Draft Code is simply intended to assist ISPs in compliance with the Act. The Act provides for legislative remedies where an ISP fails to comply with section 92A; yet the Draft Code provides for further enforcement measures where a Party fails to comply with the Draft Code, in addition to any such measures found in the Act.

Google does not agree that additional enforcement measures are necessary in a voluntary industry code of practice. Where a Party is in breach of the Draft Code, this does not either directly or indirectly indicate that the Party is in breach of section 92A. Including such enforcement measures will not assist ISPs in complying with the Act, which is the central purpose for the Draft Code. In addition, Google is particularly concerned with the severity and seriousness of the three-stage enforcement measures set out in the Draft Code in light of the Draft Code's voluntary nature.

***Recommendation:* The enforcement measures in the Draft Code (clauses 42 to 56) should also be deleted and possibly replaced with provisions more appropriate for a voluntary code of practice.**

Comparison with non-compliance provisions in other industry codes of practice

Google notes that in similar TCF-prepared industry codes of practice, enforcement measures are not included within the code. For example, the ISP Spam Code of Practice (the **Spam Code**) was



developed as a co-regulatory code of practice under the Unsolicited Electronic Messages Act 2007, building on the legislative requirements without imposing any additional specific enforcement measures. Under the Spam Code, complaints regarding breaches of the Spam Code are to be considered by independent, established dispute resolution bodies such as the Disputes Tribunal or the Court.

Google notes that in other industries, common sanctions for breaches of industry-wide codes may be contrasted with the severe sanctions set out in the Draft Code. In the advertising sector, industry participants agree to comply with the self-regulatory advertising codes and standards established by the Advertising Standards Authority (ASA). Any person may lay a complaint with the Advertising Standards Complaints Board regarding a breach of a code or standard. The remedy for an upheld complaint is the withdrawal or amendment of the advertisement that is the subject of complaint.

Costs incurred in implementing the Draft Code

There are likely to be significant costs imposed on ISPs in achieving section 92A compliance, which will be increased through the implementation of the Draft Code. These costs will have an even greater impact on smaller ISPs. The Draft Code establishes a 'Processing Fee', to be set by the TCF and paid for by rights holders to ISPs. However, there is little additional information in the Draft Code as to the structure or basis for this fee.

Google is concerned with the lack of detail included in the Draft Code in relation to the costs of implementation of the termination policy by each ISP. Google considers it important that any fee is sufficient to cover the full costs of implementation. If the fee charged is inadequate to cover these costs, there could be an unwanted pass-through of costs onto users.

Recommendation: The Draft Code should have greater transparency, including additional detail as to the basis for any fees (e.g. the total cost of implementation) and the composition of the Processing Fee.



Suggested wording for the Draft Code

As noted above on page 14, Google submits that clause 4 be amended to the following:

- 4.3 “All ISPs are required to comply with the Act (whether or not they are Parties to this Code). The fact that an ISP is or is not a Party to this Code does not, either directly or indirectly, indicate whether an ISP is complying with the Act.”

Conclusion

The above submission sets out Google’s preliminary comments on a number of the issues raised by the TCF’s Draft ISP Copyright Code of Practice and section 92A of the Copyright Act 1994.

Google would be pleased to provide further details on any of the issues raised in this submission and looks forward to constructively engaging with the New Zealand Government, the TCF and other industry players as it continues to consider these important issues for the legal and regulatory framework of copyright in New Zealand.

Kind regards

Carolyn Dalton

Senior Policy Counsel
Public Policy and Government Relations
Australia & New Zealand
Google Sydney office
E: carolyndalton@google.com

Antoine Aubert

European Copyright Policy Counsel
Google Brussels office
E: aaubert@google.com

Contact: **Mark Toner**

Partner
MGF Webb
Auckland
T: 09 970 4108
E: mark.toner@mgfwebb.com

ABC | JULY 25, 2011

Harry Potter and the Order of the Phoenix to air on ABC



Christine Nyholm

Cable TV & Celebrity Examiner

[+ Subscribe](#)

[Add a comment](#)

[Email](#)

[Print](#)



Rupert Grint, Daniel Radcliffe and Emma Watson attend the New York premiere of "Harry Potter And The Deathly Hallows: Part 2" at Avery Fisher Hall, Lincoln Center on July 11, 2011 in New York City

ABC is presenting "Harry Potter and the Order of the Phoenix," as the "ABC Movie of the Week" on Saturday, August 13, 2011 (8:00 - 10:00 p.m. ET). The Harry Potter film stars Daniel Radcliffe, Emma Watson, Rupert Grint, Ralph Fiennes, Maggie Smith and Harry Melling. "Harry Potter and the Order of the Phoenix," winner of two People's Choice Awards, will be broadcast in HDTV format with 5.1-channel surround sound.

The fifth installation in the Harry Potter film series follows Harry Potter (Daniel Radcliffe) in his fifth year at Hogwarts after The Ministry of Magic refuses to believe that Lord Voldemort (Ralph Fiennes) has returned and appoints bureaucrat Dolores Umbridge (Imelda Staunton) as a teacher





Reference Manual on Scientific Evidence: Third Edition

ISBN
978-0-309-21421-6

1038 pages
6 x 9
PAPERBACK (2011)

Committee on the Development of the Third Edition of the Reference Manual on Scientific Evidence; Federal Judicial Center; National Research Council

 Add book to cart

 Find similar titles

 Share this PDF



Visit the National Academies Press online and register for...

- ✓ Instant access to free PDF downloads of titles from the
 - NATIONAL ACADEMY OF SCIENCES
 - NATIONAL ACADEMY OF ENGINEERING
 - INSTITUTE OF MEDICINE
 - NATIONAL RESEARCH COUNCIL
- ✓ 10% off print titles
- ✓ Custom notification of new releases in your field of interest
- ✓ Special offers and discounts

Distribution, posting, or copying of this PDF is strictly prohibited without written permission of the National Academies Press. Unless otherwise indicated, all materials in this PDF are copyrighted by the National Academy of Sciences. Request reprint permission for this book

Reference Manual on Scientific Evidence

Third Edition

Committee on the Development of the Third Edition of the
Reference Manual on Scientific Evidence

Committee on Science, Technology, and Law
Policy and Global Affairs

FEDERAL JUDICIAL CENTER

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

THE NATIONAL ACADEMIES PRESS
Washington, D.C.
www.nap.edu

[About Us](#)[Management Team](#)[Contact Us](#)[Press Room](#)[Affiliate Program](#)[Careers](#)[Privacy Policy](#)[Image Resources](#)[Trademark Policy](#)[DMCA Policy](#)

Roku Affiliate Program

The Roku Affiliate Program is a fun, easy, and free way to generate revenue by advertising Roku products and referring qualified customers to the roku.com website. Affiliate partners can earn up to 5% commission with every purchase.

With over a million players sold, Roku is the market leader in streaming entertainment devices for the TV, renowned for its simplicity, variety of entertainment choices, and exceptional value. (We encourage you to read the reviews yourself.)



How do I get started?

Raise your hand, jump up and down, and shout, "I'm ready!" Only kidding. You can join the Roku Affiliate Program through [Google Affiliate Network](#) or [Commission Junction Network](#). It almost runs itself!

How does it work?

We've kept it really simple. Here's a general overview:

- 5% commission per sale
- 7.5% Conversion Rate
- Affiliates may never use the URL www.roku.com or variations of it in its destination URL of their PPC campaign.
- Affiliates may not bid on Roku brand keywords.
- Affiliates are required to land PPC traffic on their own URL first.
- Affiliates may use the brand Roku in their ad text/description.
- Affiliates may not copy or directly mimic the ad copy being used by Roku for its own PPC campaigns, using only original content they create.

Click here to join Roku's Google Affiliate Network program:

[Apply Now](#)

Click here to join Roku's Commission Junction Affiliate Network program:

[Apply Now](#)

How does it work?

For more information on the Roku Affiliate program, contact us at:

support@rokuaffiliate.com

About Roku

Roku launched the first product designed to deliver movies from Netflix instantly on TV, using

Channels

Netflix
Hulu Plus

About Us

Products
Image Resources

Connect

Blog
Refer-A-Friend

Buy Roku

BUY NOW
RETAIL LOCATIONS

the power of the Internet. Since then, our best-selling Roku streaming players have become synonymous with amazing choice, control, and value in TV entertainment. Join us—and be a part of the Roku revolution.

[Amazon Instant Video](#)
[HBO GO](#)
[NBA Game Time](#)
[Pandora](#)
[See all Channels](#)

[Management](#)
[Jobs](#)
[Press](#)
[Contact](#)
[Affiliate Program](#)

[Facebook](#)
[Twitter](#)
[Create a Channel](#)
[Developers](#)
[Customer Reviews](#)

[ACCESSORIES](#)
[REVIEWS](#)
[ROKU COUPONS](#)
[Country Selector](#)



View the online version [here](#)

[Forward to a friend](#)

Introducing Refer-a-Friend

Refer. Rejoice. Repeat.



We would like to **reward you** for the little things you already do.

Simply **share a sweet deal** on a new Roku player with friends and family. For every player purchased, you get an Amazon Instant Video rental—on us.

The more you **share**, the more you get. In fact, we think some of you may never pay for another movie rental again!



Tell your friends ▶

Happy Streaming,
Your friends at Roku

Give

\$5 off

Get

\$4 credit

amazon instant video

The diagram illustrates the 'Give and Get' cycle. On the left, a Roku player is shown with a speech bubble saying '\$5 off'. On the right, an Amazon Instant Video monitor is shown with a speech bubble saying '\$4 credit'. The Roku player is labeled 'Give' and the Amazon monitor is labeled 'Get'.

[About Us](#) | [Privacy](#) | [Support](#) | [Products](#) | [Accessories](#) | [Unsubscribe](#) | [Forward to a Friend](#)

Share this with your friends: [Facebook](#) | [Twitter](#) | [Email](#)

© 2012 Roku Roku | 12980 Saratoga Avenue | Saratoga, CA 95070 | United States

This email was sent to r. [redacted] @d [redacted] by Roku


[Back to Walmart.com](#)
[Sign in](#) to the Affiliate Member Center

WALMART.COM AFFILIATE PROGRAM

GENERAL INFORMATION

[Home](#)
[Benefits](#)
[Newsletters & Calendars](#)
[FAQs](#)
[Contact Us](#)

AFFILIATE MEMBER CENTER

[Get Banners and Text Links](#)
[Get Data Feeds](#)
[Get Top Sellers](#)
[RSS Feeds](#)
[Link Generator](#)
[Manage Your Account](#)

Walmart.com Operating Agreement for Affiliate Network

This Walmart.com Operating Agreement for Affiliate Network, including all exhibits and attachments hereto and incorporated herein by reference (the "Agreement"), contains the complete terms and conditions that apply to a party's participation as an affiliate in the Affiliate Network of Walmart.com (the "Program"). As used in this Agreement, "we", "us", or "Walmart.com" means Wal-Mart.com USA, LLC (and its parent company and related entities), and "you" means the applicant party. "Site" means a World Wide Web site and, depending on the context, refers either to the Walmart.com Site or to the portion of the Walmart.com Site that you will link to using Qualifying Links as defined in Section 2 of this Agreement.

[Enrollment in the Program](#)
[Prohibitions](#)
[Referral Fees](#)
[Referral Fee Payment](#)
[Limited License; Restrictions](#)
[Term of the Agreement](#)
[Relationship of Parties](#)
[Limitation of Liability](#)
[Independent Investigation](#)
[Publicity](#)
[Remedies to Walmart.com](#)
[Exhibit A](#) - Trademark Requirements

[Exhibit C](#) - RSS Feeds

[Links on Your Website](#)
[Order Processing](#)
[Referral Fee Schedule](#)
[Policies and Pricing](#)
[Responsibility for Your Website](#)
[Modification](#)
[Indemnification](#)
[Disclaimers](#)
[Miscellaneous](#)
[Confidentiality](#)
[Exhibit B](#) - Sub-Affiliates or Networks

1. Enrollment in the Program

You may submit a completed Program application to begin the enrollment process ("Application"). Submission of your Application to the Walmart.com Affiliate Program implies acceptance to the terms set forth in this Agreement. We will evaluate your Application and notify you of your acceptance or rejection. We may reject your Application if, in our sole discretion, we determine for any reason that your website is unsuitable for the Program.

Wal-Mart Stores, Inc., Wal-Mart.com USA, LLC and Sam's West, Inc. associates, family members of associates, and suppliers are not eligible to enroll in the Program.

Unsuitable websites include, but are not limited to, those that:

- Promote sexually explicit material;
- Promote violence or hate toward any persons or groups;
- Promote illegal activities;
- Promote alcohol, tobacco, gambling/lottery in any way;
- Promote the use of pyramid, "ponzi", or similar investment schemes;
- Promote discrimination based on race, sex, religion, nationality, disability, sexual orientation, or age;
- Contain, in our sole judgment, material that is defamatory, fraudulent, or harassing to us or any third party;
- Are known as "blogging sites", defined for purposes of this Agreement as sites that contain only blogging and no other form of informational content;
- Include "walmart", "wal-mart" or variations or misspellings thereof in their domain names;
- Otherwise violate intellectual property rights of Walmart.com, Wal-Mart Stores or its suppliers;

- Disparage Walmart.com, Wal-Mart Stores, Wal-Mart Stores affiliates, or their suppliers;
- Are under construction or not live at the time of Application;
- Are non-US based or are websites that primarily serve a non-US based audience;
- Do not clearly state an online privacy policy to its visitors; or
- Provide a portion of their Referral Fees (as defined in Section 5) to websites or organizations that would violate any of the above criteria.

In addition, Walmart.com may, at its discretion, decline to accept, require adherence to an additional set of terms and conditions, or require the posting of specific copy for any website that (1) donates, directs or transfers any portion of their Referral Fees or affiliate benefits to any charitable website, education-related website, organization or program; or (2) provides rewards back to their members in the form of points or cash-back, or conduct other similar loyalty programs in connection with purchases made by members via their websites.

Regardless of your acceptance in the Program, we may terminate this Agreement for any reason, at any time.

The terms of our acceptance criteria are subject to change at any time without prior notice.

All decisions for acceptance into the Program will be made within our sole discretion.

If your Application is not accepted, you may reapply to the Program at any time; however, you should not and may not link to our Site unless you are approved for the Program.

As a member of the Affiliate Program, you grant Walmart.com permission to distribute any email communication directly to you that Walmart.com determines is necessary communication for you to receive in order to continue as a member of the Program, regardless of your choice to opt-out from certain communication.

[Top of Page](#)

2. Links on Your Website or Third Party Websites

Upon acceptance into the Program, we will make available to you Qualifying Links that are subject to the terms and conditions of this Agreement. A "Qualifying Link" is a link from a website to our Site using one of the Universal Record Locators ("URLs") or graphic links provided by Walmart.com, The LinkShare Network or by other means selected by us for use in the Program. All Qualifying Links must link directly and exclusively to Walmart.com. Walmart.com must approve each and every website that links to our Site through a Qualifying Link. If you use a Qualifying Link to link a website to our Site without seeking explicit authorization, your continued use of that Qualifying Link shall be considered a breach of this Agreement. However, continued use of the Qualifying Link will nonetheless subject such websites to the Terms and Conditions of this Agreement.

The Qualifying Links will serve to identify you website as a member of the Program and will establish a link from a website to our Site. All Qualifying Links that you will use in the Program will be provided to you from The LinkShare Network or by other means selected by us. You also agree that you will display on the website containing the Qualifying Link only those logos, trade names, trademarks, graphic images and similar identifying material ("Licensed Materials") that are provided by us or by the LinkShare Network, and you will substitute such images with any new materials provided by us or the LinkShare Network from time to time throughout the term of this Agreement. A web widget that is pre-approved in writing by Walmart.com for use on your website may be considered a Qualifying Link for purposes of this Agreement. Accordingly, web widgets are subject to all of the Terms and Conditions of this Agreement that apply to Qualifying Links.

Only valid Qualifying Links will be tracked for purposes of determining Referral Fees that you may be eligible to receive on sales of Qualifying Products (as defined below) generated through your participation in the Program.

Only Qualifying Links may be used to link a website to areas within our Site. You may not link directly to Walmart.com without use of a Qualifying Link. You may post as many Qualifying Links to our Site as you like on a website, provided that you ensure that each website containing a Qualifying Link posted by you meets the terms of this Agreement, including without limitation, that such website does not fall into the "unsuitable website" category described in Section 1, does not fall into the "prohibitions" set forth in Section 3, and you take responsibility for all websites on which you post a Qualifying Link in

accordance with Section 10. The position, prominence and nature of links on a website shall comply with any requirements specified in this Agreement but otherwise will be in your discretion.

You acknowledge that, by participating in the Program and placing a link to Walmart.com (or any category page therein) on any website through use of a Qualifying Link, we may receive information from or about visitors to such website or communications between such website and those visitors. Your participation in the Program constitutes your specific and unconditional consent to and authorization for our access to, receipt, storage, use, and disclosure of any and all such information, consistent with the policies and procedures set forth in our Privacy Policy located in the footer of the Walmart.com Site.

[Top of Page](#)

.....

3. Prohibitions

You understand and acknowledge that this Agreement is made between you and Wal-Mart.com USA, LLC and is solely for the purpose of allowing you to link to the Walmart.com Site.

As a condition to your acceptance and participation in the Program, you agree to the following prohibitions:

A. General Prohibitions.

You may NOT:

- engineer any website containing a Qualifying Link in such a manner that pulls Internet traffic away from Walmart.com;
- publish, link to, sell, otherwise distribute, or place a Qualifying Link on the same page or in close proximity to any Objectionable Content. For purposes of this Agreement, "Objectionable Content" means any material, including textual, audio or video material, which is offensive (including hate speech or violence against a particular group of people); contains any nudity, explicit violence or sexual material; contains depictions of violent or sexual acts; is defamatory to any group or individual; or promotes alcohol, tobacco, or gambling/lottery;
- publish, link to, sell, otherwise distribute, or place a Qualifying Link on any social networking sites, including, but not limited to, Facebook, MySpace, Twitter, etc.
- attempt to modify or alter our Site in any way;
- make any representations, either express or implied, or create an appearance that a visitor to your website is visiting our Site, e.g., "framing" or "wrapping" the Site in any manner without first obtaining in advance our express written permission. Such requests must be made in writing and sent to Walmart.com, Attn: Affiliate Program Manager, 7000 Marina Boulevard, Brisbane, CA, 94005;
- "scrape" or "spider" the Site or any other websites for content (such as images, logos or text);
- participate in Yahoo's Search Submit Pro (SSP) Search Marketing Program;
- place ads on, or participate in any way in, AdNetworks or Search Content Networks;
- employ, use or place any web browser add-ons, toolbars or pop-ups on your website;
- promote, or participate in the promotion of, Walmart.com's Pharmacy(www.walmart.com/pharmacy);
- link any Qualifying Link to any website other than our Site, including, for example, your own website;
- bid on our Trademarks at any website that provides search engine services and that results in driving traffic to any website, other than our Site, including your website;
- engage in any direct or indirect relationships with ISPs and/or mobile carriers that results in the delivery or act of address bar keyword and URL error trafficking (e.g., a user mistypes a web address in the ISP's address bar or search bar, and, as a result, is redirected to a web page that contains a Qualifying Link that directs the user to sites like Walmart.com).
- employ the use of any type of software download or technology which attempts to intercept or redirect traffic or Referral Fees to or from any website;
- use any Trademark (as defined in Exhibit A), or any Licensed Materials (as defined in Section 2), provided to you as a result of your participation in the Program to advertise or engage in services which result in a sale occurring on your website, whether or not you then have the item fulfilled through Walmart.com;
- without the prior written approval of Walmart.com, use any Trademark, or any

- Licensed Material in an advertisement that is not created or provided by Walmart.com in any way that might suggest or imply or mislead or is likely to mislead a visitor to your website into believing that Walmart.com, Wal-Mart Stores or any related entity was the creator or sponsor of such advertisement;
- re-distribute Licensed Materials (as defined in Section 2) to websites which can reasonably be viewed as Walmart.com's competitors, including but not limited to, Target, Kmart, Sears, JC Penney, Kohls, Circuit City, Best Buy, Amazon, eBay, Toys R Us, eToys, and KB Kids;
 - re-distribute, display or syndicate Licensed Materials and/or Walmart.com's datafeed, including any product information set forth therein, to any third party partner, network or agency;
 - employ, use, or receive any direct or indirect benefit from, any "cookie stuffing" methods (e.g., use of "cookie stuffing" to cause LinkShare's tracking systems to conclude that a user has clicked through a Qualifying Link - and to pay commissions accordingly - even if the user has not actually clicked through any such link);
 - install spyware on another person's computer; cause spyware to be installed on another person's computer, or use a context based triggering mechanism to display an advertisement that partially or wholly covers or obscures paid advertising other content on a website in a way that interferes with a person's ability to view that website;
 - display any material on a website containing a Qualifying Link which contains viruses, Trojan horses, worms, time bombs, cancel bots or other similar harmful or deleterious programming routines;
 - without the prior written approval of Walmart.com, use any widgets on your website that: (a) include any Trademarks (as defined in Exhibit A); (b) include any Licensed Materials (as defined in Section 2); or (c) directly or indirectly send traffic to Walmart.com;
 - post, publish, link to or place a Qualifying Link on the Walmart Facebook Page;
 - forward, redistribute, or otherwise repurpose any or all Qualifying Links to any third party;
 - release Wal-Mart's sales circulars, advertisements or other information prior to their authorized release dates; or
 - purchase products or services sold or promoted on Walmart.com through a Qualifying Link for resale or commercial use of any kind.

B. Prohibitions Regarding Use of Electronic Communications

Electronic Communication includes email messages, text messages, and any other form of non-verbal communication occurring without the use of physical mail. You may NOT do any of the following using Electronic Communication unless you first obtain in advance Walmart.com's express written permission. Such requests must be made in writing and sent to Walmart.com, Attn: Affiliate Program Manager, 7000 Marina Boulevard, Brisbane, CA, 94005. These prohibitions are in addition to, and not in place of, all prohibitions and restrictions that you are bound to under the LinkShare Affiliate Membership Agreement, as amended. You may not:

- generate or use Electronic Communication using or containing Trademarks (as defined in Exhibit A), or any variation or misspelling thereof, or products, or any of the Qualifying Links or URLs provided to you as part of the Program;
- send any other Electronic Communication that in any way suggests or implies or misleads or is likely to mislead (including without limitation, via the return address, subject heading, header information or message contents) a recipient into believing that Walmart.com, Wal-Mart Stores or any related entity was the sender or sponsor of such Electronic Communication or procured or induced you to send such Electronic Communication;
- forward, redistribute, or otherwise repurpose any Electronic Communication that Walmart.com sends to its affiliates and/or customers; and
- generate or send any unsolicited Electronic Communication (spam) under this Agreement

C. Prohibitions regarding use of Trademarks (as defined in Exhibit A)

In addition to the requirements and prohibitions regarding use of the Trademarks set forth in Exhibit A, and incorporated herein by reference, you may NOT:

- use the Trademarks in any manner not expressly authorized by this Agreement.
- use the Trademarks, or any variation or misspelling thereof, in metatags, hidden text or source code, in your domain name or any other part of your URL as further detailed in Exhibit A;

- bid on keywords as further detailed in Exhibit A;
- bid on our Trademarks at any website that provides search engine services and that results in driving traffic to any website, other than our Site, including your website;
- use Walmart.com or Wal-Mart Store's vendors or suppliers' logos, trade names, trademarks, graphic images, product images, product references and similar identifying material unless expressly and specifically provided by Walmart.com for use in the Program, unless used within a keyword string (i.e. Hanes T-Shirts).

In addition, you are bound to act in compliance with all applicable federal, state and local laws and regulations, including without limitation, the CAN-SPAM Act of 2003 ("CAN-SPAM") and the Children's Online Privacy and Protection Act of 1998 ("COPPA"). You shall protect, defend, indemnify and hold harmless us and our parent and related entities from and against any claims, actions, liabilities, losses, damages, costs or expenses, including without limitation, attorneys' fees and costs of litigation, even if such claims are groundless, fraudulent or false, incurred by us or our parent or related entities arising out of any content or activity by you or on your website or resulting from or in connection with your violation of any of the terms or prohibitions contained in this Agreement or any law, rule or regulation, including without limitation, claims for violations of third party intellectual property rights, and rights of privacy, including but not limited to CAN-SPAM and COPPA.

[Top of Page](#)

4. Order Processing

We will be responsible for all aspects of order processing and fulfillment of orders placed by customers who follow your Qualifying Links to the Walmart.com Site in accordance with applicable legal requirements. We reserve the right to reject orders that do not comply with any reasonable requirements that we periodically may establish. Among other things, we will prepare orders forms; process payments, cancellations, and returns; and handle customer service. Through the LinkShare Network, you have the ability to track sales made to customers who purchase products using your Qualifying Links and you can review reports summarizing this sales activity. To permit accurate tracking, reporting, and fee accrual, you must ensure that your Qualifying Links are properly formatted. The form, content, and frequency of the reports are limited to those reports and capabilities available through The LinkShare Network and may vary from time to time in our and/or The LinkShare Network's reasonable discretion. Walmart.com is not responsible for any changes that The LinkShare Network's format, timing, or types of reports available to members of The LinkShare Network and Walmart.com's Affiliates. Walmart.com will not be responsible for improperly formatted links regardless of whether you have made amendments to the code or not. In addition, we are unable to track or provide credit for sales from customers that are referred to us with browsers that do not have their cookies setting enabled.

[Top of Page](#)

5. Referral Fees

We will pay you Referral Fees on certain product sales to third parties generated from our Site only. For a product sale to generate a Referral Fee, the customer must

- use a browser that has its cookies setting enabled;
- follow a Qualifying Link (in the format specified by Walmart.com) from a site to the Walmart.com site;
- purchase the product using our automated ordering system;
- accept delivery of the product at the shipping destination; and
- remit full payment to us.

We will pay, to LinkShare for ultimate payment to you, Referral Fees on products that are actually purchased by a customer within three (3) days after the customer has initially entered our Site ("Referral Fee Time") as long as the customer reenters our Site directly during that time (and not through another affiliate link). We will not pay Referral Fees on any products are purchased on our Site when a customer has re-entered our Site (other than through a Qualifying Link from your website) after the Referral Fee Time, even if the customer previously followed a link from your website to our Site. Referral Fees will not be earned on products where a customer's purchase of the product derived from search results driven from free or natural search; this includes results containing Qualifying Links displayed in a search engine's free/non-paid, natural, or organic search results in

response to a search query which sends customers directly to Walmart.com without the customer first being sent to an affiliate site and the customer clicking on a link to arrive at Walmart.com. Purchases from Sam's Club and Samsclub.com, Pharmacy, Travel, Financial Services, tires, optical, Wal-Mart Connect Internet Service, Gift Cards, and Online Gift Cards are not eligible to earn Referral Fees. Customer Service invoice adjustments and reorders are not eligible to earn Referral Fees. Products that are entitled to earn Referral Fees under the rules set forth above are hereinafter referred to as "Qualifying Products."

[Top of Page](#)

6. Referral Fee Schedule

You will earn Referral Fees based on the sale price of Qualifying Products (as defined above), according to fee schedules to be established by us. "Sale price" means the sale price listed on our Site and excludes costs for shipping, handling, gift-wrapping, rebates, refunds, returns, chargebacks, cancellations and taxes. The current Referral Fee Schedule is available to you through the affiliate portal available to all members of the Program ("Affiliate Portal").

Because of the volume and breadth of items that we carry on our Site and our practice of continuing to add new items over time, certain items may not be listed in the Referral Fee database at the time purchases are made through your website. In addition, we reserve the absolute right and discretion to exclude items from our Referral Fee database (see e.g., purchases from Sam's Club and Samsclub.com, Pharmacy, Travel, Financial Services, tires, optical, Wal-Mart Connect Internet Service, Gift Cards, and Online Gift Cards). Therefore, you acknowledge and agree that we cannot and do not warrant or guarantee that you will be paid a referral fee on any item(s) or that all items eligible for a referral fee will be paid in accordance with the Referral Fee Rates listed on the Affiliate Portal. You acknowledge that in such circumstances, you will accept the Referral Fee Rates and payouts actually paid to you. We reserve the right, at our sole discretion, to change, modify, add or remove portions of this Referral Fee Schedule, at any time. If you have any questions concerning whether a certain item is eligible for a referral fee, please contact affiliates@walmart.com.

[Top of Page](#)

7. Referral Fee Payment

Approximately 60 days following the end of each calendar month, you will receive a check for the Referral Fees earned on products that were shipped during that month, less any taxes that we or LinkShare are required by law to withhold from the final payment to you. If a customer returns a product that generated a Referral Fee, you will see a deduction for the corresponding Referral Fee from your next monthly payment; if there is no subsequent payment, you will receive an invoice for the Referral Fee payable within sixty (60) days of your receipt of the invoice. All determinations of Qualifying Links and whether a Referral Fee is payable will be made by The LinkShare Network and/or Walmart.com and will be final and binding.

[Top of Page](#)

8. Policies and Pricing

Customers who buy products through this Program will be deemed to be customers of Walmart.com without affecting their status as your customer. Accordingly, all Walmart.com rules, policies, and operating procedures concerning customer orders, customer service, and product sales will apply to those customers with respect to their transactions at Walmart.com. We may change our policies and operating procedures at any time consistent with applicable laws. For example, we will determine the prices to be charged for products sold under this Program in accordance with our own pricing policies. Product prices and availability may vary from time to time. You may include current price information in your product descriptions only if such information is provided to you by Walmart.com, provided that any price information must be accompanied with a statement on your website indicating to the user that in the event of any price difference between your website and Walmart.com, the price listed on Walmart.com will govern. We will use commercially reasonable efforts to present current and accurate information, but we cannot guarantee the availability or price of any particular product.

[Top of Page](#)

9. Limited License; Restrictions

We grant you a limited, nonexclusive, nontransferable, revocable right to access our Site through the Qualifying Links solely in accordance with the terms of this Agreement and solely in connection with the Licensed Materials (as defined in Section 2), only as provided to you by us, through The LinkShare Network, or by other means selected by us, and solely for the purpose of identifying your website as a Program participant and to assist in generating the sale of Walmart.com products.

You acknowledge that this Agreement does not provide you with any intellectual property rights in the Licensed Materials other than the limited rights contained herein. We reserve all of our rights in the Licensed Materials and of our other proprietary rights. You may not sublicense, assign or transfer any such licenses for the use of the Licensed Materials, and any attempt at such sublicense, assignment or transfer is void. We may terminate your license to use the Licensed Materials for any reason at any time in our sole and absolute discretion. You agree to follow our Trademark Requirements in Exhibit A, as those may change from time to time. We may revoke your license at any time by giving you written notice.

[Top of Page](#)

10. Responsibility for Your Website or a Third Party Website on which you place a Qualifying Link ("Third Party Site")

You will be solely responsible for the development, operation, and maintenance of your website and for all content that appears on your website. For example, you will be solely responsible for:

- the technical operation of your website and all related equipment;
- creating and posting product descriptions on your website or a Third Party Site and linking those descriptions to our Site;
- updating product information, content and item descriptions (including, but not limited to, product price and availability) within 24 hours of any update of such product information, content and/or item description at Walmart.com or from datafeed content provided through Linkshare;
- the accuracy, timeliness and appropriateness of content posted on your website (including, among other things, all product-related materials);
- ensuring that materials posted on your website or a Third Party Site do not violate or infringe upon the rights of any third party (including, for example, copyrights, trademarks, privacy, or other personal or proprietary rights), or any term of this Agreement;
- monitoring your website content and the content of a Third Party Site to ensure your website or the Third Party Site does not publish, link to, sell or otherwise distribute Objectionable Content (as defined in Section 3);
- **removing** any Licensed Materials and Trademarks from your website or a Third Party Site as soon as any Objectionable Content appears on the website
- **notifying** us and The LinkShare Network of any Objectionable Content that appears or appeared on your website or a Third Party Site at any time during your participation in the program, within six (6) hours of its appearance, even if you immediately removed the Objectionable Content per the requirements of this Agreement or for any other reason;
- ensuring that content posted on your website or a Thrid Party Site is not libelous or otherwise illegal; and
- notifying us and The LinkShare Network of any malfunctioning of the Qualifying Links or other problems with your participation in the Program in accordance with the terms of this Agreement.

We disclaim all liability for these matters. Further, you will indemnify and hold us harmless from all claims, damages, and expenses (including, without limitation, attorneys' fees) relating to the development, operation, maintenance, and contents of your website.

[Top of Page](#)

11. Term of the Agreement

The term of this Agreement will begin upon our acceptance of your Application and will end when terminated by either party. You may terminate this Agreement at any time, with or without cause, by giving us (five) 5 days prior written notice of termination. We

may terminate this Agreement immediately at any time, with or without cause, by giving you written notice of termination. Upon termination, all Walmart.com related content and links shall be promptly removed from your website. You are only eligible to earn Referral Fees on sales of Qualifying Products occurring during the term, and fees earned through the date of termination will remain payable only if the related orders are not canceled or returned. In the event overpayment is made by us, you agree to promptly remit such excess payment upon notification by us. We may withhold your final payment for a reasonable time to ensure that the correct amount is paid.

[Top of Page](#)

12. Modification

We may modify any of the terms and conditions contained in this Agreement, at any time and in our sole discretion, by posting a change notice or a new agreement on our Site. We will also make commercially reasonable efforts to notify you of such changes prior to or upon implementation. Modifications may include, for example, changes in the scope of available Referral Fees, Referral Fee Schedules, payment procedures, and Program rules. IF ANY MODIFICATION IS UNACCEPTABLE TO YOU, YOUR ONLY RECOURSE IS TO TERMINATE THIS AGREEMENT. YOUR CONTINUED PARTICIPATION IN THE PROGRAM FOLLOWING OUR POSTING OF A CHANGE NOTICE OR NEW AGREEMENT ON OUR SITE AND/OR SENDING YOU THE CHANGE NOTICE WILL CONSTITUTE BINDING ACCEPTANCE OF THE CHANGE.

[Top of Page](#)

13. Relationship of Parties

You and we are independent contractors, and nothing in this Agreement will create any partnership, joint venture, agency, franchise, sales representative, or employment relationship between the parties. You will have no authority to make or accept any offers or representations on our behalf. You will not make any statement, whether on your website or otherwise, that reasonably would contradict anything in this Section.

[Top of Page](#)

14. Indemnification

You acknowledge that by entering into and performing its obligations under this Agreement, we do not assume and should not be exposed to the business and operational risks associated with your business, or any aspects of the operation or content of your website(s). Accordingly, in addition to any other indemnification obligations contained in this Agreement, you shall protect, defend, hold harmless and indemnify us and our parent or related entities from and against any and all claims, actions, liabilities, losses, costs and expenses, even if such claims are groundless, fraudulent or false (including court costs and reasonable attorneys' fees) incurred as a result of claims of customers or other third parties against us and our affiliates, licensors, suppliers, officers, directors, employees and agents arising from or connected with any of the content or activities of your website (including without limitation any activities or aspects thereof or commerce conducted thereon) or related business, or your misuse, unauthorized modification or unauthorized use of the services or materials provided by us hereunder.

[Top of Page](#)

15. Limitation of Liability

We will not be liable for indirect, special, or consequential damages (or any loss of revenue, profits, or data) arising in connection with this Agreement or the Program, even if we have been advised of the possibility of such damages. Further, our aggregate liability arising with respect to this Agreement and the Program will not exceed the total Referral Fees paid or payable to you under this Agreement.

[Top of Page](#)

16. Disclaimers

We make no express or implied warranties or representations with respect to the Program or any products sold through the Program (including, without limitation, warranties of

fitness, merchantability, non-infringement, or any implied warranties arising out of a course of performance, dealing, or trade usage). In addition, we make no representation that the operation of our Site will be uninterrupted or error-free, and we will not be liable for the consequences of any interruptions or errors; however, we will make commercially reasonable efforts to correct errors or interruptions promptly.

[Top of Page](#)

17. Independent Investigation

YOU ACKNOWLEDGE THAT YOU HAVE READ THIS AGREEMENT AND AGREE TO ALL ITS TERMS AND CONDITIONS. YOU UNDERSTAND THAT WE MAY AT ANY TIME (DIRECTLY OR INDIRECTLY) SOLICIT CUSTOMER REFERRALS ON TERMS THAT MAY DIFFER FROM THOSE CONTAINED IN THIS AGREEMENT OR OPERATE WEB SITES THAT ARE SIMILAR TO OR COMPETE WITH YOUR WEB SITE. YOU HAVE INDEPENDENTLY EVALUATED THE DESIRABILITY OF PARTICIPATING IN THE PROGRAM AND ARE NOT RELYING ON ANY REPRESENTATIONS, GUARANTEE, OR STATEMENT OTHER THAN AS SET FORTH IN THIS AGREEMENT.

[Top of Page](#)

18. Miscellaneous

This Agreement will be governed by the laws of the United States and the state of Arkansas, without reference to rules governing choice of laws. Any action relating to this Agreement must be brought in the federal or state courts having jurisdiction and venue in or for Benton County, Arkansas and you irrevocably consent to the jurisdiction of such courts. You may not assign this Agreement, by operation of law or otherwise, without our prior written consent. Subject to that restriction, this Agreement will be binding on, inure to the benefit of, and enforceable against the parties and their respective successors and assigns. Our failure or agreement not to enforce your strict performance of any provision of this Agreement in a given instance will not constitute a waiver of our right to subsequently enforce such provision or any other provision of this Agreement.

[Top of Page](#)

19. Publicity

You shall not create, publish, distribute, make or permit any public announcement of this Agreement or the relationship contemplated hereunder, (including, but not limited to, any press release, client list, screen shot, advertisement or any promotional material) without first submitting such material to us and receiving our written approval, which we may withhold in our sole discretion.

[Top of Page](#)

20. Confidentiality

Except as otherwise provided in this Agreement or with our prior written consent, you agree that all information including, without limitation, the terms of this Agreement, our business and financial information, our customer lists and purchase history, and our pricing and sales information, shall remain strictly confidential and shall not be utilized, directly or indirectly, by you for your own business purposes or for any other purpose except and solely to the extent that any such information is generally known or available to the public through a source or sources other than you or your affiliates.

Notwithstanding the foregoing, you may deliver a copy of any such information (a) pursuant to a subpoena issued by any court or administrative agency, (b) to your accountants, attorneys, or other agents on a confidential basis, and (c) otherwise as required by applicable law, rule, regulation or legal process, upon written notification to Walmart.com.

[Top of Page](#)

21. Remedies to Walmart.com

Violation of any of the terms or prohibitions contained in this Agreement may result in, among other things, (a) the immediate termination of this Agreement; (b) the withholding of Referral Fees due to you; or (c) the commencement of an action by Walmart.com against

you seeking, without limitation, injunctive relief, recovery of actual, statutory or punitive damages.

We have the right in our sole and absolute discretion to monitor your website at any time and from time to time to determine if you are in compliance with the terms of this Agreement, and you agree to provide us with unrestricted access to your website for such purpose.

Last Updated: October 29, 2010

[Top of Page](#)

Exhibit A - Trademark Requirements

These requirements apply to your use of Walmart.com and other trademarks and service marks belonging to Wal-Mart.com USA, LLC, Wal-Mart Stores, Inc. or other related entities (the "Trademarks") in content that has been approved by us.

1. You may use the Trademarks only for purposes expressly authorized by us.
2. You may not modify the Trademarks in any manner. For example, you may not change the proportion, color, or font of the Trademarks.
3. You may not display the Trademarks in any manner that implies endorsement of your website or business by Walmart.com outside of your involvement in the Program.
4. You may not use the Trademarks to disparage Walmart.com, its products or services, or in a manner which, in our reasonable judgment, may diminish or otherwise damage our good will in the Trademarks.
5. Each Trademark must appear by itself, with reasonable spacing (at least the height of the Trademark) between each side of the Trademark and any other graphic or textual image. You may place the Walmart.com name or logo adjacent to competitive brands, subject to the requirements of this Agreement, including prohibitions against objectionable material and websites.
6. You must use the TM symbol next to the trademarks. You must use the SM symbol next to the service marks.
7. You must include the following statement in your materials that include the Trademarks: "WALMART.COM SM is a service mark of Wal-Mart.com USA, LLC and Wal-Mart Stores, Inc." You must include similar statements for any other Trademarks used on an ongoing basis in your materials.
8. You acknowledge that all rights to the Trademarks are our exclusive property and all goodwill generated through your use of the Trademarks will inure to our benefit.
9. YOU MAY NOT USE THE TRADEMARKED NAMES, WAL-MART, WALMART.COM, WAL-MART STORES, OR ANY VARIATIONS OR MISPELLINGS THEREOF, IN ANY MANNER INCLUDING KEYWORD BIDDING ON SEARCH ENGINES; YOU MAY NOT USE WAL-MART, WALMART.COM, OR ANY VARIATION OR MISPELLINGS THEREOF, IN METATAGS OR TO DIRECT TRAFFIC TO ANY WEBSITE OTHER THAN OUR SITE; YOU MAY NOT USE WAL-MART, WALMART.COM, OR ANY VARIATIONS OR MISPELLINGS THEREOF, IN HIDDEN TEXT OR SOURCE CODE ; YOU MAY NOT USE WAL-MART, WALMART.COM, OR ANY VARIATIONS OR MISPELLINGS THEREOF, IN YOUR DOMAIN NAME OR ANY OTHER PART OF YOUR UNIVERSAL RECORD LOCATOR.
10. You may not bid on any keyword or on any Pay per Click Search Engines (PPCSEs) where such keyword is one of our Trademarks or any variation or misspelling of one of our Trademarks (see the non-exclusive list of examples set forth below in Section 15). Further, you may not bid on any word or term that is confusingly similar to any of our Trademarks standing alone. You may, however, bid on keyword strings that incorporate our Trademarks (e.g., "Walmart.com Electronics", "Wal-Mart Toys"); provided, however, that such permissible keyword strings must be used in a manner that directs traffic only to our Site. You may not bid on keywords strings that contain the terms "Rollbacks" or "Advertised Values".
11. You may not employ any "fat finger" domains or typosquatters redirecting web traffic to your website. A typosquatter for "fat finger" domain is any domain that amounts to misspellings of any registered or unregistered Trademarks.
12. You may not bid on any keyword or on any PPCSEs that is one of our competitors' trademarks (or a derivation of a competitor's trademark), or any other word or term that is likely to cause confusion regarding its affiliation with the competitor. Examples of these keywords include, but are not limited to: "Target", "Kmart", "Sears", "JC Penney", "Toys R Us", "Amazon", "Circuit City", "eToys", and "KB Kids".
13. You may not bid on restricted manufacturer brand terms, including but not limited to "MagicJack", or any derivatives thereof that are likely to cause confusion regarding its affiliation with Walmart.com, its affiliates or you, in any paid search.

14. You may not use the Trademarks alongside or in conjunction with the following terms: "percent (%) off", "sale", or "coupons".
15. Walmart.com may, in its sole discretion, terminate you or withhold payment of your Referral Fees for the days that we determine that you were bidding in violation of the keyword bidding requirements above.
16. The list below sets forth examples of impermissible keywords, "fat-finger" domains, and variations of Trademarks that you may **not** bid on. The list is for example purposes only and is not a complete list of prohibited words which infringe a Trademark, and therefore violate a term of this Agreement.

walmart	walmart.com	wal-mart	wal mart	www.walmart.com
walmart	wal-mart.com	wall mart	walmart.com	www.wallmart.com
wal mart.com	Walmart.com	Walmart.c	www.wal-mart.com	www.walmart
walmart stores	wal-mart store	walmarts	www.wal mart.com	wal mart stores
wall-mart	walmart supercenter	site:www.walmart.com walmart	wal-mart supercenter	super Walmart
walmart stores	super wal mart	walmart.com.	walmart store	walt mart
walmart.com	wal mart.com	walmart online	wal mart.com	www.wal-mart
Walmart.c_om	Walmart.c	site:walmart.com	-	-

We reserve the right in our sole discretion to modify these requirements at any time.

Last Updated: March 25, 2010

[Top of Page](#)

Exhibit B - Networks & Sub-Affiliates

These requirements apply to your use of the Program as a Network. A "Network" is defined as any affiliate that operates one or more websites as sub-affiliates through one (1) Application to the Program.

1. You agree to all the terms and conditions of this Agreement, including all attached Exhibits, on behalf of all sub-affiliates operating under the Network and are responsible for any action or inaction by such sub-affiliates.
2. You must seek prior written approval from Walmart.com before choosing to operate as a Network. Such requests must be made in writing and sent to Walmart.com, Attn: Affiliate Program Manager, 7000 Marina Boulevard, Brisbane, CA, 94005.
3. You must provide to Walmart.com a list of all sub-affiliates working within your Network within twenty-four (24) hours of any request from Walmart.com.
4. Any violation of this Agreement by any sub-affiliate may result in immediate termination of the entire Network from the Program.

[Top of Page](#)

Exhibit C - RSS Feeds

Walmart.com offers free RSS (Really Simple Syndication) feeds for personal, noncommercial use. Any other uses, including without limitation the incorporation of advertising into or the placement of advertising associated with or targeted towards the content of our RSS feeds, are prohibited. You must use the RSS feeds as provided by Walmart.com, and you may not edit or modify the text, content or links supplied by Walmart.com. You must always have the most current RSS feed visible. You must accompany all price information with a statement on your website indicating to the user that in the event of any price difference between your website and Walmart.com, the price listed on Walmart.com will govern.

The RSS feeds may be used only with those platforms from which a working link is made available that, when accessed, takes the viewer directly to the display of the full product offer on Walmart.com. Any display of the content of the RSS feeds must permit successful

linking to, redirection to or delivery of the applicable Walmart.com web page. You may not insert any intermediate page, splash page or other content between the RSS link and the applicable Walmart.com web page.

Walmart.com retains all ownership and rights in the content of our RSS feeds.

Walmart.com reserves the right to discontinue providing any or all of the RSS feeds at any time and to require you to cease displaying, distributing or otherwise using any or all of the RSS feeds for any reason including, without limitation, your violation of any provision of this Agreement.

Last Updated: March 25, 2010

[Top of Page](#)

[Security & Privacy](#)

| [Terms & Conditions](#)

| [Store Finder](#)

| [Site Directory](#)

| [About Walmart.com](#)

