

Amazon Web Services Blog

Amazon Web Services, Products, Tools, and Developer Information...



RECENT POSTS

Amazon EC2 Beta

I'll have a Lemonade and Some Links, Por Favor...

Openfount S3 Manager

Friday Wrapup

S3Fox Organizer for Amazon

Blueorganizer

New Version of Amazon Products Feed

Java.util.Map Implementation for Amazon S3

Thought-Provoking Series of S3 Posts

Amazon S3 and SmugMug

RECENT COMMENTS

Fred Domke on Amazon EC2 Beta

Leitsystemseiten on Amazon EC2 Beta

Glenn Fleishman on Amazon EC2 Beta

Bryan Kennedy on Amazon S3 and SmugMug

John on Openfount S3 Manager

BPAndrew on Openfount S3 Manager

Bill Donahue on Openfount S3 Manager

Jonathan on Friday Wrapup

Hugo Dalgleish on Improved Amazon Shopping in Second Life

rikomatic on Improved Amazon Shopping in Second Life

RECENT LINKS

Brian "Krow" Aker's Idle Thoughts - S3, JungleDisk, and where did I put my tape drive?

Dr. Dobb's | Synchronize Now! | August 2, 2006

tsandho's boring crap: Backup Goodness

Make Money Online - Alexa traffic rank: what it is and why you should care (or not)

birdDog digital Blog & Podcast - Mac Sales & Support for Small Businesses & Creative Firms: von Kempelen - Home - Great Idea

I hate hard drives - Julio Garcia dot org - Now with more 'Whaaaaaaas'?

Stairways Software: How Do I Make An Amazon S3 Hosted Website Using Interarchy

akatsuki.co.uk: Backup in the Age of the Terabyte

Chip Griffin: Pardon the Disruption: 10 Ways Web 2.0 Promises to Change the Way We Live and Work

Mechanical Turk Monitor: Happy Monday Turkers!

WEB SERVICES WISH LIST

amazon.com.

DEVELOPERWORKS : BUILD YOUR OWN FEED : ARTICLES AND TUTORIALS FOR WEB SERVICES AND XML

Achieve semantic interoperability in a SOA

DB2 9 XML performance characteristics

Create vector graphics in the browser with SVG

Exposing an Apache Derby stored procedure as a Web service

Enable C++ applications for Web service using XML-RPC

Amazon EC2 Beta

Innovation never takes a break, and neither do I. From the steaming hot beaches of Cabo San Lucas I would like to tell you about the Amazon Elastic Compute Cloud, or Amazon EC2, now open for limited beta testing, with more beta slots to open soon.

Amazon EC2 gives you access to a virtual computing environment. Your applications run on a "virtual CPU", the equivalent of a 1.7 GHz Xeon processor, 1.75 GB of RAM, 160 GB of local disk and 250 Mb/second of network bandwidth. You pay just 10 cents per clock hour (billed to your Amazon Web Services account), and you can get as many virtual CPUs as you need. You can learn more on the [EC2 Detail Page](#). We built Amazon EC2 using a virtual machine monitor by the name of [Xen](#).

Amazon EC2 works in terms of AMIs, or Amazon Machine Images. Each AMI is a pre-configured boot disk -- just a packaged-up operating system stored as an [Amazon S3](#) object. There are web service calls to create images, and to assign them to virtual CPUs to run your application. If your application consists of the usual web server, business logic, and database tiers, you can build distinct AMIs for each tier, and then spawn one or more instances of each type based on the load.

In a previous post, [Sometimes You Need Just a Little...](#), I alluded to the new world of scalable, on-demand web services. In that post I talked about the fact that sometimes a little bit of storage is all you need.

Sometimes you need a lot of processing power, and sometimes you need just a little. Sometimes you need a lot, but you only need it for a limited amount of time. Perhaps you are doing some number crunching, some in-depth text processing, some scientific research, or your end-of-month accounting. Or perhaps you want to experiment with some radical new parallel processing algorithm for a week or two. In any of these situations, acquiring sufficient hardware to accommodate the high-water mark of your usage would definitely not be economical. There are already some interesting examples of this in the [Amazon EC2 Discussion Forums](#). For example:

- Daniel Drucker [says](#) "We're planning on using it for functional MRI analysis. We have large datasets which, when they're being processed, require a cluster of 15-20 machines... but we only need those machines for a couple hours every few weeks."
- In the same thread, a user by the name of spanglu says "Let's say your back-office inventory app is web-based, but that you're only using it from 7am to 7pm. You can cut your server costs by 50%. Take this to its logical conclusion - only start up an instance when you actually need your inventory app..."

Put another way, time is another interesting axis of scalability.

Before the advent of Amazon EC2, you had to buy or rent sufficient servers to cover your present needs, and you also had to be able to anticipate, forecast, and pay (in advance) for enough hardware, storage, and network bandwidth to accommodate organic growth as well as bursts of traffic brought upon by popular sites such as [Digg](#) or [Slashdot](#). If you are too generous with your planning, hardware sits idle. Too frugal, and your chance at fame and fortune may very well pass, as thousands of would-be users are greeted with a "site too busy" message.

With Amazon EC2, you don't need to acquire hardware in advance of your needs. Instead, you simply turn up the dial, spawning more virtual CPUs, as your processing needs grow. During the beta you can run up to 20 virtual servers per account, or more by special arrangement.

Returning to our hypothetical three-tier application, Amazon EC2 gives you the ability to control network access on a very fine-grained basis. For example, you can allow the outside world to talk to the web servers, but not to the business logic or to the database server. You can allow the web server to talk to the business logic, and the business logic to talk to the database, and that's it. You also get free, fast-path access to [Amazon S3](#), making S3 a natural place to store your raw data and your results.

Ok, developers, now that you have access to this computing resource, how can you take advantage of it?

For starters I would recommend that you get your hands on a copy of Cal Henderson's new book, [Building Scalable Web Sites](#). Cal was one of the engineers behind [Flickr](#), and his real-world experience in building a high-traffic site are aptly recounted and generalized in his book.

Second, start looking at higher level packages that will let you decompose immense computational tasks into a form suitable for parallel processing. It is no secret that one of [Google's](#) tricks is a software framework that they call [MapReduce](#). This framework simplifies the task of performing similar processing steps on millions or even billions of pieces of data. Doug Cutting's [open source version of MapReduce](#) is probably worth investigating, as is [Starfish](#), an implementation of MapReduce in [Ruby](#).

Third, consider what you can do to help other developers use Amazon EC2. What about building specialized AMIs and then selling them to other developers? Preload an AMI with a popular open source stack (being careful to respect any and all redistribution prohibitions in the software licenses). What about an advanced monitoring system that spools up additional machine images in times of heavy load, then safely winds them down after the load goes away? There are a lot of ways that you can add value on top of what's already there.

Finally, the phrase "leveling the playing field" is used quite often with regard to the full line of Amazon's Web Services, but I think that it is truly appropriate here. I find it worthwhile to imagine that there is a developer (maybe its you) sitting in a dorm room or a spare bedroom somewhere, burning the midnight oil and creating the next killer web site or search algorithm. Without services like Amazon EC2, Amazon SQS, and Amazon S3, you'd need to invest a lot of capital, max out your credit cards, re-mortgage your house, and take a pretty big financial risk just to see if your ideas will work in the real world. With these new services your risk is a lot smaller, but your potential reward is in no way diminished.

That's about all I have to say. You might find Jon Udell's [S3 screencast](#) interesting, and you can always find something entertaining on the requisite [Slashdot thread](#).

-- Jeff,

August 25, 2006 in [Announcements](#) | [Permalink](#) | [Comments \(3\)](#) | [TrackBack \(0\)](#)

I'll have a Lemonade and Some Links, Por Favor...

I simply couldn't wait until next week to post a couple of new items!

In [S3 Meets R3 \(Reliability, Robustness, and Resilience\)](#), the authors benchmark Amazon S3 against the venerable SCP (Secure Copy) protocol. You can read the entire article to see the details, but the conclusion pretty much sums it up:

Amazon's S3 Services provide a standards-based and high performance mechanism for managing content. S3's performance characteristics make it ideal for end-user as well as enterprise applications where cost-effective data back up is desired. During our regression we never had a problem with reading or writing files. The S3 service had enterprise-class availability and reliability.

In addition to declaring that S3 is enterprise-ready, the authors noted that the entire battery of tests cost just 52 cents to perform.

Also in Dr. Dobb's, be sure to read about [Synchronizing Files with .Net 2.0, S3, and FTP](#).

Meanwhile, Martin Kochanski continues his [informative series of S3 articles](#).

August 23, 2006 in [Thought Pieces](#) | [Permalink](#) | [Comments \(0\)](#) | [TrackBack \(0\)](#)

Openfount S3 Manager

I am posting from the beach to tell you about the [Openfount S3 Manager](#).

To get this running you simply click [here](#).

Once you do this, an invisible applet is loaded to gain access to the file system and the GUI is displayed. AT this point you need to enter your Amazon keys. After you do this you get a two-pane view -- local files on the left and your S3 buckets on the right. There's also a complete ACL manager.

Ok, I had better put the laptop away before I get caught!

-- Jeff,

This post was updated to reflect the fact that the local copy of the Java policy file is no longer needed.

August 22, 2006 in [Cool Sites](#) | [Permalink](#) | [Comments \(3\)](#) | [TrackBack \(0\)](#)

Friday Wrapup

I'm just about to sign off in anticipation of a vacation week devoid of computers, Blackberries, web services, blogs, and even Second Life. Before I go, here are a few things to keep you busy in the week to come.



Photo by Jeff Atwood

RSS FEED

Subscribe to this site



Enter your email address:

Subscribe

Delivered by FeedBurner

USEFUL LINKS

[Amazon Web Services](#)

[AWS Developer Forums](#)

[Amazon Associates JP](#)

[A9 Developer Blog](#)

[Media/Analyst Inquiries](#)

Add me to your TypePad People list

AUGUST 2006

Sun	Mon	Tue	Wed	Thu	Fri	Sat
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31		

ARCHIVES

[August 2006](#)

[July 2006](#)

[June 2006](#)

[May 2006](#)

[April 2006](#)

[March 2006](#)

[February 2006](#)

[January 2006](#)

[December 2005](#)

[November 2005](#)

USEFUL BOOKS

 Jason Bloomberg: Service Orient or Be Doomed! : How Service Orientation Will Change Your Business

 Tom Standage: The Turk : The Life and Times of the Famous Eighteenth-Century Chess-Playing Machine

 Theodore W. Leung: Professional XML Development with Apache Tools : Xerces, Xalan, FOP, Cocoon, Axis, Xindice (Wrox Professional Guides)

 Jason Levitt: The Web Developer's Guide to Amazon E-Commerce Services: Developing Web Applications Using Amazon Web Services and PHP

 Will Inverson: Real World Web Services

 John Mueller: Mining Amazon Web Services: Building Applications with the Amazon API

 Paul Bausch: Amazon Hacks: 100 Industrial-Strength Tips and Tools

 Rachel Andrew: The CSS Anthology : 101 Essential Tips, Tricks and Hacks



In addition to releasing that new version of [Blueorganizer](#), Alex Iskold just published a very nice article -- [Amazon, the Real Web Services Company](#). As Alex says, "Amazon has been quietly, and recently not so quietly, building the blocks of the next generation software platform." While we don't generally talk about our "big picture" strategy, this article does a nice job of synthesizing an accurate model of that strategy based on our products. If you want the bird's-eye view of our entire lineup of services, this is a must-read.

[Martin Kochanski](#) has apparently been typing and blogging non-stop this week. He's released the first 5 items in his series of "S3 in Business" posts. Here's what he has so far:

1. [S3 in Business: 1-Introduction](#)
2. [S3 in Business: 2-Tunesafe](#)
3. [S3 in Business: 3-Tunesafe.com](#)
4. [S3 in Business: 4-My Tunesafe](#)
5. [S3 in Business: 5-Money in Golden Buckets](#)

I'll be among the last to ready next week's series, unless I get "web withdrawal" and sneak off to an internet cafe for a quick fix.

And with that I'm off to the Silicon Valley [TechCrunch party](#), a quick flight home tomorrow, and then on to [Mexico](#) on Monday.

-- Jeff,

August 18, 2006 in [Web Services News](#) | [Permalink](#) | [Comments \(1\)](#) | [TrackBack \(0\)](#)

S3Fox Organizer for Amazon


The [S3Fox Organizer for Amazon](#) is another Firefox extension!

After you install it and perform the usual Firefox restart, your Tools menu will contain an entry for "Amazon S3." Select this entry and you will see the S3Fox user interface, which consists of two panes -- local file system and S3 account -- and a number of ancillary controls. Start by getting an [S3 account](#) (if you don't have one already) and clicking the button labeled "Manage Accounts." Select a name for your account and enter your Access Key and your Secret Key.

Once you've done this, the right pane will reveal the contents of your S3 account. From there you can select items in the local (left side), and click the right-pointing arrow to copy them over to S3, or vice versa. The transfers are queued up (and shown at the bottom) and performed one at a time. You can also right-click on the items in the S3 account to set the permissions, delete them, copy the URL of the item, and so forth.

The extension is available as a free download, but the author would probably appreciate a donation or two, via the [download page](#).

-- Jeff,



August 18, 2006 in [Cool Sites](#) | [Permalink](#) | [Comments \(0\)](#) | [TrackBack \(0\)](#)

Blueorganizer

Alex from [Adaptive Blue](#) sent me a note about about a new release of their flagship product, the [Blueorganizer](#).


Blueorganizer is an intelligent bookmark manager. Behind the scenes it uses [Amazon S3](#) to store special types of bookmarks called bluemarks. If you install the product on more than one computer, your Bluemarks will be transparently shared and made available on all of them. Your Bluemarks are shown in a sidebar panel called the Blueorganizer, complete with ratings and tags. Once a product or site is marked, there are many ways to pivot around the information -- finding other products by the same author, in the same product categories, with the same tags, and so forth.

Bluemarks can represent products (from Amazon and a number of other e-commerce sites) or web sites. There's complete context-menu integration and easy access to the Blueorganizer using a toolbar button. The organizer has a myriad of well-organized options for sorting, filtering, and searching.

You can also share your selections through an RSS feed (stored in S3, of course), and you can publish them on a site using a Bluebadge.

All in all, very functional and very nicely done!

-- Jeff,



August 18, 2006 in [Cool Sites](#) | [Permalink](#) | [Comments \(0\)](#) | [TrackBack \(0\)](#)

New Version of Amazon Products Feed

[MrRat](#) has released an update to his popular [Amazon Products Feed](#) script.

The new version includes three major features, all of which came about in response to requests from users. The new features include the ability to filter out adult products, creation of a breadcrumb trail as the user browses through the catalog, and display of alternate versions of a book. Browse node support has been improved, with a special focus on the DE (German) locale.

MrRat has been a very long-term developer and supporter of AWS and it has been very enjoyable to watch him build his [community](#) and refine his product over the past couple of years.

You may also enjoy reading his recent article, [The Profit in Altruism](#).

-- Jeff,



August 16, 2006 in [Announcements](#) | [Permalink](#) | [Comments \(0\)](#) | [TrackBack \(0\)](#)

Java.util.Map Implementation for Amazon S3

[Tom White](#) dropped me a note to tell me that he had created an [implementation of java.util.Map](#) on top of [Amazon S3](#). As Tom notes, "Normally we strive for an implementation that is efficient in terms of speed, however in this case I was very aware of the monetary cost of each method."

Tom took care to document the number of S3 calls made by each method in his implementation, and, being a faithful S3 user, also stored the [JavaDoc documentation for S3Map](#) in S3! This table is quite interesting and should be very helpful to any developer using the class.

Way to go Tom!

-- Jeff,

August 14, 2006 in [Developer Tools](#) | [Permalink](#) | [Comments \(0\)](#) | [TrackBack \(0\)](#)

Thought-Provoking Series of S3 Posts

I met [Martin Kochanski](#) developer of [Cardbox](#) in London last month. We met at the [Athenaeum Club](#) and had a very pleasant working lunch

As we talked, it was really clear to me that Martin held a number of interesting opinions about all sorts of subjects and we talked about blogging as an information sharing vehicle. Today, I see that Martin has posted [\\$3 in Business: 1 - Introduction](#). This is the first in a two week series of articles. The first part is interesting and thought-provoking, and I'm looking forward to the rest of the series.

-- Jeff,

August 14, 2006 in [Thought Pieces](#) | [Permalink](#) | [Comments \(1\)](#) | [TrackBack \(0\)](#)

Amazon S3 and SmugMug

SmugMug CEO Don MacAskill describes how they use [Amazon S3](#) in his newest blog post, [Amazon S3 = The Holy Grail](#). In the post, Don reveals that SmugMug stores 500 million images which collectively occupy 300 terabytes of storage.

Don goes on to reveal how S3 allowed SmugMug to take their architecture to the next level of safety, security, economy, and speed. This sentence in particular caught my eye:

Perhaps even more importantly, our cash-flow situation is vastly improved. Instead of paying \$25,000 for a handful of terabytes of redundant storage up-front, even before they're used, we now pay \$0.15/GB/month as we use it.

An entire business school case study should be written around that one sentence!

By providing infrastructure components on an as-needed, pay as you go basis, I believe that we are changing the economics of building large-scale web applications. Don notes that S3 is a playing-field leveler, lowering barriers to entry. Indeed, the chasm between concept and reality is now significantly narrower. Ventures that once required angel funding to get off the ground can now be built from the comfort of a dorm room or a kitchen table. As long as you have a business model that provides you with a return as soon as you have primed the pump, you can get started without making those up-front \$25K investments on a regular basis.

-- Jeff,

August 14, 2006 in [Thought Pieces](#) | [Permalink](#) | [Comments \(2\)](#) | [TrackBack \(1\)](#)

