

# Amazon Simple Storage Service (S3) SOAP Interface

Mike Gartrell

October 5, 2006



# Agenda

- S3 functionality, pricing, and design requirements
- Storage concepts
- Authentication
- Operations
- SOAP API examples
- Demo

# Functionality

- Write, read, and delete objects from 1 byte to 5 gigs in size. Number of objects is unlimited.
- Each object is stored and retrieved via a unique key, assigned by the developer.
- Authentication mechanisms are provided to control access. Objects can be made private or public, and rights can be granted to specific users.
- REST and SOAP interfaces are provided. Only SOAP will be presented here.
- HTTP (default) and BitTorrent download protocols are supported. Other protocols may be added in the future.

# Pricing

- \$0.15 per GB-month of storage used
- \$0.20 per GB of data transferred
- No other fees

# Design Requirements (1 of 2)

- Scalable: Amazon scales the service by simply adding additional nodes (servers) to the system. Scales in terms of storage, request rate, and users.
- Reliable: Data stored “durably”. Amazon claims 99.99% availability and no single points of failure.
- Fast: Fast enough to support high-performance applications. Any bottlenecks can be fixed by adding more nodes.
- Inexpensive: Built from inexpensive commodity hardware.
- Simple: REST and SOAP interfaces should be easy to use.

## Design Requirements (2 of 2)

- S3 system must support the needs of both internal Amazon applications and external developers. Must be fast and reliable enough to run Amazon.com's websites, while also permitting any developer to use it for any data storage application.
- Implies that Amazon uses S3 for their internal and customer-facing applications.

# S3 Storage Concepts

- Objects
- Buckets
- Keys

# S3 Storage Concepts: Objects

- Objects are the fundamental storage entities in S3. They are composed of object data and metadata, and range in size from 1 byte to 5 gigabytes.
- Metadata is a set of name-value pairs associated with the object. Default metadata includes date last modified, and standard metadata such as Content-Type. Custom metadata can be specified at the time the object is stored.



# S3 Storage Concepts: Buckets

- A bucket is a container for objects. Every object is contained within a bucket.
- The bucket provides a unique namespace for management of objects contained in the bucket.
- No limit on the number of objects per bucket.
- Bucket namespace is global. Each user is limited to 100 buckets.

# S3 Storage Concepts: Keys

- A key is a unique identifier for an object within a bucket. Every object has exactly one key.
- The bucket name + key uniquely identify each object.
- The full unique address for each object consists of service endpoint + bucket name + key. Example: “<http://s3.amazonaws.com/testbucket/2006-09-30/photo.jpg>”, where “testbucket” is the name of the bucket, and “2006-09-30/photo.jpg” is the name of the key.

# Authentication (1 of 2)

- Most requests to S3 require authentication to verify that the requestor is authorized to perform the requested action. If the owner of an object has specifically granted anonymous access to an object or bucket, then authentication isn't required.
- When you register for S3, you will be provided with a AWS Access Key ID and a AWS Secret Access Key.

## Authentication (2 of 2)

- To authenticate a request, the following authentication information is included in the SOAP request:
  - AWSAccessKeyId
  - Timestamp
  - Signature: The HMAC-SHA1 digest of the concatenation of “Amazon S3 + OPERATION + Timestamp”, using your AWS Secret Key as the signing key.

# S3 Operations (1 of 2)

- Common operations:
  - Create a bucket (CreateBucket)
  - Write an object (PutObjectInline)
  - Read an object (GetObject)
  - Deleting an object (DeleteObject)
  - Listing keys (ListBucket) – list keys contained in a bucket

# S3 Operations (2 of 2)

- All SOAP API operations:
  - Service
    - ListAllMyBuckets
  - Buckets
    - CreateBucket, DeleteBucket, ListBucket, GetBucketAccessControlPolicy, SetBucketAccessControlPolicy, GetBucketLoggingStatus (beta), SetBucketLoggingStatus (beta)
  - Objects
    - PutObjectInline, PutObject, GetObject, GetObjectExtended, DeleteObject, GetObjectAccessControlPolicy, SetObjectAccessControlPolicy

# SOAP API Examples: CreateBucket

- Create a bucket named “quotes”:

- Sample request:

```
<CreateBucket xmlns="http://s3.amazonaws.com/doc/2006-03-01/">  
  <Bucket>quotes</Bucket>  
  <AWSAccessKeyId>1D9FVRAYCP1VJS767E02</AWSAccessKeyId>  
  <Timestamp>2005-01-31T23:59:59.183Z</Timestamp>  
  <Signature>Iuyz3d3P0aTou39dzbq7RrtSFmw=</Signature>  
</CreateBucket>
```

- Sample response:

```
<CreateBucketResponse xmlns="http://s3.amazonaws.com/doc/2006-03-01/">  
  <CreateBucketResponse>  
    <Bucket>quotes</Bucket>  
  </CreateBucketResponse>  
</CreateBucketResponse>
```

# SOAP API Examples: PutObjectInline (1 of 3)

- Write some text and metadata into the “Nelson” object in the “quotes” bucket, and make the object readable by anonymous users.



# SOAP API Examples: PutObjectInline (2 of 3)

- Sample request:

```
<PutObjectInline xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>quotes</Bucket>
  <Key>Nelson</Key>
  <Metadata>
    <Name>Content-Type</Name>
    <Value>text/plain</Value>
  </Metadata>
  <Metadata>
    <Name>family</Name>
    <Value>Muntz</Value>
  </Metadata>
  <Data>aGEtaGE=</Data>
  <ContentLength>5</ContentLength>
  <AccessControlList>
    <Grant>
      <Grantee xsi:type="Group">
        <URI>http://acs.amazonaws.com/groups/global/AllUsers</URI>
      </Grantee>
      <Permission>READ</Permission>
    </Grant>
  </AccessControlList>
  <AWSAccessKeyId>1D9FVRAYCP1VJS767E02</AWSAccessKeyId>
  <Timestamp>2005-01-31T23:59:59.183Z</Timestamp>
  <Signature>Iuyz3d3P0aTou39dzbq7RrtSFmw=</Signature>
</PutObjectInline>
```

# SOAP API Examples: PutObjectInline (3 of 3)

- Sample response:

```
<PutObjectInlineResponse xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <PutObjectInlineResponse>
    <ETag>&quot;828ef3fdfa96f00ad9f27c383fc9ac7f&quot;</ETag>
    <LastModified>2005-11-18T00:27:11.991Z</lastModified>
  </PutObjectInlineResponse>
</PutObjectInlineResponse>
```

# SOAP API Examples: GetObject (1 of 2)

- Get the “Nelson” object from the “quotes” bucket:

- Sample request:

```
<GetObject xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>quotes</Bucket>
  <Key>Nelson</Key>
  <GetMetadata>true</GetMetadata>
  <GetData>true</GetData>
  <InlineData>true</InlineData>
  <AWSAccessKeyId>1D9FVRAYCP1VJS767E02</AWSAccessKeyId>
  <Timestamp>2005-01-31T23:59:59.183Z</Timestamp>
  <Signature>Iuyz3d3P0aTou39dzbq7RrtSFmw=</Signature>
</GetObject>
```

# SOAP API Examples: GetObject (2 of 2)

- Sample response:

```
<GetObjectResponse xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <GetObjectResponse>
    <Status>
      <Code>200</Code>
      <Description>OK</Description>
    </Status>
    <Metadata>
      <Name>Content-Type</Name>
      <Value>text/plain</Value>
    </Metadata>
    <Metadata>
      <Name>family</Name>
      <Value>Muntz</Value>
    </Metadata>
    <Data>aGEtaGE=</Data>
    <LastModified>2005-11-18T01:25:07.994Z</LastModified>
    <ETag>&quot;828ef3fdfa96f00ad9f27c383fc9ac7f&quot;</ETag>
  </GetObjectResponse>
</GetObjectResponse>
```

# SOAP API Examples: DeleteObject

- Delete the “Nelson” object from the “quotes” bucket:

- Sample request:

```
<DeleteObject xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>quotes</Bucket>
  <Key>Nelson</Key>
  <AWSAccessKeyId> 1D9FVRAYCP1VJS767E02</AWSAccessKeyId>
  <Timestamp>2005-01-31T23:59:59.183Z</Timestamp>
  <Signature>Iuyz3d3P0aTou39dzbq7RrtSFmw=</Signature>
</DeleteObject>
```

- Sample response:

```
<DeleteObjectResponse xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <DeleteObjectResponse>
    <Code>200</Code>
    <Description>OK</Description>
  </DeleteObjectResponse>
</DeleteObjectResponse>
```

# SOAP API Examples: ListBucket (1 of 2)

- List up to 40 keys in the “quotes” bucket that have the prefix “N” and occur lexicographically after “Ned”:

- Sample request:

```
<ListBucket xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>quotes</Bucket>
  <Prefix>N</Prefix>
  <Marker>Ned</Marker>
  <MaxKeys>40</MaxKeys>
  <AWSAccessKeyId>1D9FVRAYCP1VJS767E02</AWSAccessKeyId>
  <Timestamp>2005-01-31T23:59:59.183Z</Timestamp>
  <Signature>Iuyz3d3P0aTou39dzbq7RrtSFmw=</Signature>
</ListBucket>
```

# SOAP API Examples: ListBucket (2 of 2)

– Sample response:

```
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/"> <Name>quotes</Name>
  <Prefix>N</Prefix>
  <Marker>Ned</Marker>
  <MaxKeys>40</MaxKeys>
  <IsTruncated>>false</IsTruncated>
  <Contents>
    <Key>Nelson</Key>
    <LastModified>2005-11-17T07:13:48Z</LastModified>
    <ETag>"828ef3fdfa96f00ad9f27c383fc9ac7f"</ETag>
    <Size>5</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>bcaf1ffd86f41caff1a493dc2ad8...</ID>
      <DisplayName>webfile</DisplayName>
    </Owner>
  </Contents>
  <Contents>
    <Key>Neo</Key>
    <LastModified>2005-11-17T07:13:48Z</LastModified>
    <ETag>"828ef3fdfa96f00ad9f27c383fc9ac7f"</ETag>
    <Size>4</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>bcaf1ffd86f41caff1a493dc2ad8c2c281e37522a640e161ca5fb16fd081034f</ID>
      <DisplayName>webfile</DisplayName>
    </Owner>
  </Contents>
</ListBucketResult>
```

## S3 Libraries for SOAP API

- Amazon provides libraries in the following languages for working with the SOAP API:
  - Ruby, Perl, Python, Java, and C#
- See <http://developer.amazonwebservices.com/connect/kbcategory.jspa?categoryID=47> to download libraries



# Sources

- S3 Homepage (<http://aws.amazon.com/s3>)
- S3 Developer Guide  
(<http://docs.amazonwebservices.com/AmazonS3/2006-03-01/>)
- S3 Getting Started Guide  
(<http://developer.amazonwebservices.com/connect/servlet/KbServlet/download/123-102-104/s3-gs-20060301.zip>)

# Demo

- S3Driver Java class