



STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2015

# What You Need to Know on Cloud Storage

**David Slik, Technical Director, NetApp, Inc.**

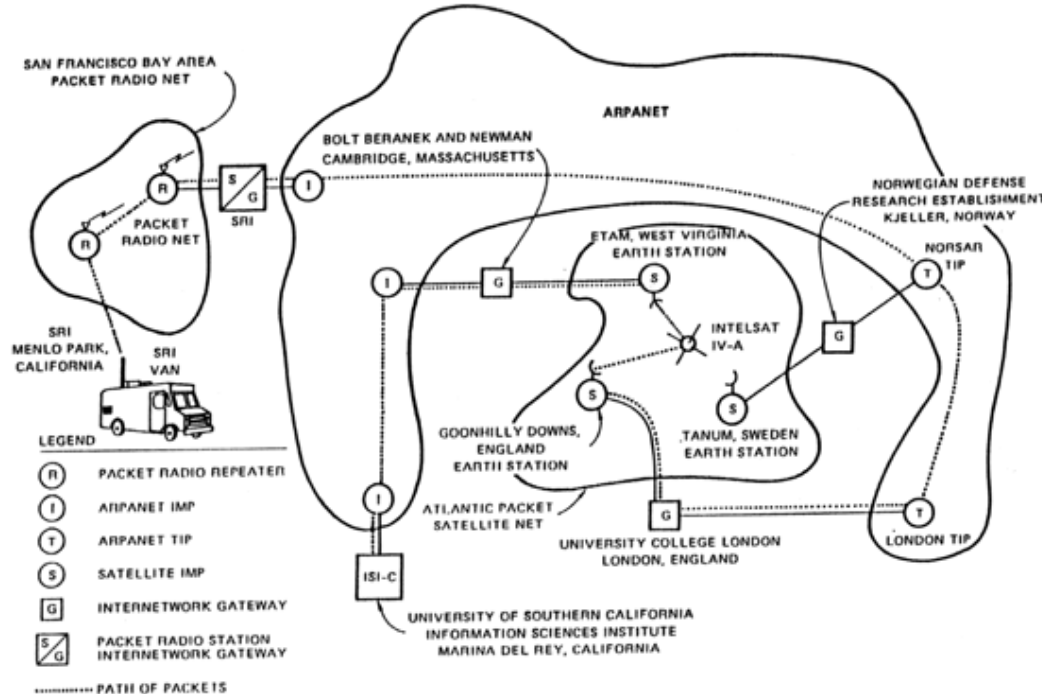
**Mark Carlson, Principal Engineer, Industry Standards, Toshiba**

# Session Agenda

- Part 1
  - What is Cloud Storage?
  - Introduction to HTTP
  - Introduction to REST
- Break
- Part 2
  - Introduction to CDMI

# What is Cloud Storage?

- The use of the term “cloud” arose from network architecture drawings that abstracted away details that didn’t matter so that could be ignored.



The first known “Cloud” diagram was used in 1977 to abstract away the internal structure of ARPANET.

# What is Cloud Storage?

- ❑ Thus, “Cloud Storage” is the abstraction of the complexity of storage to the point where you can largely ignore it’s implementation, location and operation.
- ❑ Some key differentiators:
  - ❑ Typically operated as a service (internally or externally)
  - ❑ Single access point (an HTTP URL is typical)
  - ❑ Self-service, self-provisioning and billed on usage

# What is Cloud Storage?

- ❑ Examples of uses of Cloud Storage have grown to encompass virtually all storage applications, but are especially common for:
  - ❑ Internet-based media storage (photos, videos, etc.)
  - ❑ Web document repositories
  - ❑ Backup
  - ❑ Sync and share
  - ❑ Third platform applications

# What is Cloud Storage?

- ❑ What does a typical cloud storage experience look like?
  - ❑ 1. Select provider
  - ❑ 2. Sign up (provide credit card)
  - ❑ 3. Create endpoint using provider web interface
  - ❑ 4. Test using Cloud Storage API to store and retrieve
  - ❑ 5. Integrate Cloud API with application
  - ❑ 6. Application uses Cloud API for store and retrieve
  - ❑ 7. Pay based on usage
  
- ❑ This can be done in less than a day (or even an hour)!

# What is Cloud Storage?

Click,  
pay,  
and use!



This is radical simplicity compared to traditional IT storage procurement and provisioning

The screenshot shows a payment method selection form. The heading is "Please Select a Payment Method". Below the heading, there is a paragraph of text: "All your future AWS service charges will be billed to the payment method you select. Should you need to do so, you will be able to change to another payment method in the future." Below this text, there is a section titled "Enter a Credit Card Below". This section contains four input fields: "Credit Card" (a dropdown menu with "Visa" selected), "Credit Card Number" (a text input field), "Expiration Date" (two dropdown menus for month and year, with "01" and "2009" selected), and "Cardholder's Name" (a text input field). A yellow "Continue" button with a mouse cursor over it is located below the input fields.

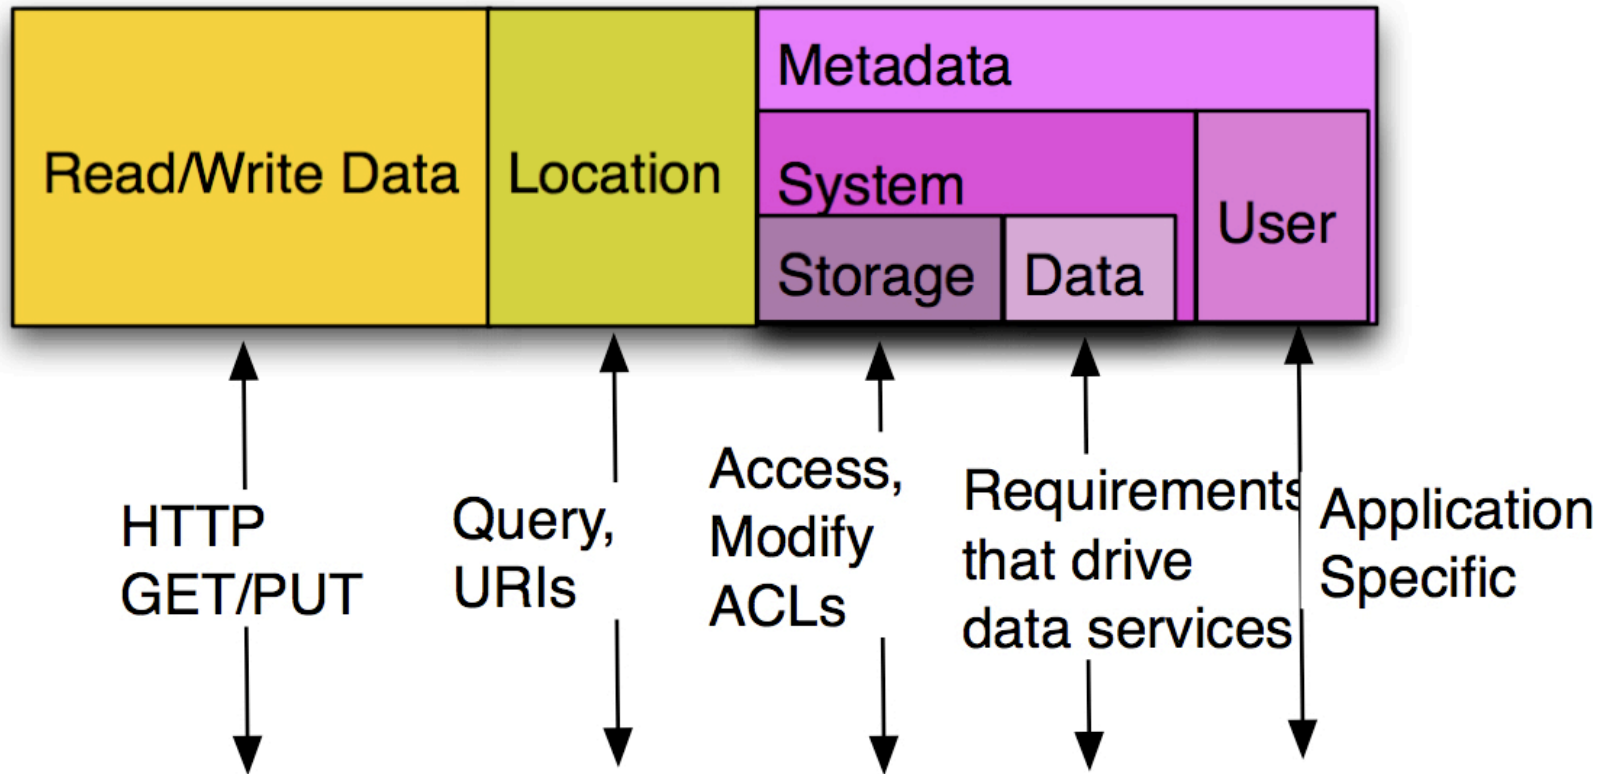
# What is Cloud Storage?

- ❑ What does a typical cloud storage experience look like?
  - ❑ 1. Select provider
  - ❑ 2. Sign up (provide credit card)
  - ❑ 3. Create endpoint using provider web interface
  - ❑ 4. Test using Cloud Storage API to store and retrieve
  - ❑ 5. Integrate Cloud API with application
  - ❑ 6. Application uses Cloud API for store and retrieve
  - ❑ 7. Pay based on usage
  
- ❑ This can be done in less than a day (or even an hour)!



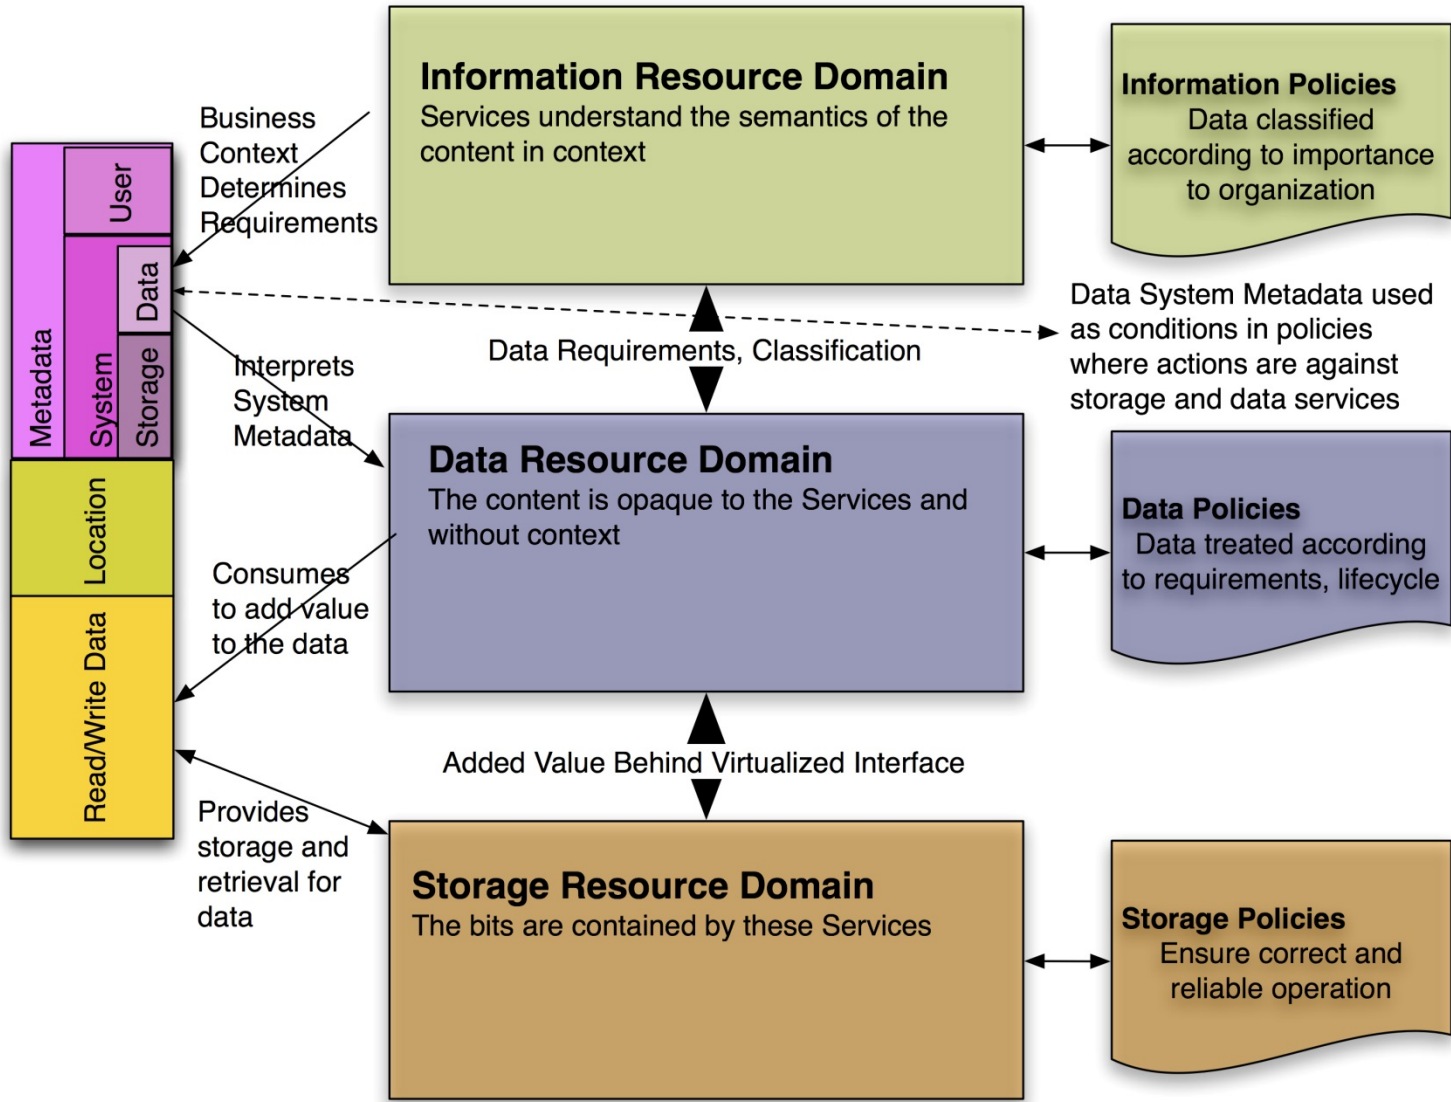
# What is Cloud Storage?

- How does this map to the standard storage model?



# What is Cloud Storage?

- ❑ Some important concepts
  - ❑ Metadata – Information about the data stored.
  - ❑ Examples of metadata include:
    - ❑ The date a photo was taken on
    - ❑ The author of a document
    - ❑ The number of times a file was accessed
    - ❑ Who is allowed to access a document
  - ❑ Policies – Directives controlling how storage is managed



# Introduction to HTTP

- ❑ HTTP is the protocol that underpins the web
- ❑ HTTP runs over TCP/IP, and is based around three main concepts:
  - ❑ Verbs – Commands that tell a server what to do
  - ❑ Headers – Key/Value pairs used to provide control and metadata signaling
  - ❑ Bodies – The contents of the operation
- ❑ In HTTP, every operation (verb) is performed against a “resource”, which exists at a named location (URL).
  - ❑ Operations involve a “representation” of a given resource.
  - ❑ For example, a stored image may be retrieved as a JPEG or PNG (two representations of the same resource)

# Introduction to HTTP

## DEMONSTRATION OF BASIC HTTP CONCEPTS

# Introduction to REST

- ❑ **REpresentation State Transfer**
  - ❑ Started with a dissertation by Roy Fielding outlining the principles
- ❑ **Addressability**
  - ❑ Every object (resource) is addressable through a unique identifier
- ❑ **Uniform, Constrained Interface**
  - ❑ Use HTTP verbs and model other semantics in the data model
  - ❑ Allows for Familiarity (low learning curve), Interoperability and Scalability
- ❑ **Representation Oriented**
  - ❑ Complexity is in the representations
- ❑ **Communicate Statelessly**
  - ❑ No persistent client-server connections

# Introduction to REST

- ❑ Why are RESTful APIs widely used?
  - ❑ Simplicity!
  - ❑ Safety — no side effects and operation idempotency.
  - ❑ Commonly available and built-in infrastructures in many programming languages and for many platforms
  - ❑ Low learning curve leads to developer adoption
  - ❑ Developer adoption creates eco-system around API
  - ❑ Eco-system eases adoption by vendors and customers
  - ❑ Well suited for scale-out implementations

# SDC 2015 Cloud Storage Primer

BREAK



# A Brief Overview of CDMI

- ❑ CDMI (Cloud Data Management Interface) is a RESTful API for accessing and managing cloud storage.
- ❑ The major cloud storage APIs are:
  - ❑ Amazon S3
  - ❑ CDMI
  - ❑ Microsoft Azure
  - ❑ Swift API (part of OpenStack)
- ❑ CDMI is widely implemented
  - ❑ >30 server implementations
  - ❑ CDMI gateways, OpenStack support

# A Brief Overview of CDMI

- ❑ 2009: SNIA Cloud Technical Working Group founded to explore API standardization
- ❑ 2011: CDMI 1.0 ratified as a US Technical Architecture
  - ❑ CDMI 1.0.1 errata released in late 2011
  - ❑ CDMI 1.0.2 errata released in mid 2012
- ❑ 2012: CDMI 1.0.2 becomes ISO/IEC 17826
- ❑ 2013: CDMI 1.1 under active development
  - ❑ 18 Extensions submitted
- ❑ 2015: CDMI 1.1.1 Submitted to ISO/IEC
  - ❑ 13 Extensions submitted, 5 incorporated



# A Brief Overview of CDMI

- ❑ Why does CDMI Matter?
  - ❑ Simple and easy to implement
    - ❑ Start with HTTP and add functionality, few mandatory parts
  - ❑ Advanced functionality not found in other APIs
    - ❑ Provides a foundation for next generation cloud services, such as federation
  - ❑ Open industry standard
    - ❑ Not controlled by any one vendor, protection against patents
  - ❑ Well defined formal standard
    - ❑ Enables interoperability, testing, and cross-vendor support
  - ❑ Widespread government support and adoption

# A Brief Overview of CDMI

- ❑ CDMI Standardizes:
  - ❑ CRUD operations (Create/Read/Update/Delete)
  - ❑ Data, Container, Queue and Domain objects
  - ❑ Identity and access control model
  - ❑ Metadata (including client and vendor extensibility)
  - ❑ Query and Notifications
  - ❑ Versioning
  - ❑ Serialization and Deserialization
  - ❑ Interoperability with other NAS and cloud protocols

# A Brief Overview of CDMI

- ❑ The CDMI standard can be downloaded from SNIA's web site:
  - ❑ <http://www.snia.org/cdmi>
- ❑ Extensions are defined to extend CDMI
  - ❑ [http://www.snia.org/tech\\_activities/publicreview/cdmi](http://www.snia.org/tech_activities/publicreview/cdmi)
- ❑ We will be referring to the CDMI specification for the remainder of this session, so if you have your laptop with you, you'll want to download the CDMI 1.1.1 technical position document to follow along.

# Introduction to HTTP

## DEMONSTRATION OF BASIC CDMI CONCEPTS

Thank you!

Questions

[dslik@netapp.com](mailto:dslik@netapp.com)

[Mark.Carlson@taec.toshiba.com](mailto:Mark.Carlson@taec.toshiba.com)