



Education

CLOUD STORAGE

The New Paradigm for Accessing Storage as a Service

Ashvin Kamaraju, Symantec
Philippe Nicolas, KerStore

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced without modification
 - ◆ The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the Author nor the Presenter is an attorney and nothing in this presentation is intended to be nor should be construed as legal advice or opinion. If you need legal advice or legal opinion please contact an attorney.
- The information presented herein represents the Author's personal opinion and current understanding of the issues involved. The Author, the Presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

➤ Cloud Storage – A new paradigm for accessing storage as a service

- ◆ Cloud, Cloud Computing, Cloud Services or Cloud Storage, just to name a few variations, has the attention of many IT professionals, developers, marketers, press and analysts. Why? Simply because it is perceived as a major shift in IT trends.

This tutorial will provide an overview of the market dynamics causing the paradigm shift to cloud services, describe the predominant use cases for cloud storage and access methods using REST, SOAP and WSDL.

This session will appeal to data center managers, development managers, application developers and those that are seeking a fundamental understanding of a service oriented approach to accessing storage.

- Why this Paradigm Shift ?
- Definitions & Conventions
- Cloud Storage
- Conclusion

Why this Paradigm Shift ?

IT Challenges and Business Needs

What is driving the paradigm shift to Cloud Services?

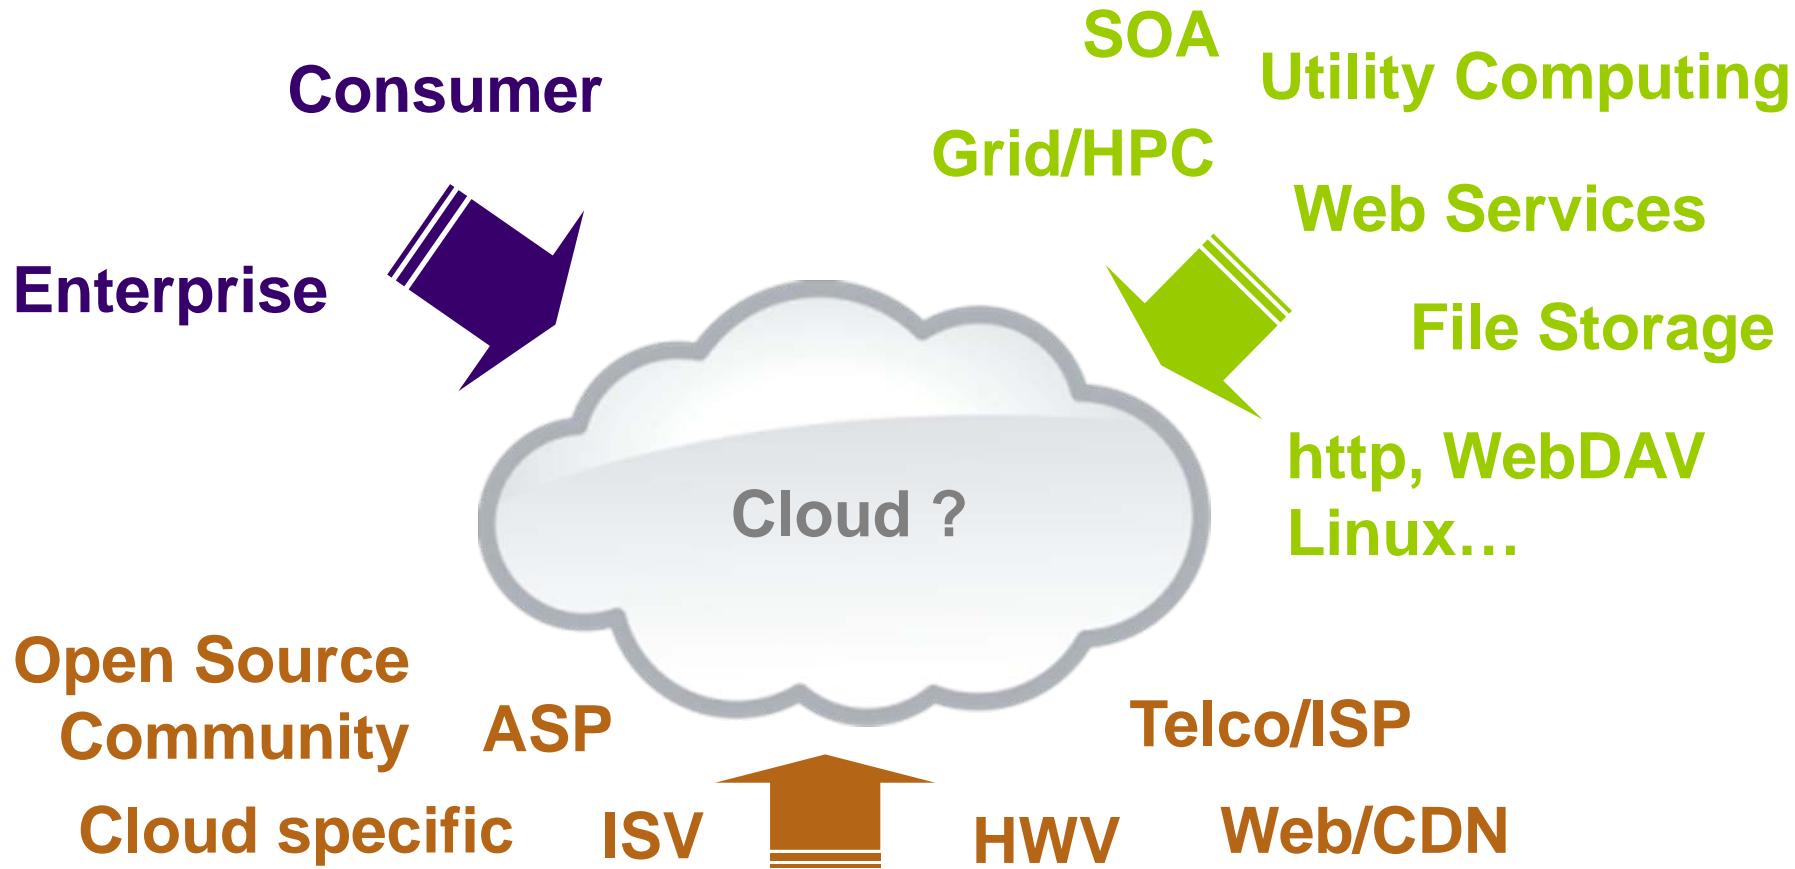
- Exponential growth of unstructured data
- Compelling economics (low cost)
- On demand access
- Compliance
- Ability to bootstrap development very quickly
- Focus on core (business) competency
- Green IT

Definitions & Conventions

Market dynamics & Root concept
What it is ? What it is not ?

Market dynamics

- **New buzz word in IT industry and on the web**
 - ◆ Several announcements every day, hundreds of online communities and forums
 - ◆ Vendors rename their product/solutions/strategy
 - ◆ New players
- **Obvious inflection point**
 - ◆ Many IT users like the flexibility to switch from “traditional IT purchasing and maintenance” to “doing more with IT services (SLA/SLO)”
 - ◆ Shift in IT spending and huge move in the next 5-7 years
- **Many examples...**
 - ◆ eBay, Amazon, YouTube, Salesforce.com, Gmail, Paypal, Facebook, Zoho, Flickr, Skype, Google, XDrive...



A World of Convergence

- Cloud is a metaphor for the internet
- Similar to services such as cable TV and power, cloud services are:
 - ◆ Accessed over a network
 - ◆ Consumed according to demand
 - ◆ Metered and billed for the amount consumed
- Typical services are:
 - ◆ Applications
 - ◆ Storage
 - ◆ Compute

**Cloud Computing is a model
to deliver scalable, reliable, cost-effective
and on-demand
« IT as a service »
using Internet technologies**

Ultimate mode of Virtualization

- SLA-based management
- Shift from CapEx to OpEx
- Massively scalable and geo distribution
- Global, universal and ubiquitous information access
- Resiliency and security
- Data center abstraction with virtualization technologies

Key denominators: Simple, Fast, Reliable





Cloud Storage

Models, Topologies Applications & Technologies

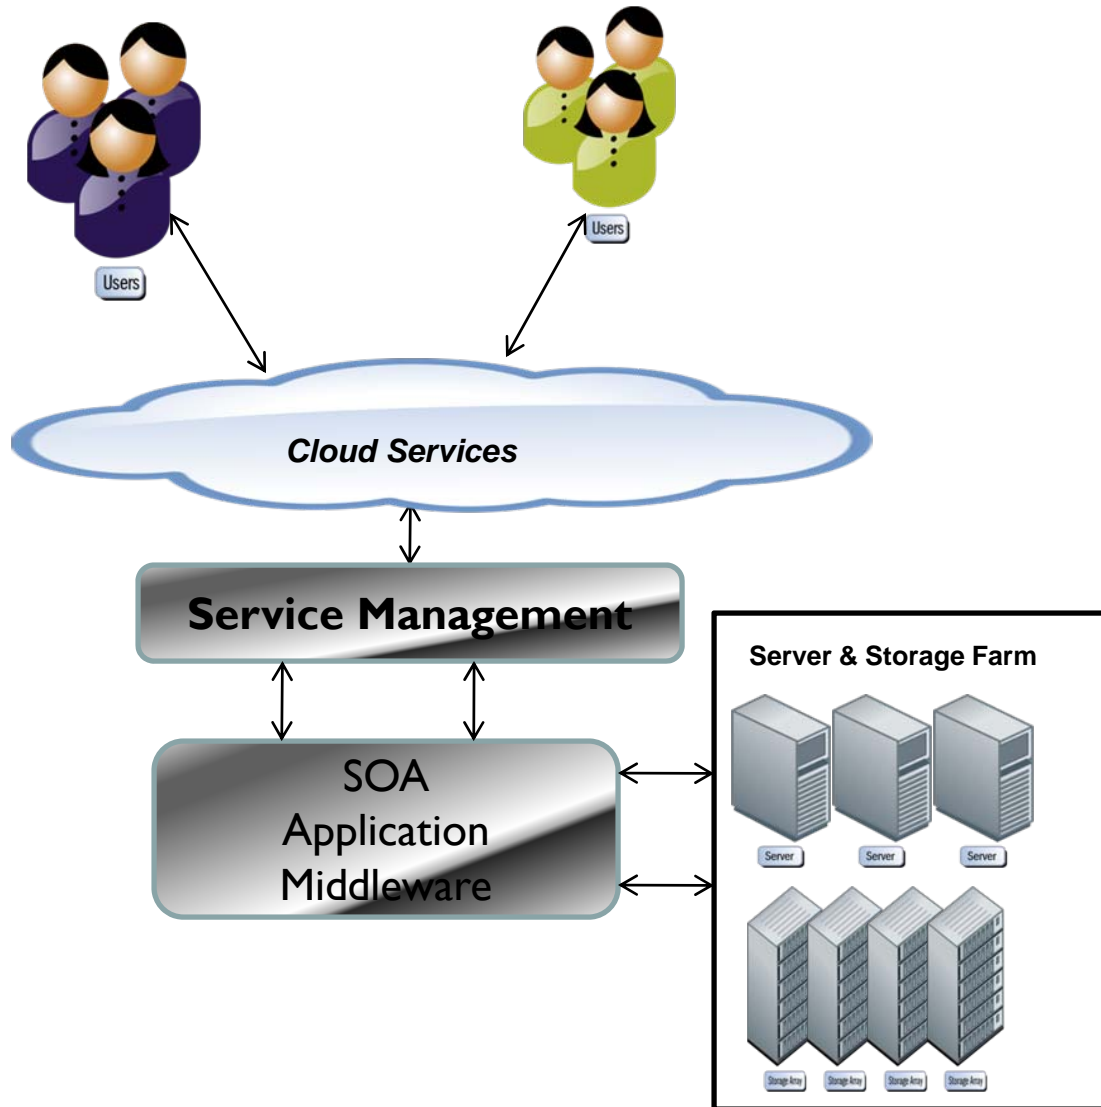
Cloud Storage

- Located in a public or private data center, separate from primary storage.
- Implemented using a Service Oriented Architecture (SOA)
- Accessed as a service:
 - ◆ Directly as blocks or files
 - ◆ Indirectly, through applications that are also co-located with the storage.
- Different from Cloud Computing where virtual machines/operating systems & applications are accessed as a service.

Cloud Storage Models

<p>Storage Cloud (cloud-attached storage)</p>	
<p>Application + Storage in the same cloud</p>	
<p>Application + Storage (On premise) + Secondary Storage for Backup, Archiving or DR (Off Premise)</p>	
<p>Application Cloud connected to Storage Cloud</p>	

Cloud Storage Use Cases - 1



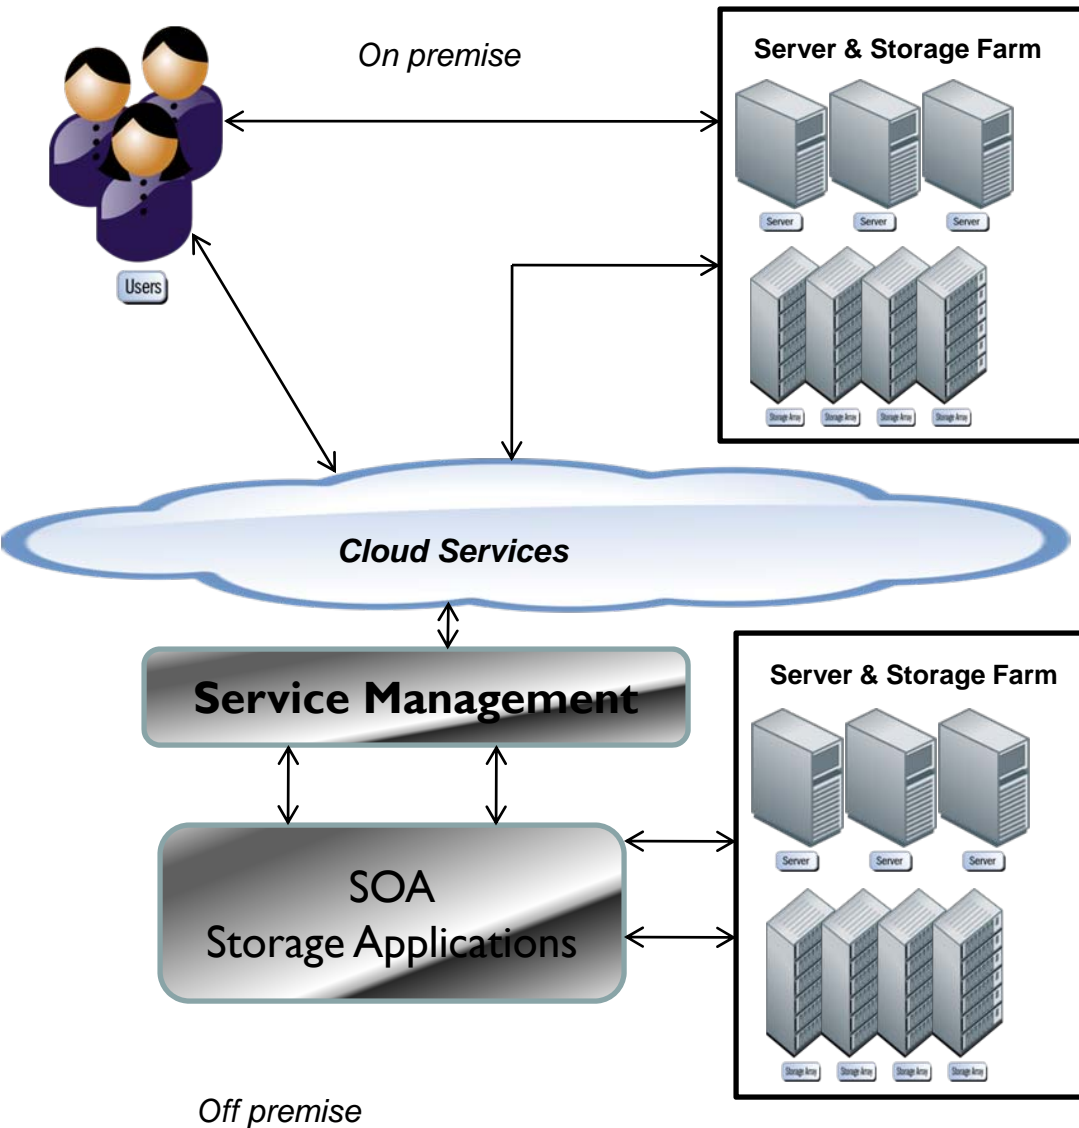
Applications co-located with storage in same cloud

- Applications that require high performance storage access
 - Enterprise applications e.g. CRM, HR, Financial
 - High performance tier-1 storage
- Applications that manage consumer digital storage
 - Pictures, videos , music
 - Lower performance tier-2/tier3 storage
- Software as a Service (SaaS)



Complete application stack offload

Cloud Storage Use Cases - 2



Storage intensive applications – On site/Off Site Data

- Applications co-located with off site storage with thin SW agents on site.

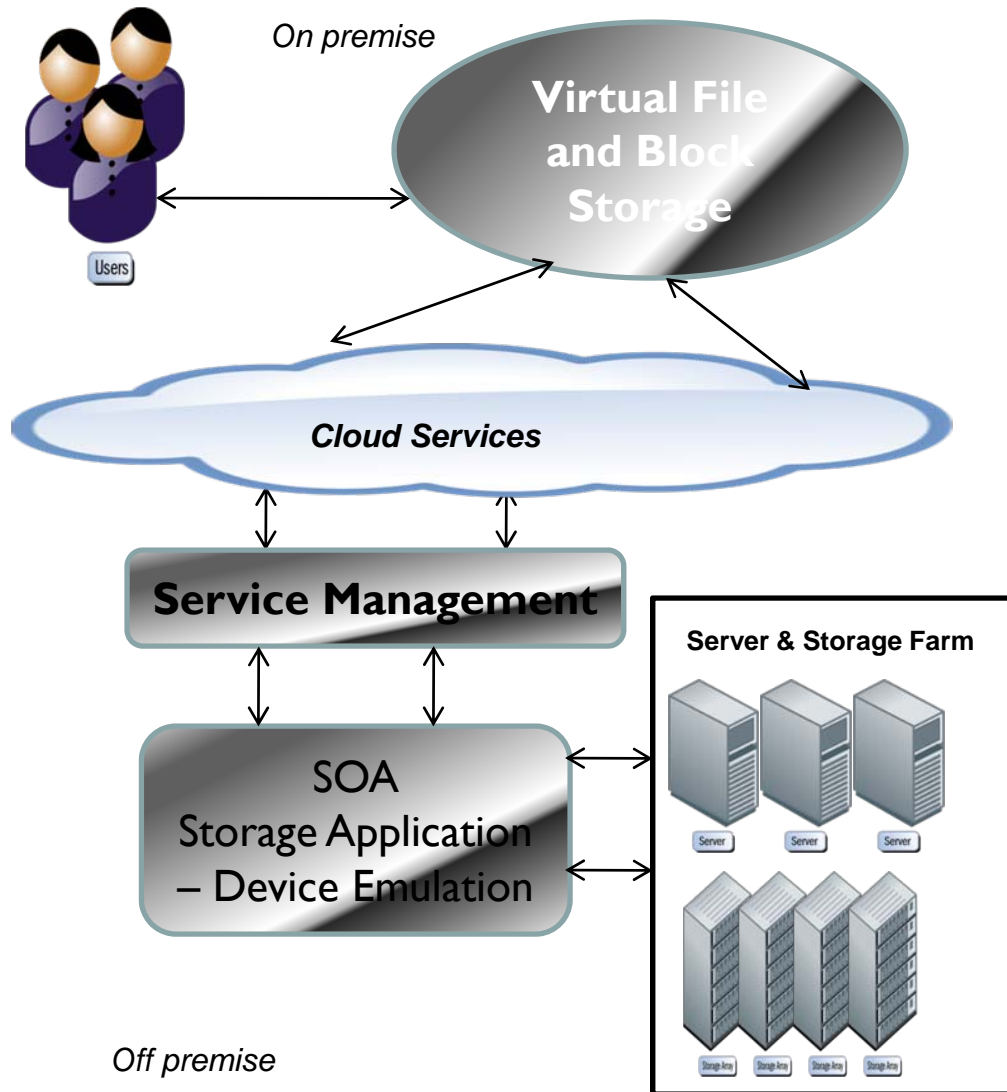
- Backup, archiving, DR, medical images

- Less performant tier-2/tier3 storage

- Software as a Service (SaaS)

- Mission critical data remains on site. Reduces costs for secondary storage.

Cloud Storage Use Cases - 3



File or Block Storage offered as a service

- Applications are NOT co-located with storage
- “Z drive” accessed using standard file or block access protocols.
- Description of storage capabilities with simple access methods for web applications (SOAP/REST)
- Tier3 storage
- Storage as a Service

Cost, Convenience and Capacity on demand.

- 1 Cheap & abundant power for storage & server farms and cooling
- 2 Network bandwidth to meet the service demands
- 3 Commodity off-the-shelf servers and storage
- 4 Ability to survive disasters (secondary sites)
- 5 Low administrative costs

Low and competitive burdened cost

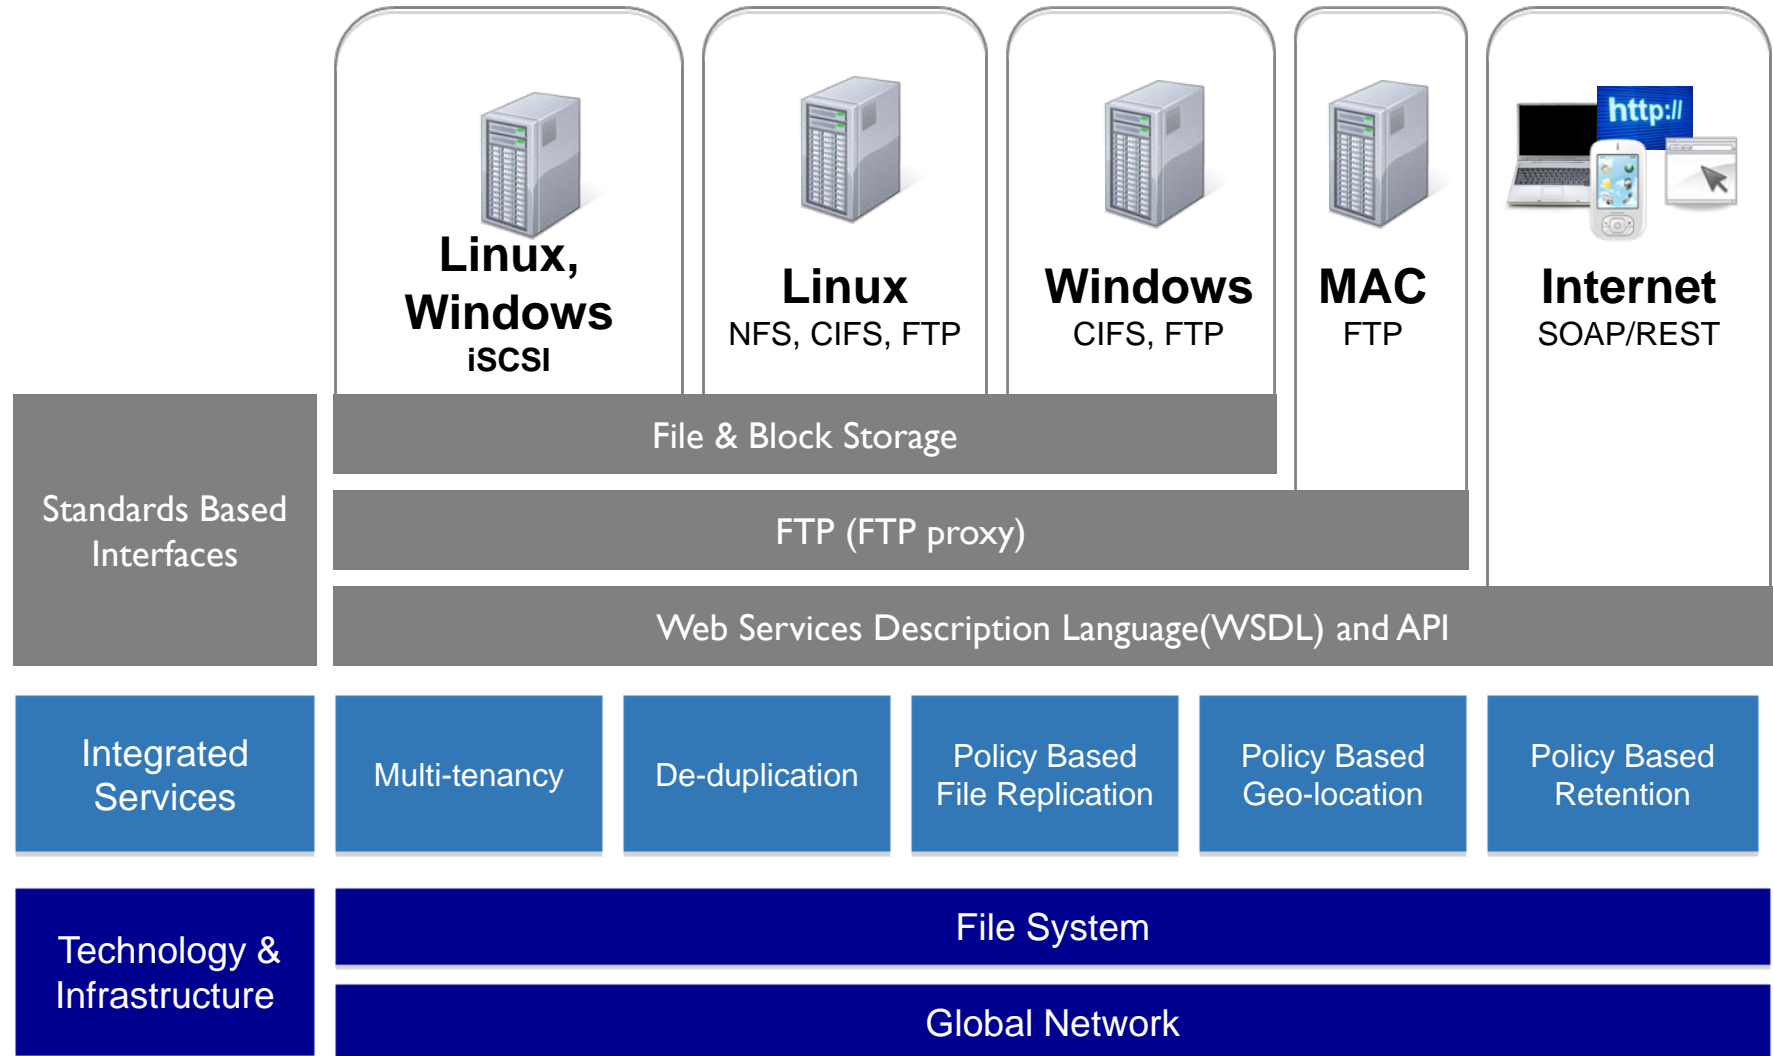
Burdened cost = Acquisition cost + Operating Cost (\$/GB/Month)

- 1 Scalable distributed/cluster file storage
Scalable block storage
- 2 Global namespace
- 3 Widely accepted data access protocols
- 4 Policy-based data management



Strong SLAs and simplified administration
SLA = Resiliency + Ubiquity + Performance by Simplification

Cloud Storage Access Protocols



➤ Representational State Transfer (REST)

- ◆ Network architecture principles or rules.
- ◆ Application state & functionality abstracted into “resources”.
- ◆ Protocol defines operations on resources and data formats.
- ◆ Application state changed by reading a resource.
- ◆ Stateless (= Scalable)
- ◆ HTTP to transfer information.

Stock quote example using REST

Resource: Stock Symbol for XYZ corporation.

Resource locator: <http://stockprice.com/stockprice/xyz>

Operation: Getquote

Messages:

The request

GET /StockPrice/XYZ HTTP/1.1

Host: stockprice.com

Accept: text/xml

Accept-Charset: utf-8

The response:

HTTP/1.1 200 OK

Content-Type: text/xml; charset=utf-8

Content-Length: nnn

<?xml version="1.0"?>

<s:Quote xmlns:s="http://stockprice.com/stock-service">

<s:TickerSymbol>XYZ</s:TickerSymbol>

<s:StockPrice>45.25</s:StockPrice>

</s:Quote>

➤ Simple Object Access Protocol (SOAP)

- ◆ Specification for exchanging structured information
- ◆ Relies on an application layer protocol (e.g. RPC or HTTP).
- ◆ XML for message format.
- ◆ Language, Platform agnostic.

Stock quote example using SOAP

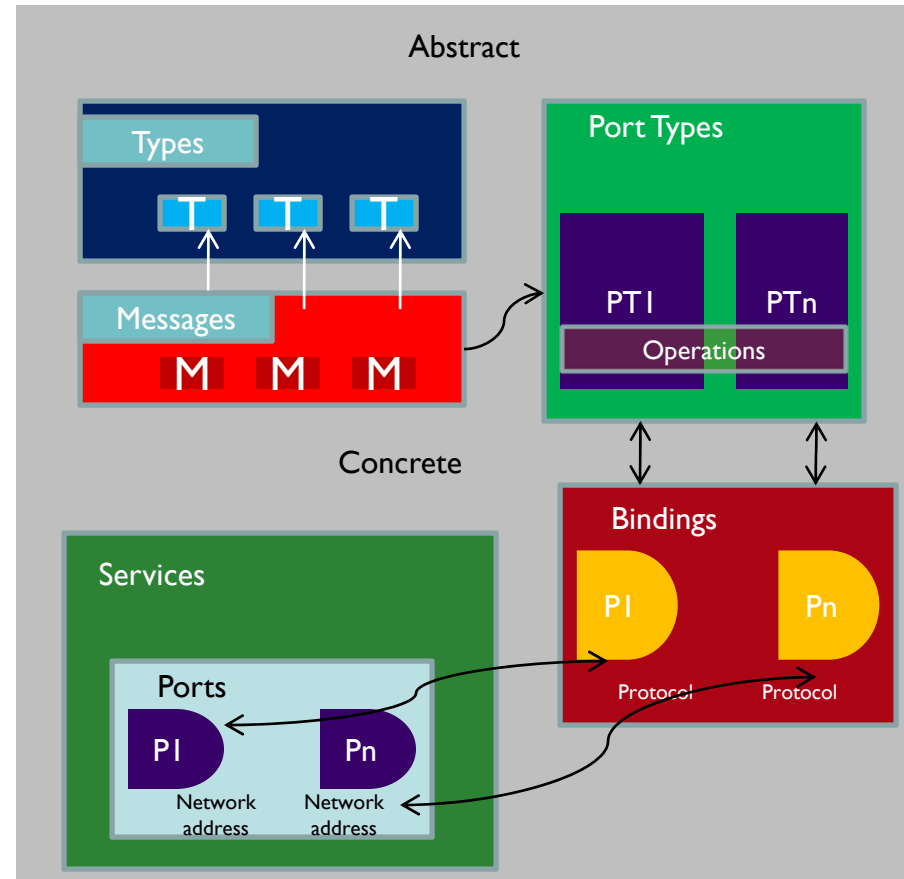
The request:
GET /StockPrice HTTP/1.1
Host: stockprice.com
Content-Type: application/soap+xml; charset=utf-8
Content-Length: nnn
<?xml version="1.0"?>
<env:Envelope xmlns:env="http://www.w3.org/2003/05/soap-envelope"
xmlns:s="http://www.stockprice.com/stock-service">
 <env:Body>
 <s:GetStockQuote>
 <s:TickerSymbol>XYZ</s:TickerSymbol>
 </s:GetStockQuote>
 </env:Body>
</env:Envelope>

The response:
HTTP/1.1 200 OK
Content-Type: application/soap+xml; charset=utf-8
Content-Length: nnn
<?xml version="1.0"?>
<env:Envelope xmlns:env="http://www.w3.org/2003/05/soap-envelope"
xmlns:s="http://www.stockprice.com/stock-service">
 <env:Body>
 <s:GetStockQuoteResponse>
 <s:StockPrice>45.25</s:StockPrice>
 </s:GetStockQuoteResponse>
 </env:Body>
</env:Envelope>

WSDL – Describing the web services

➤ Web Services Description Language (WSDL):

- ◆ An XML based language that provides a model for describing Web Services.
- ◆ Services are defined in terms of:
 - › **Port Type, Message, Data Type , Bindings and Service elements**
- ◆ Typically used with SOAP protocol, although Restful architectures can be supported with WSDL2.0.



➤ Sample Storage Service

➤ Storageservice.com

- ◆ Common operations:
 - > Create a Container
 - > List Container Content
 - > Write a File
 - > Read a File
 - > Delete a File

Get file example

Get file "photo345" from container "myimages" hosted by service on storageservice.com

REST:

*GET /photo345 HTTP/1.1
Host: myimages.storageservice.com*

SOAP:

```
<GetFile
xmlns="http://dsc.storageservice.com/ver1">
<Container>myimages</Container>
<Name>photo345</Name>
</GetFile>
```

Barriers to adoption

- Cost savings not significant enough
- Data privacy concerns
 - ◆ E.g. Shared Tenancy Issues
- Data Migration/Integration Issues
- Data availability and SLAs that meet enterprise requirements

- Cloud storage has real benefits for
 - ◆ Online File Storage, Online Backup and Data Archiving, DR... A new tier of Storage etc. (Applications that are not performance critical)
- However, due to performance limitations OLTP/Transactional applications have to be co-located with storage in the cloud.
- Security and encryption must be rock solid for enterprise customers to overcome their concerns.
- Industry standards are needed for Cloud Interoperability, Management, Data Migration/Exchange and Resources Provisioning...
- “Try and Adopt” with a multi-phase approach with a mix of Private and Public Clouds



Check out SNIA Tutorial:
The File Systems Evolution



Check out SNIA Tutorial:
Exploiting Multi-Tier File Storage Effectively



Check out SNIA Tutorial:
Virtualizing Servers, Virtualizing Storage – What’s the Difference?



Check out SNIA Tutorial:
Virtualization I - What, Why, Where and How?



Check out SNIA Tutorial:
Home Storage Explosion: Understanding the Forces Driving the Home Storage Market



Check out SNIA Tutorial:
Massively Scalable File Storage



Check out SNIA Tutorial:
pNFS, Parallel Storage for Grid and Enterprise Computing

- Please send any questions or comments on this presentation to SNIA
 - ◆ trackapplications@snia.org (Applications and Infrastructure)

**Many thanks to the following individuals
for their contributions to this tutorial.**

- SNIA Education Committee

**Phillipe Nicolas
Ashvin Kamaraju
Nenad Caklovic
Paul Massiglia
Jim Zierick
Clod Barrera
Andrew Reichman**