



Education

# WAN Optimization and Cloud Computing

Josh Tseng, Riverbed

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - ◆ Any slide or slides used must be reproduced in their entirety without modification
  - ◆ The SNIA must be acknowledged as the 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, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant 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.**

# Agenda Topics

- How the WAN is integral to Cloud Computing
- WAN-related challenges to Cloud Computing
- How WAN optimization addresses latency and bandwidth bottlenecks
- Issues with deploying WAN optimization into The Cloud

# What is a Cloud?

**Gartner**

Cloud computing is a style of computing where scalable and elastic IT-enabled capabilities are delivered as a service to external (Public) or to internal (Private) customers using Internet technologies.


**Bank of America**  
**Merrill Lynch**  
 RESEARCH

The Cloud is shared servers (resources) located within a data center delivering applications over internet technologies.



Cloud computing is a style of computing in which dynamically scalable and often virtualized resources are provided as a service over the Internet.


# The Beer Factory (Werner Vogels, CTO Amazon)



# Cloud/Utility Computing




Your Business



Shared Power  
Utility

## The “Cloud”



➤ The WAN is a core component of The Cloud

# Electrons and Photons

- The transmission “Network” allows you to place shared resources anywhere you want
  - ◆ Away from cities to avoid pollution/fallout
  - ◆ Near cheap resources
  - ◆ Where Gov’t regulations permit
- Shared data centers may be far away from end-users
  - ◆ WAN allows distant users to access central resources
  - ◆ Larger data centers can serve many users across distant geographies



# Elasticity and Economic Benefits

- **Cloud provider realizes quantities of scale**
  - ◆ Cloud provider gets large serving multiple customers
  - ◆ Per-unit provisioning costs decrease with size
- **Expand IT resources on demand**
  - ◆ Cloud provider provisions resources
  - ◆ Procurement of IT resources hidden from consumer
- **Cloud provider provides focus and expertise**
  - ◆ Focus creates superior expertise delivering IT resources
  - ◆ Superior expertise leads to improved IT performance and efficiency

# Private vs. Public Clouds

- **Private Cloud – You own the Data Center**
  - ◆ You own and manage the infrastructure and the data
  - ◆ You are responsible for security
  - ◆ You handle cost accounting
- **Public Cloud – 3<sup>rd</sup> party owns the Data Center**
  - ◆ Cloud Provider owns the infrastructure; you own the data
  - ◆ Cloud Provider provides for your security in their DC
  - ◆ Cloud Provider handles cost accounting

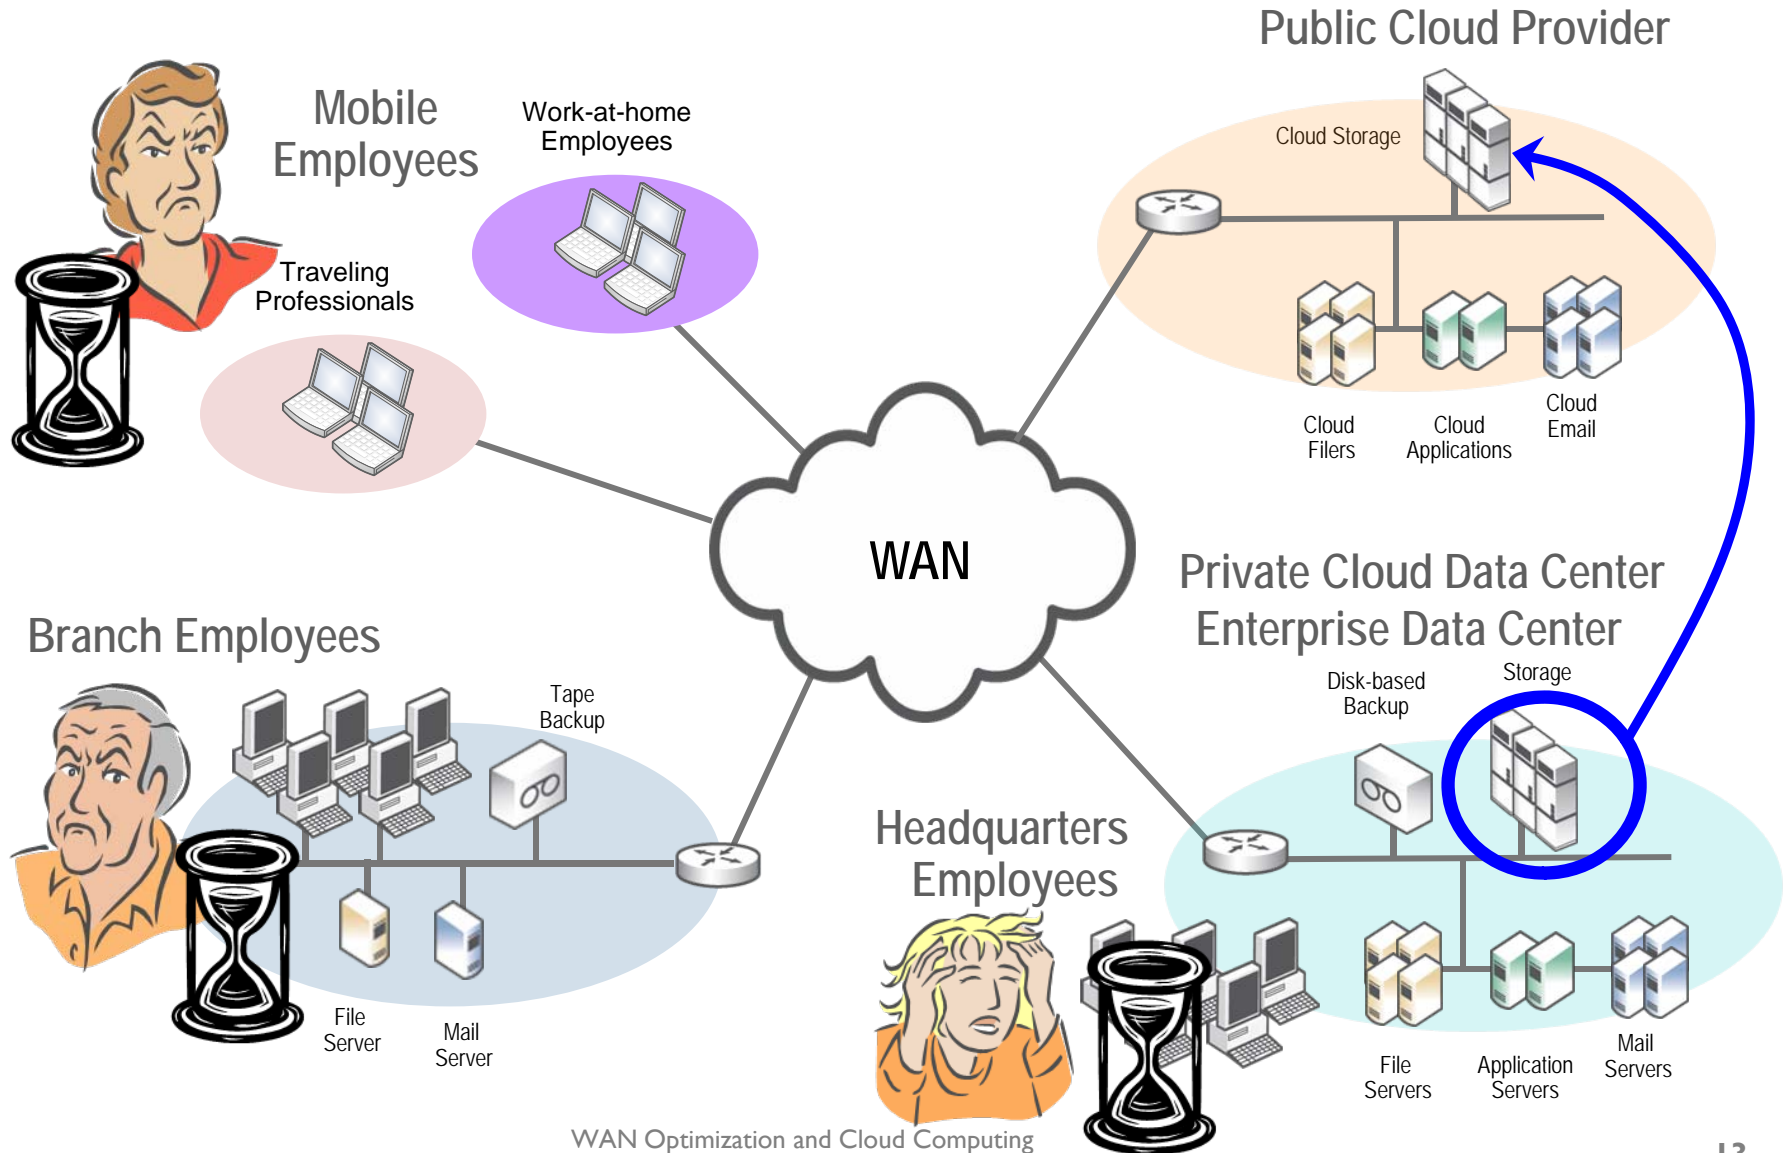
# Enabling Technologies for Cloud

- Consolidation & Virtualization
  - ◆ Sharing of hardware resources
- Billing and Provisioning
  - ◆ Appropriate cost accounting
- Security and Availability
  - ◆ Ensure your data remains accessible and safe
- WAN connectivity
  - ◆ Delivery of cloud applications to the end-user

# WAN-related challenges to Cloud

- Largely same physical WAN infrastructure used to access Cloud
  - ◆ Same WAN links with bandwidth and latency
- Cloud data centers are “far away”
  - ◆ Lower per-unit cost of large data centers driven by scale
- All employees will be “remote” from their data
  - ◆ Even single-location companies will be remote from their data
- Many legacy applications will continue use chatty protocols, even if moved to the Cloud
  - ◆ WAN latency will impact performance for many Cloud applications

# Building a Cloud Infrastructure



# WAN performance key to The Cloud

- Slow performance jeopardizes Cloud adoption
  - ◆ Reduced productivity
  - ◆ Employee frustration
- Cloud benefits negated by slow performance
  - ◆ Employees may associate Cloud Computing with slow performance
  - ◆ May reject Cloud deployment along with its benefits
- WAN performance must be addressed for Cloud adoption

# WAN Performance Bottlenecks


## The "distance" problem

- Latency and application chattiness
- Slower application access



## "Skinny Straw" problem

- Thousands of companies
- Millions of users
- Varied bandwidth



THE EXTREMELY UNFUNNY PART – UP TO 20x SLOWER


Action/ Latency	~ LAN (<10ms)	WAN (~100ms)	WAN (~200ms)
Login	2.08	23.08	34.54
"Reports" Tab	1.19	9.80	22.62
Upload 3MB doc	26.74	88.54	109.03
Download	10	38	40

Public and Private Clouds Share a Fundamental Pain

# Bottleneck #1: Bandwidth Limitations

- Lots of data needs to be sent over limited WAN bandwidth
- Congestion problems lead to miserable performance

- Files
- Email
- Web Apps
- Database
- Data Backup
- VOIP






# Fixing Bottleneck #1: Bandwidth Limitations

## ➤ Disk-based deduplication technology

- ◆ Identify redundant data at the byte level, not application (e.g., file) level
- ◆ Use disks to store vast dictionaries of byte sequences for long periods of time
- ◆ Use symbols to transfer repetitive sequences of byte-level raw data
- ◆ **Only** deduplicated data stored on disk

# Disk-based Data Reduction

**60 to 90 percent data reduction**



# Bottleneck #2: Application “Chattiness”

GET GIFs

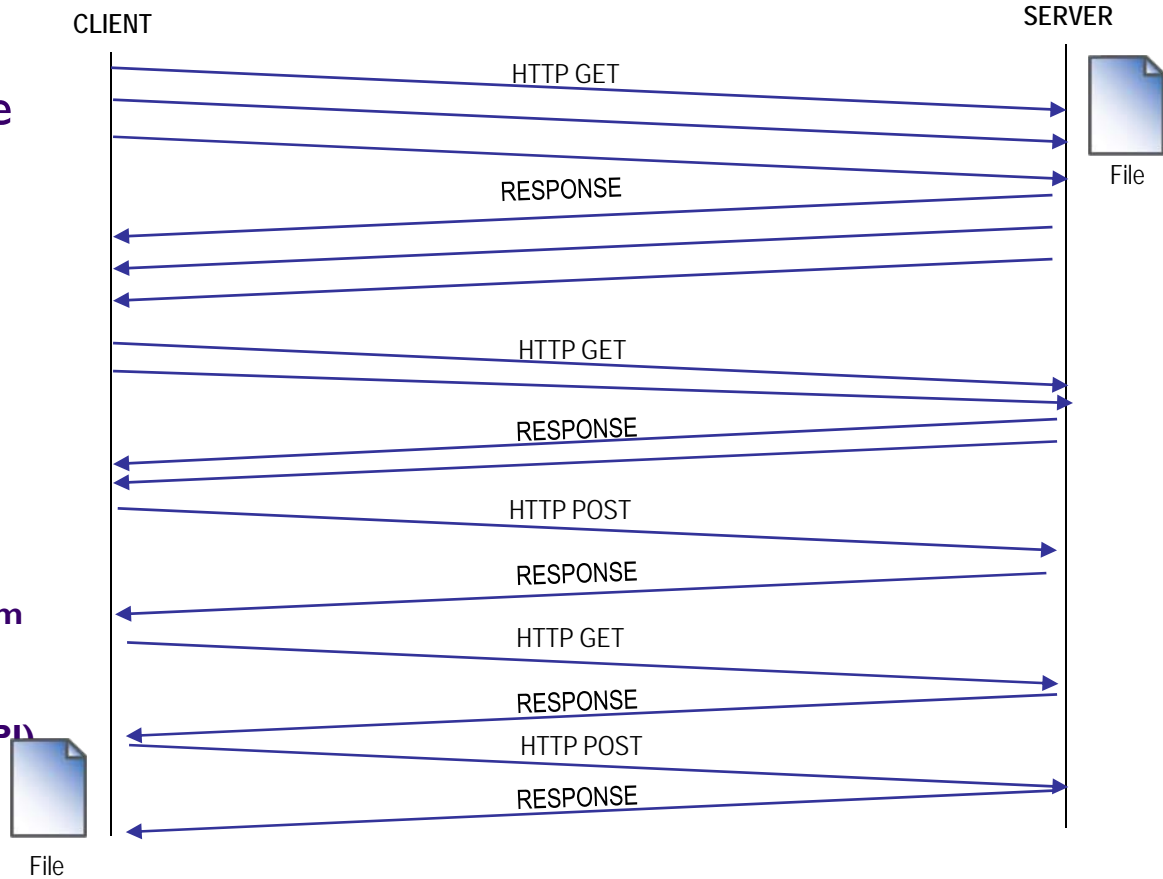
➤ Interactive apps, underlying protocols require 100s or 1000s of round trips for one operation!

- ◆ **Web-based applications**

- › CRM
- › Document Management
- › Call Center Apps
- › Project Mgmt Apps
- › Accounting Apps
- › Other Custom Apps

- ◆ **Legacy apps:**


- › Common Internet File System (CIFS)
- › Messaging Application Programming Interface (MAPI)
- › UNIX File Sharing (NFS)



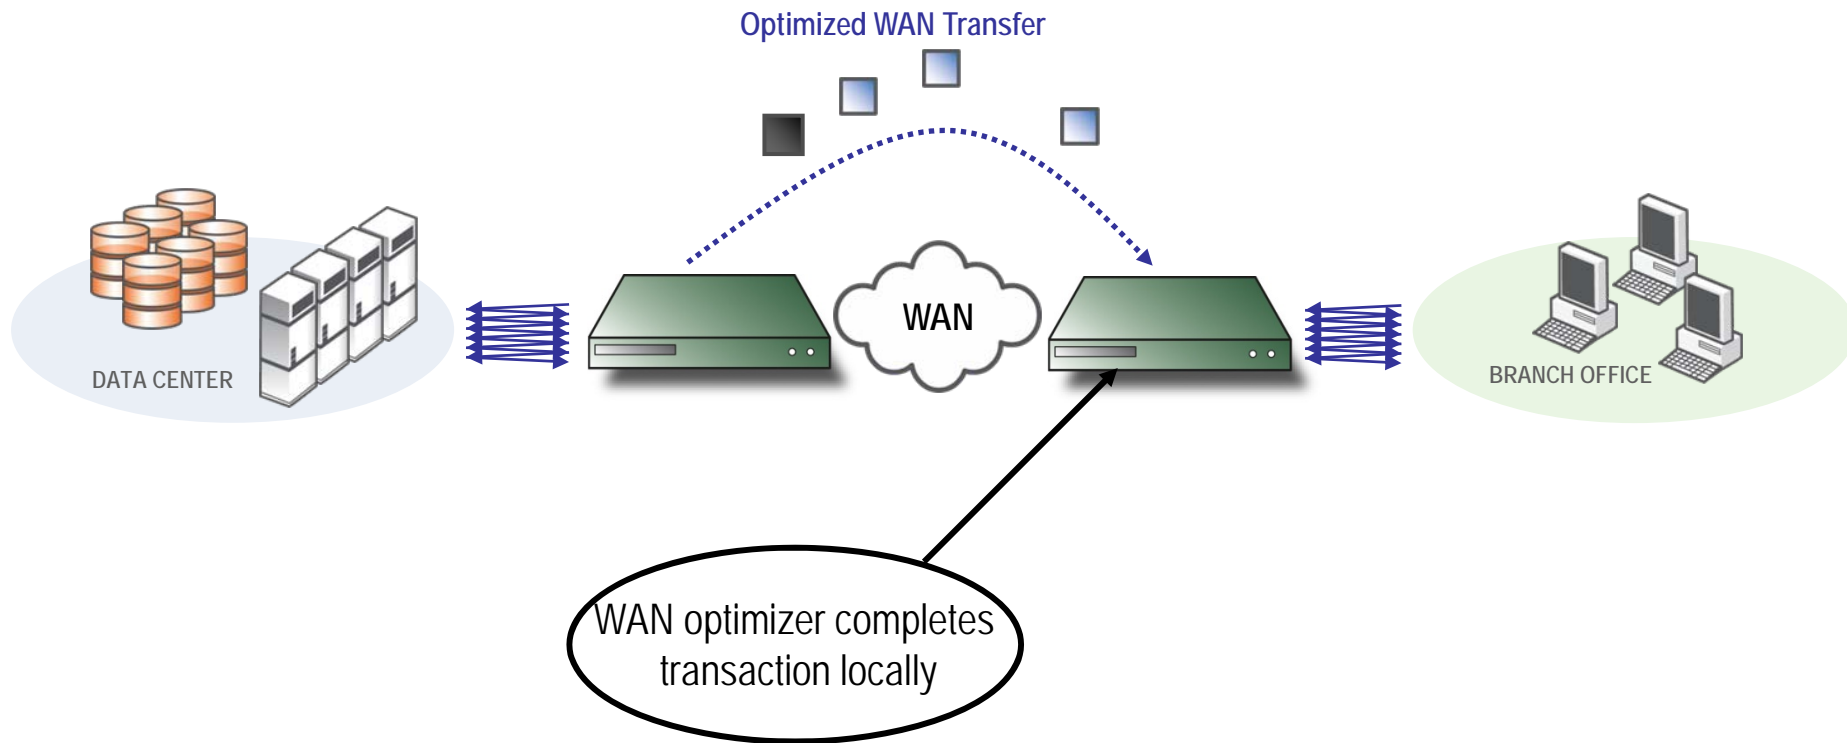
# Fixing Bottleneck #2: Application-Level Chattiness

- Application-specific chattiness mitigation modules
  - ◆ HTTP, CIFS, MAPI, MAPI2003, NFS, SQL, etc...
- Protocol-compliant read-aheads to pre-fetch data
  - ◆ Pipeline delivery of all application data
  - ◆ Eliminate chattiness over the WAN

# Addressing Application-Level Chattiness

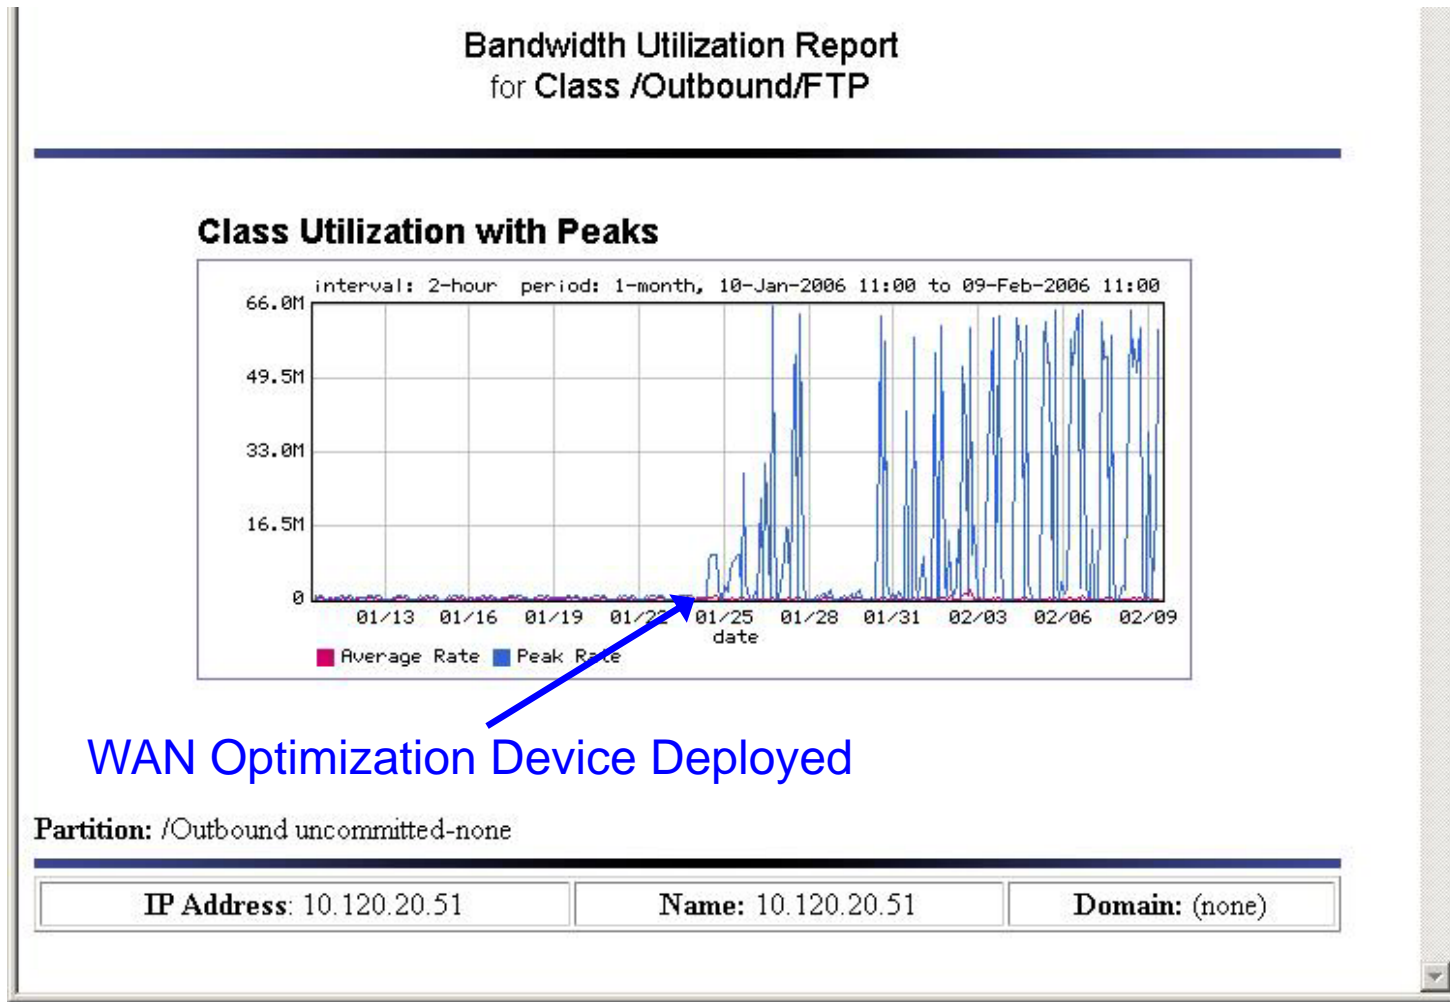


# Addressing Application-Level Chattiness




# WAN Optimization LAN-like performance

Atlanta to India E1 (2 Mbps) WAN connection (~150ms RT latency)



# WAN Optimization Bandwidth Reduction

## Traffic Summary



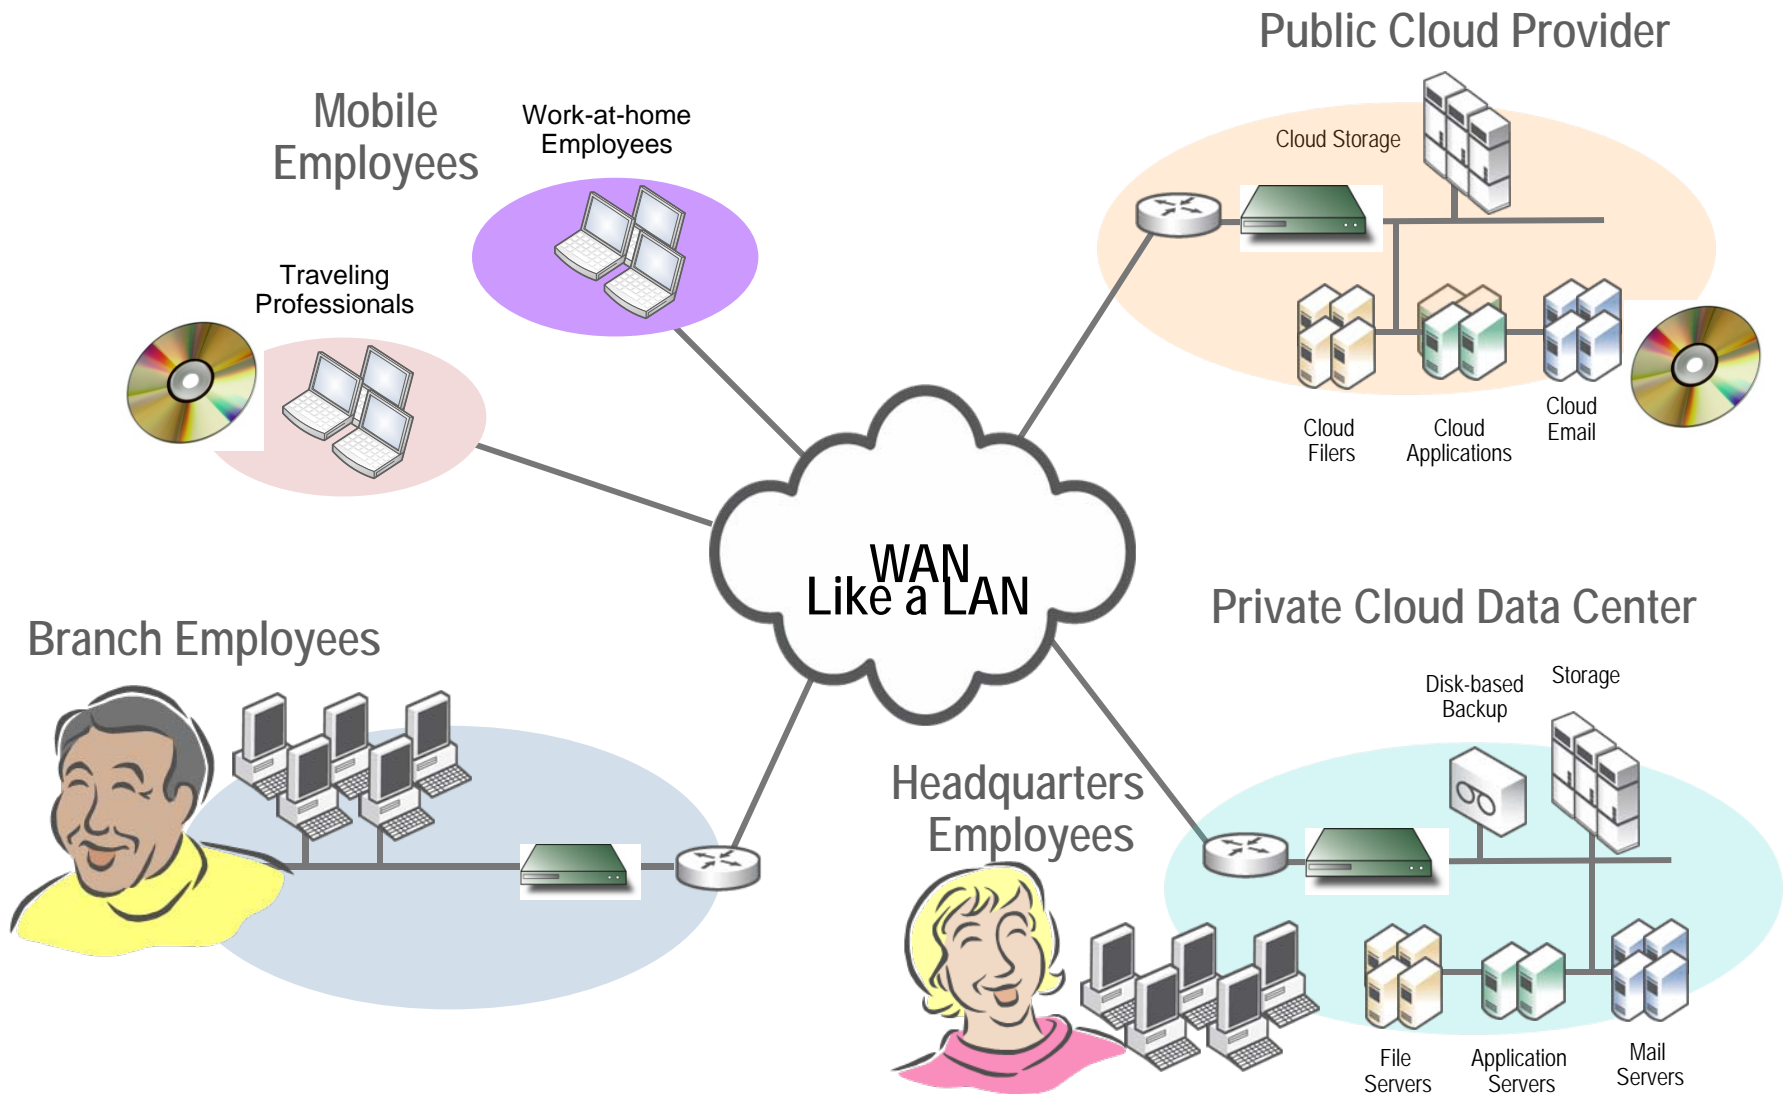
Port	Reduction	LAN	WAN	Traffic %
Total Traffic	--	78.9 GB	12.7 GB	--
HTTP (80)	(95.07%)	34.3 GB	1.6 GB	43.50%
JDE (8011)	(89.33%)	13.7 GB	1.4 GB	17.39%
email (1352)	(57.57%)	9.6 GB	4 GB	12.16%
DB (1521)	(60.84%)	7.3 GB	2.8 GB	9.34%
JDE (8003)	(89.47%)	4.7 GB	511.6 MB	6.01%
JDE (85)	(89.38%)	2.5 GB	279.8 MB	3.26%

SQL:TDS (1433)	(47.38%)	2.3 GB	1.2 GB	2.95%
JDE (8021)	(96.69%)	1.4 GB	48.5 MB	1.82%
Asset (8300)	(94.68%)	1001 MB	53.2 MB	1.24%
Unknown (1565)	(88.96%)	375.9 MB	41.4 MB	0.46%
JDE (8005)	(52.80%)	327.3 MB	154.5 MB	0.40%
CIFS:TCP (445)	(45.92%)	312.7 MB	169.1 MB	0.39%
SMTP (25)	(85.99%)	234.8 MB	32.9 MB	0.29%
FTP (21)	(84.02%)	150.1 MB	23.9 MB	0.19%
Unknown (1112)	(66.80%)	112 MB	37.2 MB	0.14%
Other	(67.6%)	371.0 MB	120.4 MB	0.43%

79GB of data was reduced to 13GB (83% reduced)  
66GB of data was removed from the International links at Malaysia

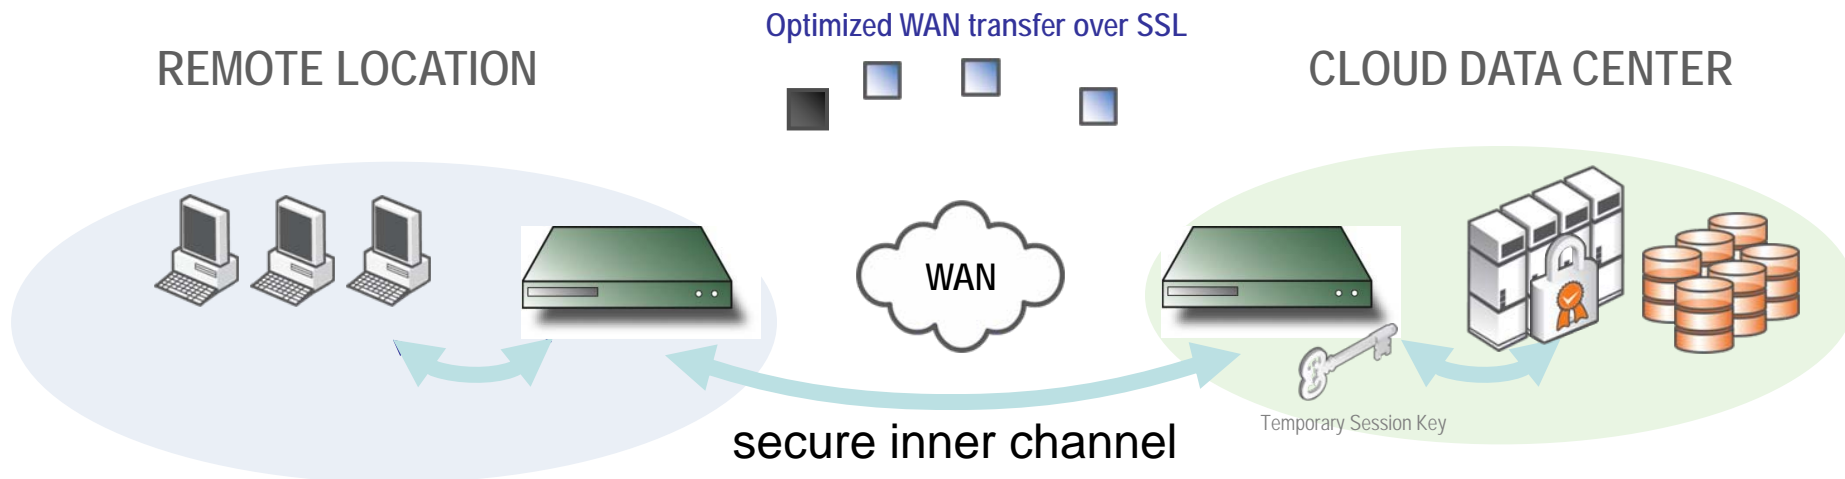


# Solving the WAN Performance Problem



# SSL Encryption

- Many Cloud Applications use SSL
  - ◆ Encrypted data cannot be deduplicated effectively
- Some WAN optimization products can securely optimize SSL-encrypted traffic



# Public Cloud DC Deployment

- No physical access to public cloud facilities, so how to deploy WAN optimization device?
- Several deployment possibilities
  - ◆ Cloud provider deploys WAN optimization (SP owns the WAN optimization devices)
  - ◆ Virtual WAN optimization software deployed over network (you own the WAN optimization devices)
  - ◆ Other approaches (TBD)
- Public Cloud deployment for WAN optimization is not a hurdle

# Metering, Cost Accounting, and Licensing SNIA

- Private Cloud: Cost accounting measures business unit performance
  - ◆ Existing cost accounting processes
    - Count bytes and sessions
  - ◆ Alternatively, no cost accounting at all
- Public Cloud: Cost accounting key to extracting profit from WAN optimization service
  - ◆ Per-user/Time-based costing and licenses
  - ◆ Similar licensing processes as Cloud Applications



**Check out SNIA Tutorial:**

**Sunshine user Base  
Accountability with your  
Cloud Storage**

# Key Points

- The WAN is core component of The Cloud
- WAN optimization is essential to Cloud adoption
- WAN optimization is proven technology
- WAN optimization can be easily deployed to support The Cloud

- Please send any questions or comments on this presentation to SNIA: [trackcloudtechnologies@snia.org](mailto:trackcloudtechnologies@snia.org)

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

**- SNIA Education Committee**

**Mark Day  
Rob Peglar  
Wendy Betts**