



Education

The Role of WAN Optimization in Cloud Infrastructures

Josh Tseng, Riverbed

- ◆ The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- ◆ 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.

- Cloud computing promises to deliver potential gains in efficiency and cost savings by consolidating IT assets into public and private data centers. Virtualization technologies allow migration of virtual machines and data into and out of Cloud infrastructures.
- However, cloud-hosted information assets must still be accessed by users over existing WAN infrastructures, where there are performance issues due to bandwidth and latency constraints.

Agenda Topics

- Importance of the WAN and the Cloud
- What is WAN Optimization and how does it work?
- Deploying WAN optimization to the Cloud
- How WAN optimization facilitates both Cloud Compute & Cloud Storage
- Key requirements for WAN optimization for the Cloud



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

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

Issues with migrating to the Cloud

- Largely same existing physical WAN infrastructure used to access Cloud
 - ◆ Same WAN links with bandwidth and latency
- Cloud data centers are potentially “far away”
 - ◆ Cloud infrastructure supports many enterprises
 - ◆ Large scale drives lower per-unit cost for data center services
- All employees will be “remote” from their data
 - ◆ Even single-location companies will be remote from their data
 - ◆ HQ employees previously local to servers, but not with Cloud model

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 (Sec)	2.08	23.08	34.54
"Reports" Tab (Sec)	1.19	9.80	22.62
Upload 3MB doc (Sec)	26.74	88.54	109.03
Download (Sec)	10	38	40

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”

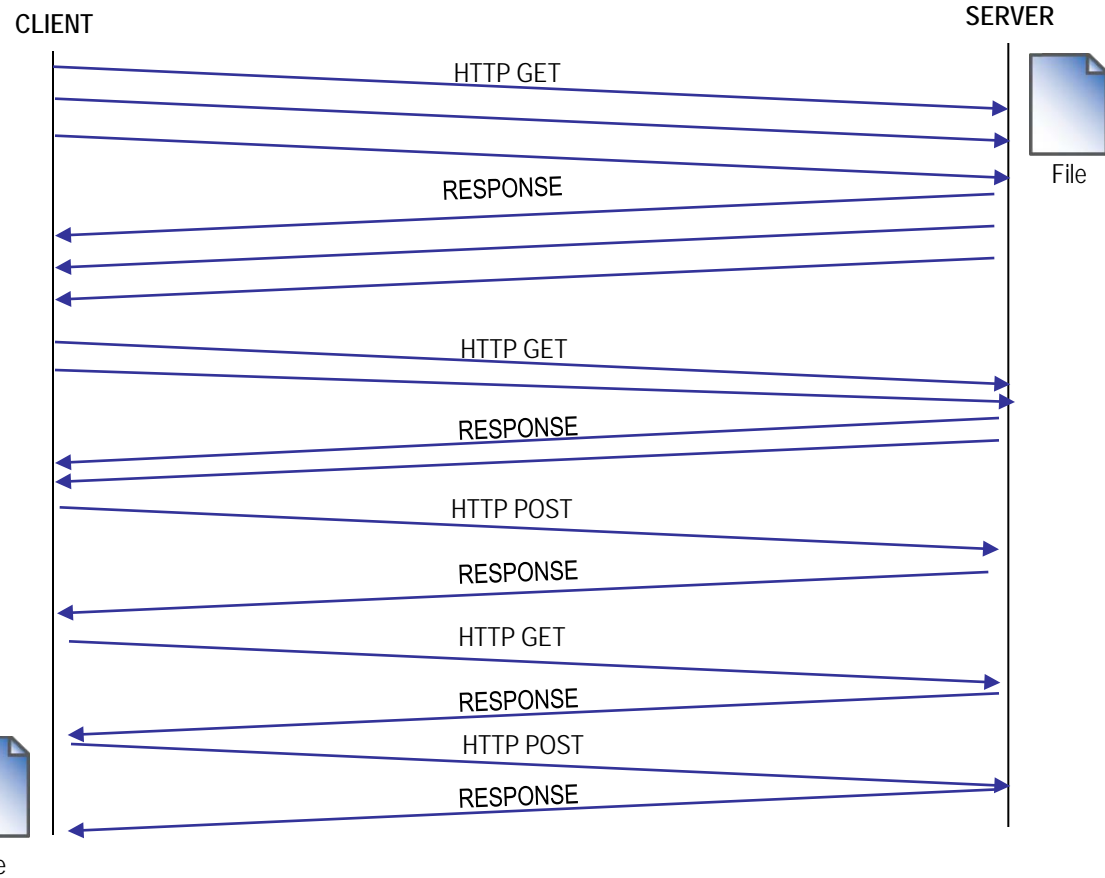
➤ Interactive apps, underlying protocols require 100s or 1000s of round trips for one operation!

- ◆ **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)



File


GET GIFs



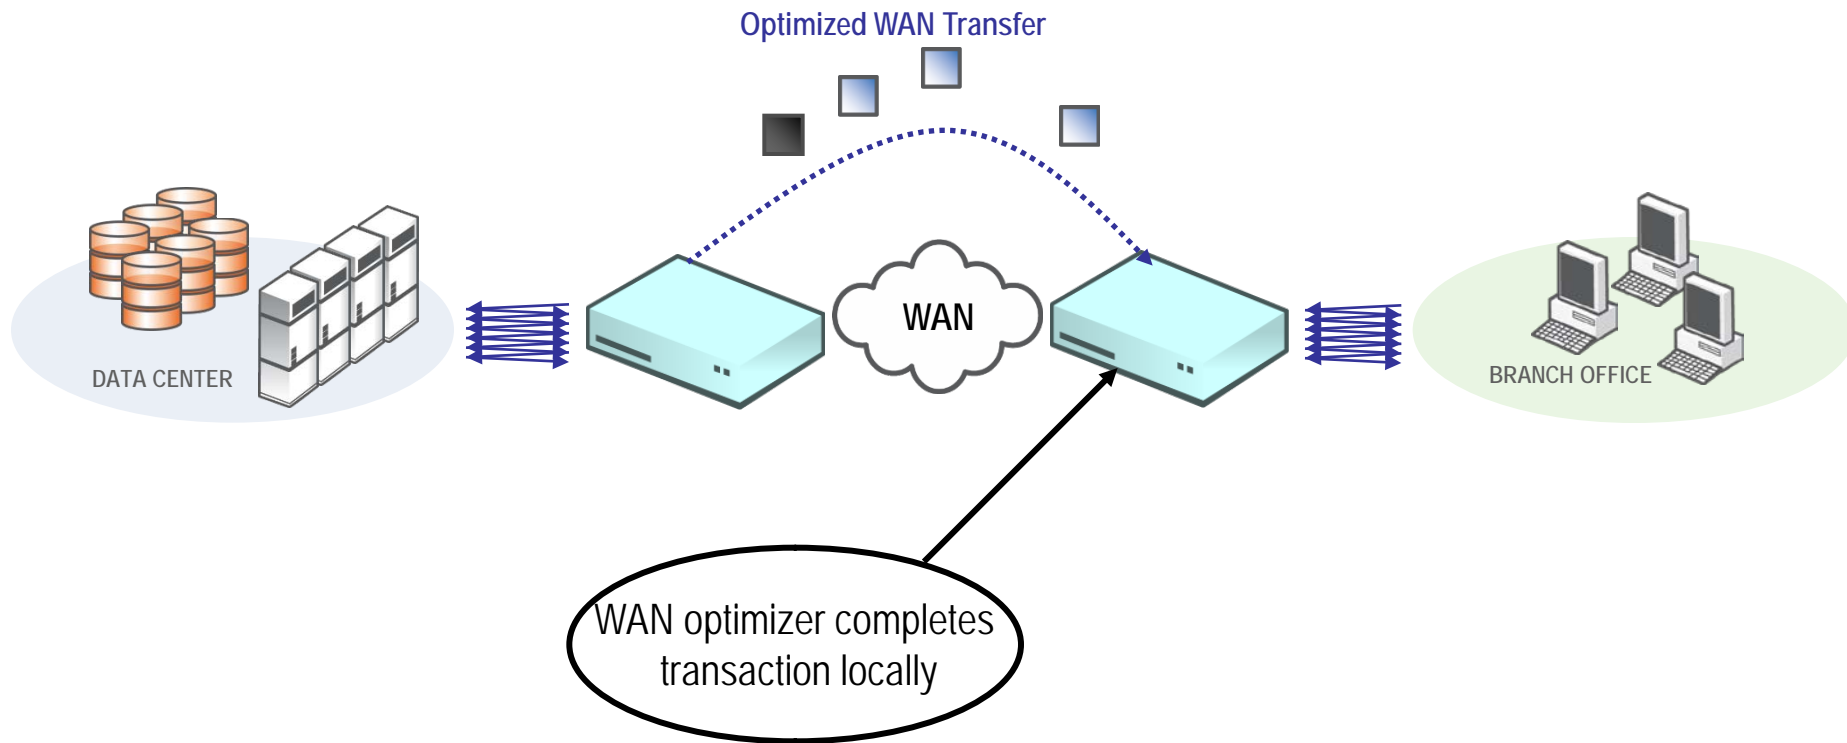
Fixing Bottleneck #2: Application-Level Chattiness

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

Addressing Application-Level Chattiness



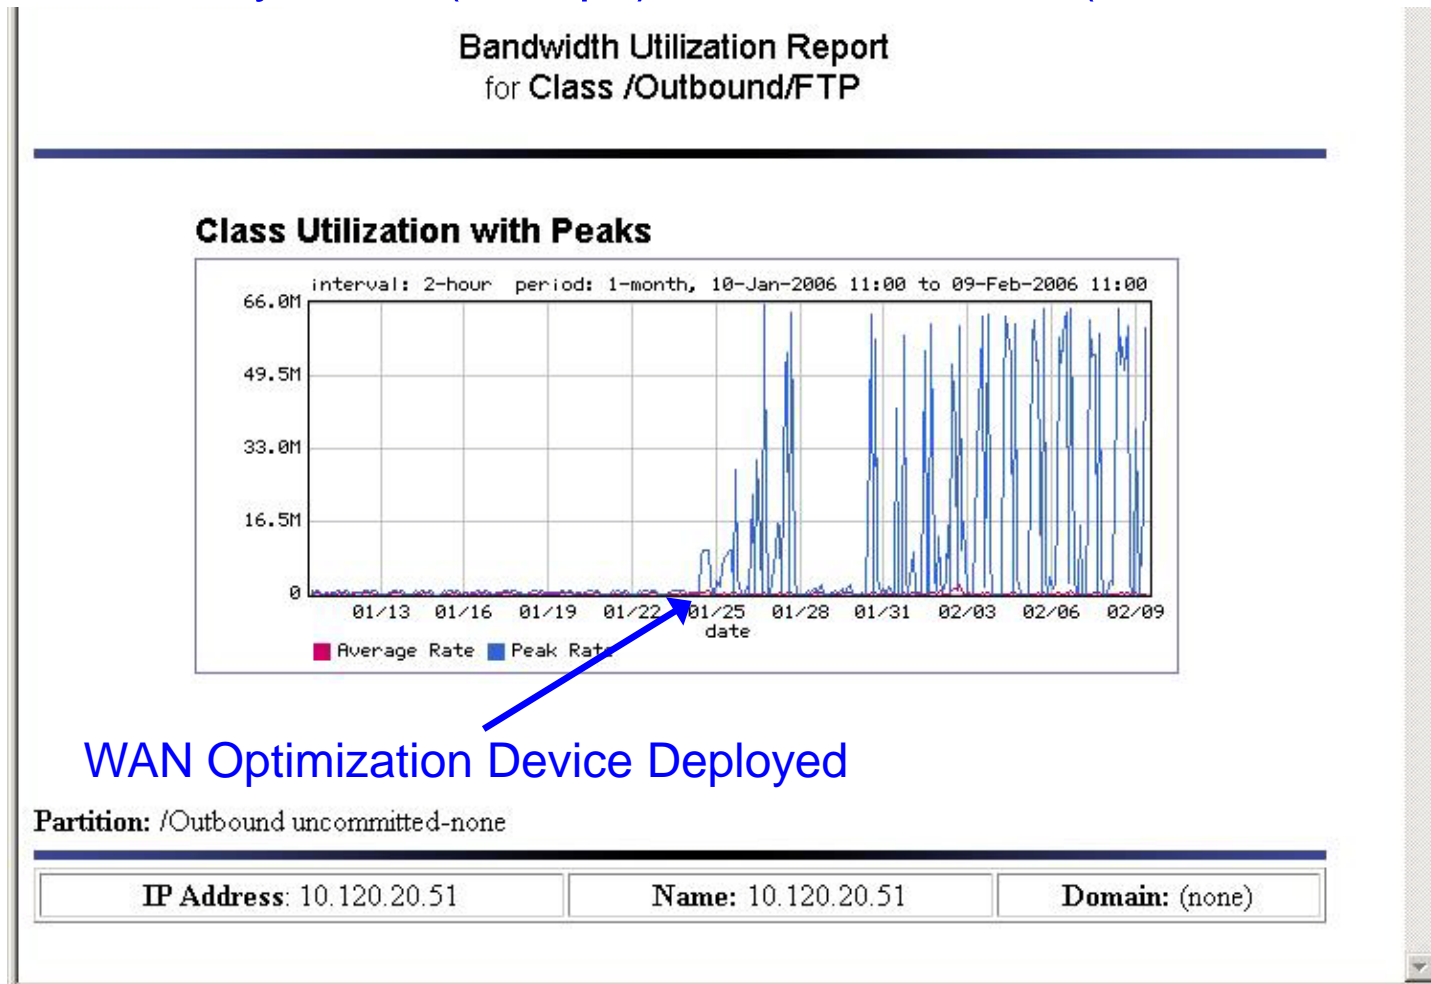
Addressing Application-Level Chattiness



Throughput achieved through WAN optimization

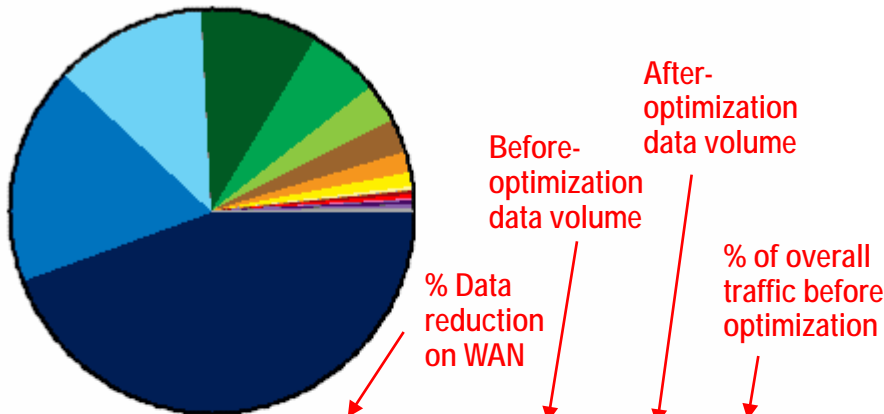
Actual Results:

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



Bandwidth reduction achieved through WAN optimization

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


WAN optimization solution components

➤ Physical appliances




- ◆ Dedicated hardware devices
- ◆ Maximize performance/cost

➤ Virtual appliances



- ◆ Shared hardware
- ◆ Improved flexibility for deployment

➤ Software clients



- ◆ WAN optimization from anywhere

➤ Traffic interception options



- ◆ Load balancing, scaling & clustering
- ◆ Different for Public vs. Private Cloud

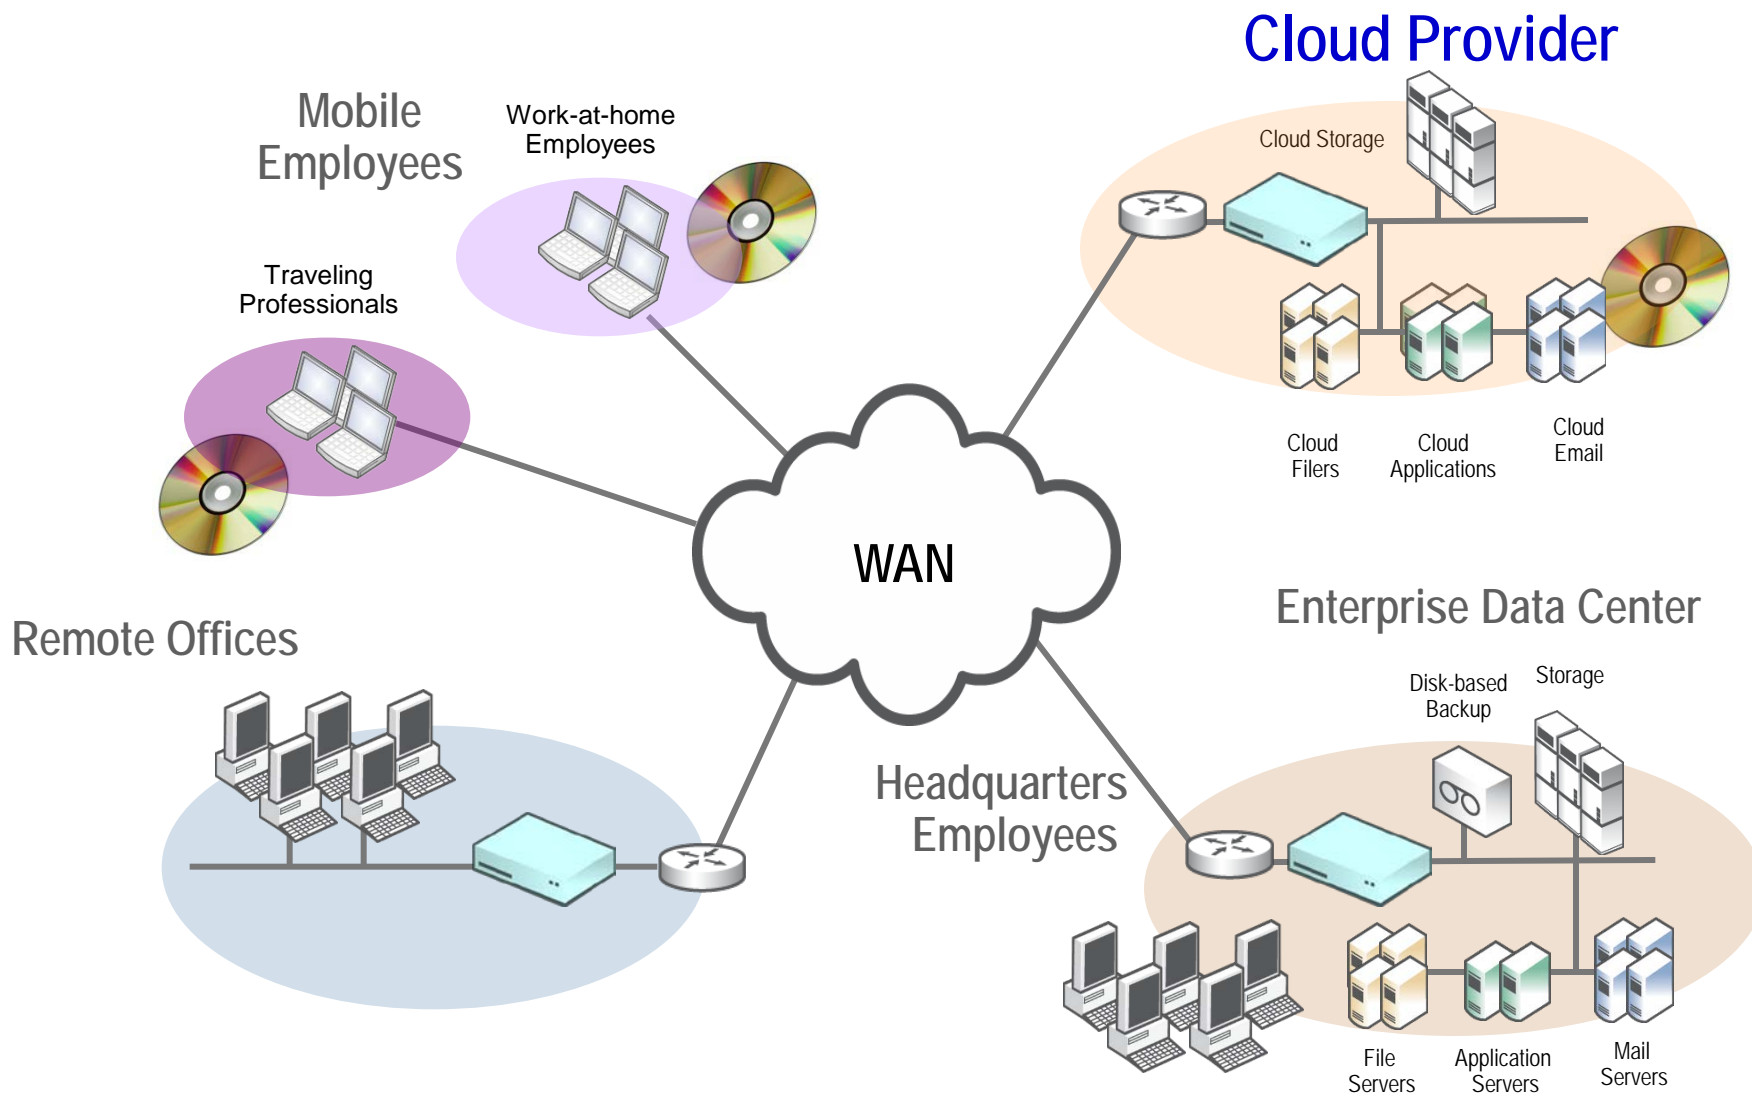
Deploying WAN Opt in the Cloud

- Deploy your own WAN optimization devices into Cloud provider's infrastructure
- Install own dedicated virtual WAN optimization software on your rented virtual machine (IaaS)
- Purchase WAN optimization from Cloud provider services catalog
 - ◆ Purchased just like software applications in SaaS model
 - ◆ Time-based and user capacity-based licensing
 - ◆ Licenses purchased as needed through provider portal interface using credit card

Deploying WAN Opt in the Enterprise **SNIA**

- Physical or virtual appliances at all locations; software clients for mobile users
 - ◆ Asymmetric (i.e., single-ended) solutions don't adequately address WAN-related performance issues
- Provisioned directly from WAN opt equipment vendor or through service provider
- Enterprise deployment must match products in Cloud
 - ◆ WAN optimization relies on proprietary technology
 - ◆ No interoperability between products from different vendors

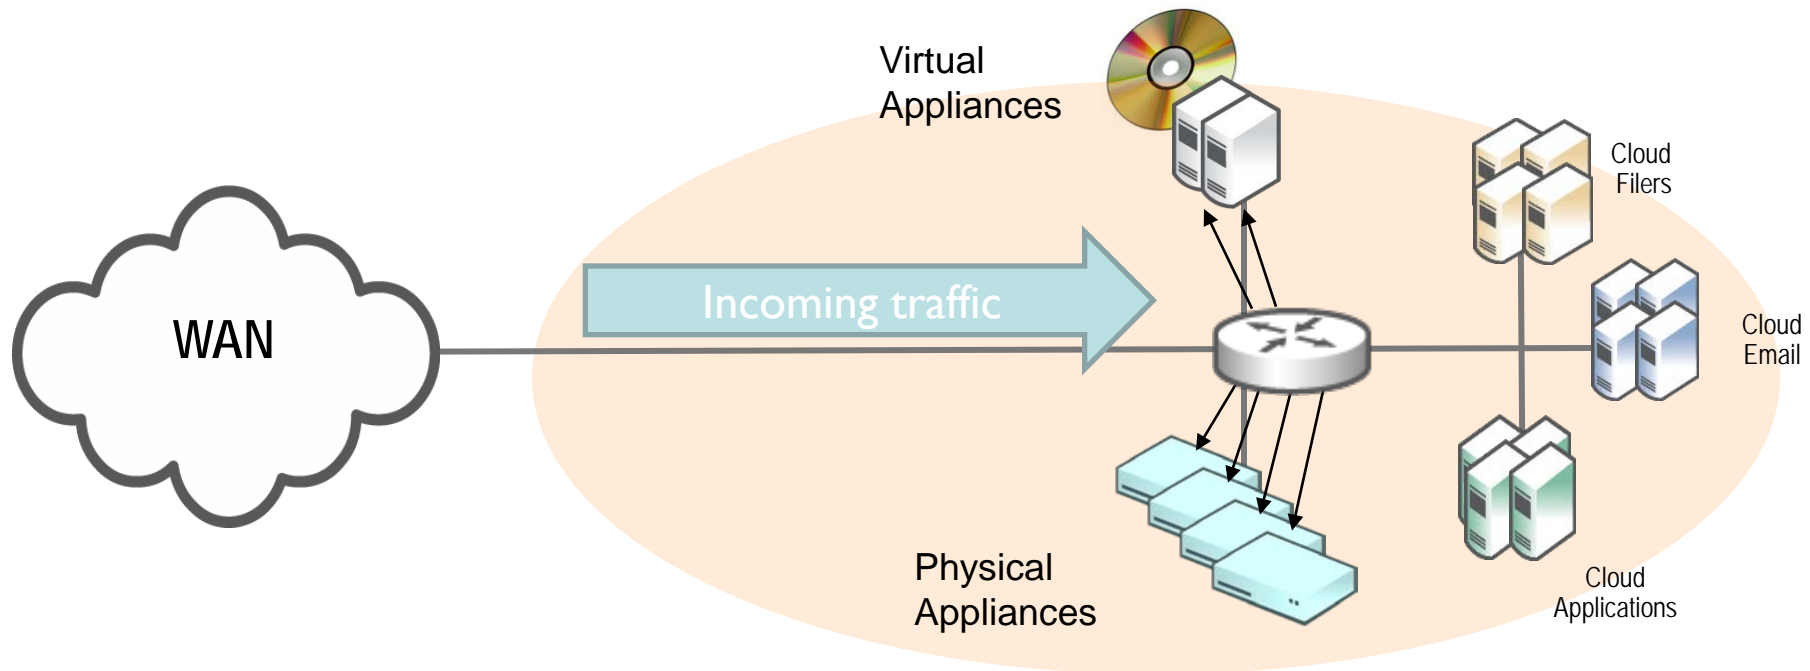
Deploying WAN Optimization



Traffic Interception in Private Cloud


➤ Private Cloud Traffic Interception Methods:

- ◆ In-path deployment of physical appliances
- ◆ Dedicated load balancer appliance
- ◆ WCCP or PBR in LAN switches



Traffic Interception in Public Cloud

- Not possible: WCCP, PBR, load balancer, In-path
 - ◆ No access to physical LAN in a Public Cloud
- Traffic Agents hosted on Application VM's
 - ◆ Redirects traffic to WAN optimization devices



➤ WAN optimization benefits for Cloud Computing

- ◆ IT consolidation/operational efficiency
- ◆ Maintain fast data access for applications



➤ WAN optimization benefits for Cloud Storage

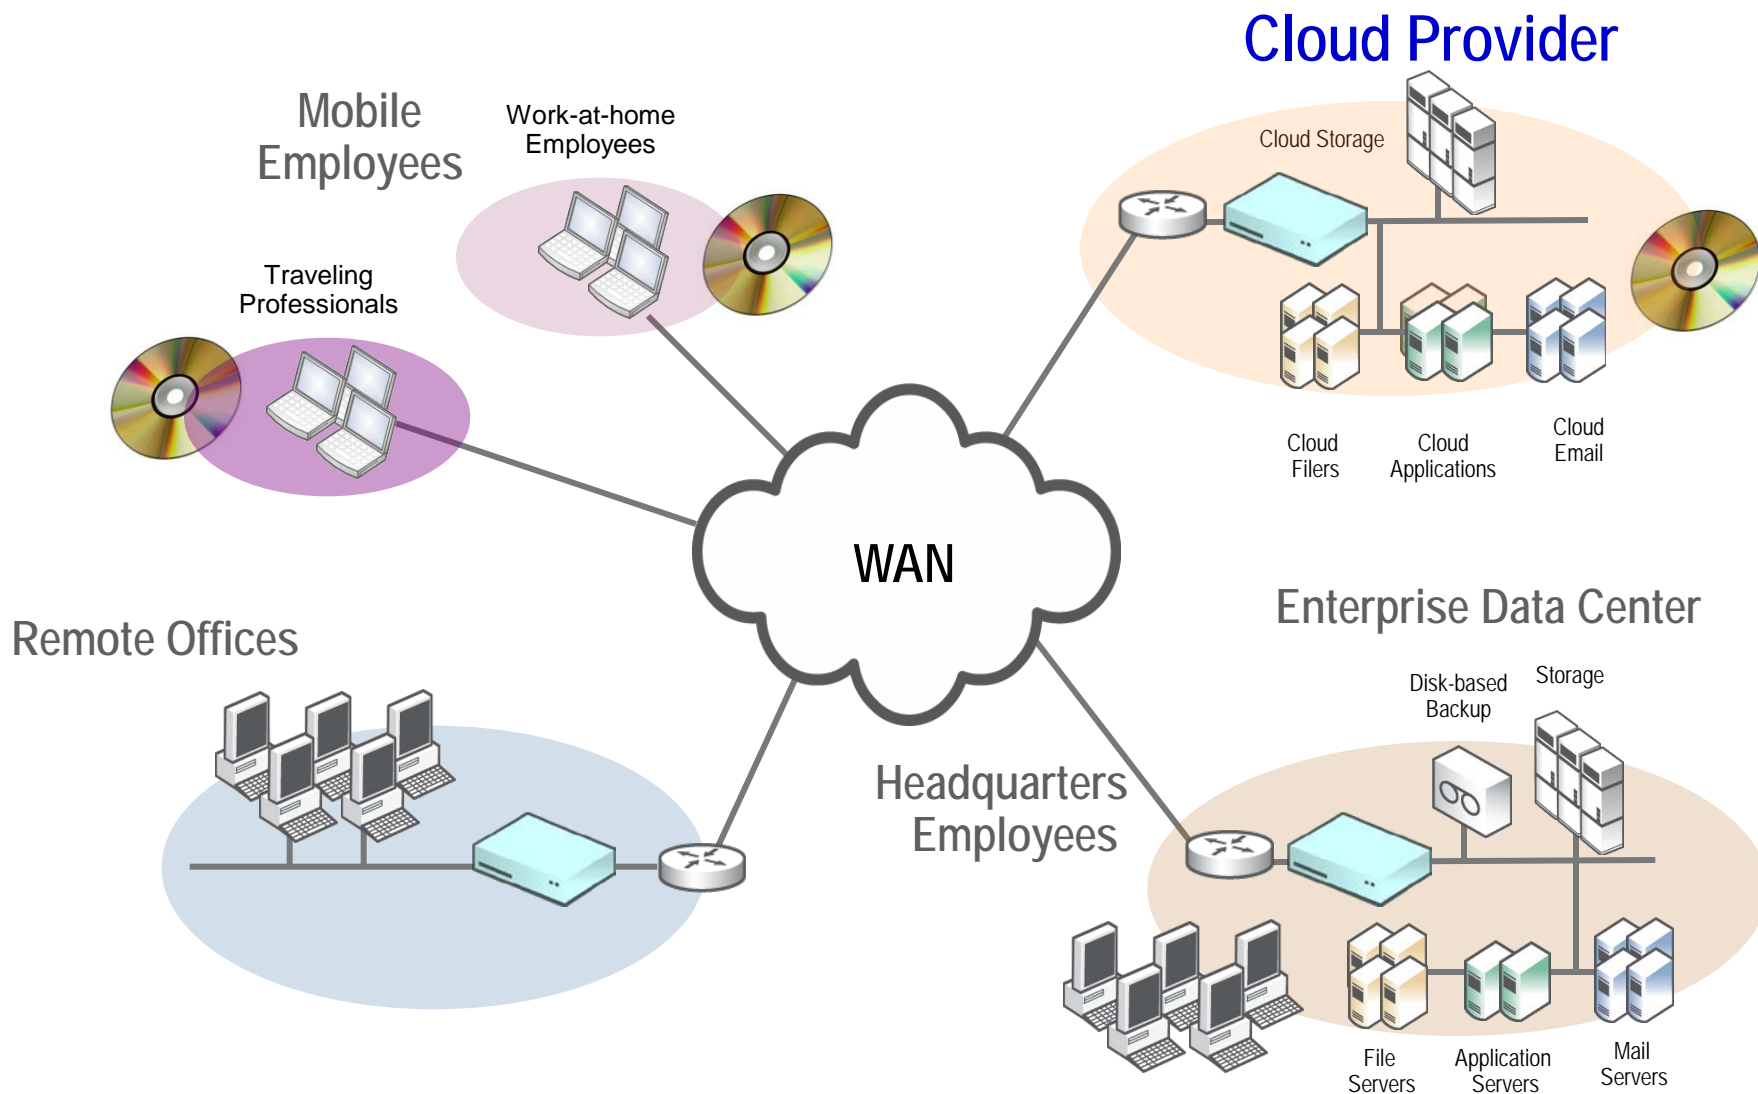
- ◆ Backup/archival of data
- ◆ Disaster recovery



Move compute resources into Cloud (i.e., SaaS, PaaS, & IaaS)

- Consolidate applications into the Cloud
- IaaS: Deploying or migrating applications to Cloud
 - ◆ VMDK files can be very large; WAN opt needed to facilitate transfer
- SaaS, PaaS, & IaaS: Facilitate access to Cloud applications
 - ◆ Maintain performance for accessing application data
 - ◆ Minimize bandwidth consumption from accessing Cloud applications
 - ◆ Many applications accessed using chatty protocols even when in the Cloud


Move compute resources into the Cloud



- **Primarily backup and/or archival of data to the Cloud**
 - ◆ Optional use of data deduplication to reduce Cloud storage consumption
 - ◆ WAN throughput must support size of the backup/archival data set
- **WAN optimization is cost-effective approach to meet throughput/performance requirements**
 - ◆ Achieve performance requirements at a fraction of cost
 - ◆ E.g., OC-3 for the cost of a DS-3

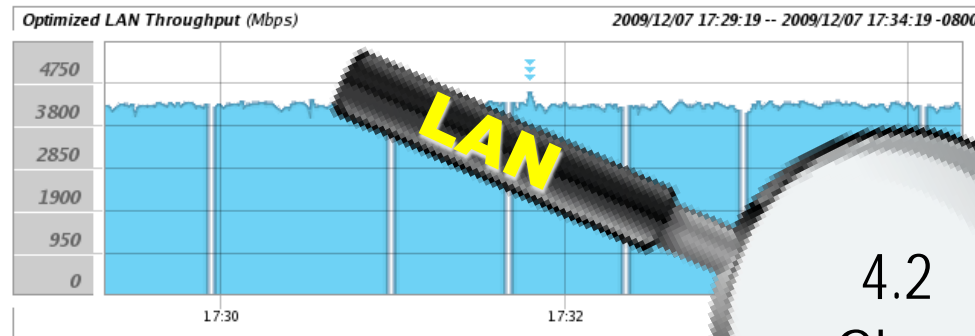
Backup/Archival of data to Cloud

- Potentially large volumes of data moved to Cloud
 - ◆ Move as much as 10's or 100's of TB of data
 - ◆ Data may or may not already be deduplicated in array or other device




Disaster Recovery: Throughput achieved through WAN optimization

Reports > Optimization > Optimized Throughput ?



- Peak LAN Throughput At 17:31:48 on 2009/12/07
- 95th Percentile LAN Throughput
- Average LAN Throughput



Period: Last 5 Minutes
Application: All
Traffic: Bi-Directional

- Peak WAN Throughput At 17:31:48 on 2009/12/07
- 95th Percentile WAN Throughput
- Average WAN Throughput

Key WAN Optimization Requirements for Cloud Infrastructures

➤ Scaling

- ◆ What works for a private infrastructure may not work for the Cloud infrastructure
- ◆ Public Cloud providers must scale to support multiple enterprises

➤ Interoperability

- ◆ WAN optimization solutions from different vendors are not interoperable
- ◆ Standardize on a vendor solution used by Cloud providers or deploy your own through IaaS

➤ Performance and throughput

- ◆ Cloud infrastructures must support data center transfer throughputs

- Utility computing allows consolidation and sharing of IT resources
 - ◆ Similar to consolidation and sharing of power/other resources
- High-performance WAN infrastructure key to facilitating utility computing
 - ◆ WAN optimization increases capacity and performance of the WAN infrastructure
 - ◆ But WAN optimization must scale to meet Cloud provider infrastructure requirements

- Please send any questions or comments on this presentation to SNIA: tracktutorials@snia.org

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

- SNIA Education Committee

Joseph White

Mark Day

Rob Peglar

Wendy Betts

Steve Riley