



Education

Fundamental Approaches to WAN Optimization

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

- Defining the WAN performance problem for distributed enterprises
- Issues impacting application performance over the WAN
- The Pros and Cons of traditional approaches
- New Wide-Area Data Services (WDS) approaches to WAN Acceleration

Distributed Enterprise Challenges

- **Branch Office Users**
 - ◆ Lengthy delays accessing data from data center
- **Traveling Users**
 - ◆ Lengthy delays accessing data from home or hotel
- **Server/Storage Consolidation**
 - ◆ Distributed servers are difficult to manage
 - ◆ Data stored in remote offices is not secure
- **Disaster Recovery**
 - ◆ Backup windows are too long

Impacts of Poor WAN Performance



Poor Wide-Area Application Performance: Three Root Causes

Bandwidth limitations


Transport protocol chattiness

Application protocol inefficiencies

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



Bottleneck #2: TCP “Chattiness”

Send a 40 MB file across the WAN



With *unlimited bandwidth & cross country latency* data transfer would *still* take 60 seconds due to TCP based round trips

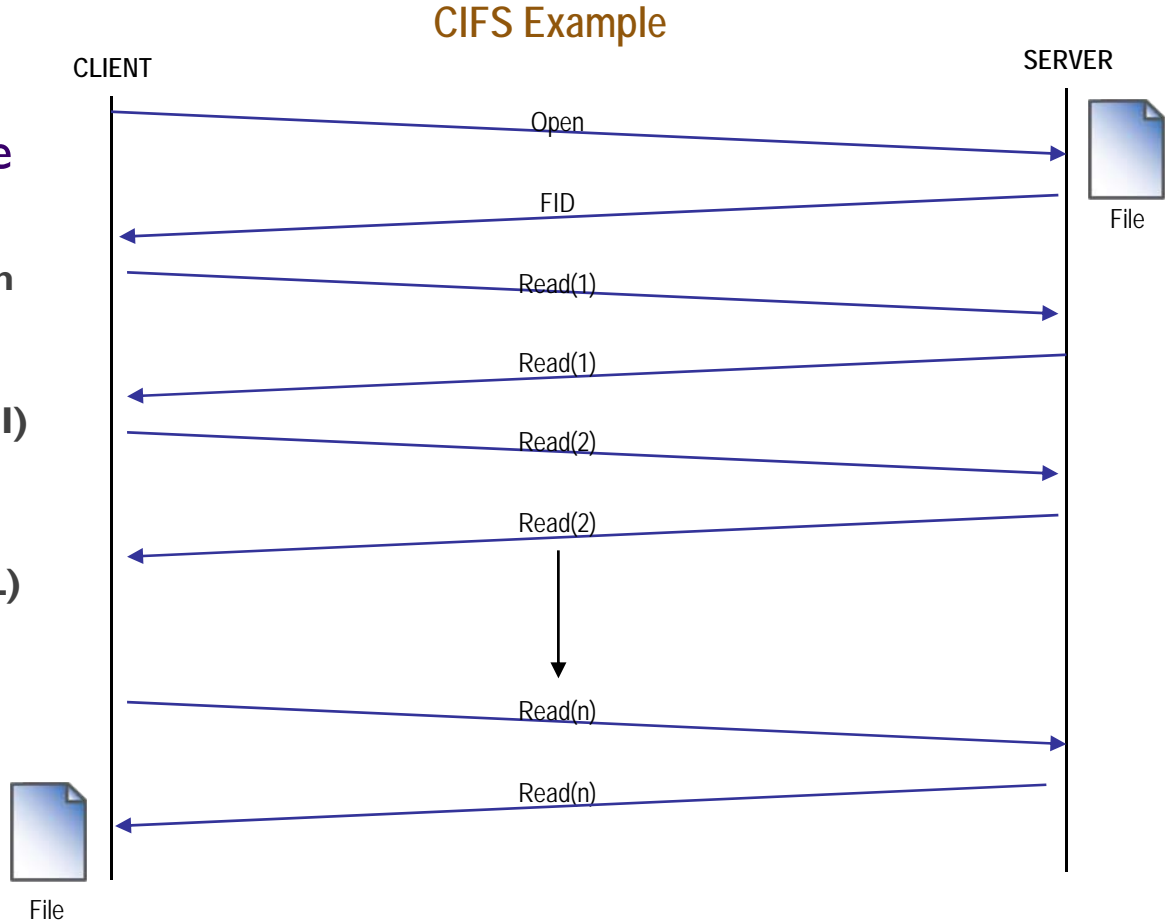
Divide traffic and send 64 KB at a time across the WAN



Bottleneck #3: Application “Chattiness”

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

- ◆ **Common Internet File System (CIFS)**
- ◆ **Messaging Application Programming Interface (MAPI)**
- ◆ **UNIX File Sharing (NFS)**
- ◆ **CRM (SQL)**
- ◆ **Document Management (SQL)**
- ◆ **Call Center Apps (SQL)**
- ◆ **Project Mgmt Apps (SQL)**
- ◆ **Accounting Apps (SQL)**
- ◆ **CAD/CAM Mgmt Apps (SQL)**
- ◆ **Custom Apps (SQL)**



The Holy Grail: LAN-like Performance Over the WAN

Reduce hard costs

Move file servers, mail servers, web servers, and tape backup systems to a central location

Increase productivity

Employee collaboration regardless of location. Order entry tasks, file transfers, and other data exchanges completed instantly

Improve data protection


Remote data backup in minutes vs. hours

Legacy WAN Acceleration Approaches

- Add WAN Bandwidth
- Compression
- QoS
- Caching/Data Prepositioning


Problem and Solution Analogy

- You must pick up 100 suitcases for your guests at the airport and take them to your hotel resort
- Your car only carries 4 suitcases at a time
- The road between airport and hotel has only one lane in each direction



Legacy Solution #1: Add WAN Bandwidth/Build More Lanes

- More freeway lanes will help...sometimes
- Does a 20-lane highway let you move these bags 20 times faster?
 - ◆ No, because you still can only carry 4 bags on each trip
 - ◆ You still have to make 25 trips



Adding Bandwidth: Pros & Cons

- Adding WAN bandwidth helps w/congestion
 - ◆ Scarce bandwidth constrains throughput

- However, WAN bandwidth doesn't address TCP and application-level chattiness
 - ◆ Applications still take the same number of round-trips
 - ◆ Speed-of-light dictates a minimum time required for each round-trip

Solution 2: Compression/Use Smaller Cars

- Require everyone use miniature cars
- Squeeze cars so each one is $\frac{1}{4}$ the size
- Highway can hold 4x more cars!
- But... No improvement in trip time at all
 - ◆ Still need 25 trips to move 100 suitcases



Compression: Pros & Cons

- Similar to adding WAN bandwidth
 - ◆ Helps to address congestion issues
- Doesn't address TCP, app-level chattiness
- Limited performance improvement if application exhibits chatty behavior

Solution 3: Quality-of-Service/Car Pool Lanes

- Does having a carpool lane between airport and hotel help deliver 100 bags of luggage?
 - ◆ Only if you have special access
 - ◆ Those without access must wait
 - ◆ You still can only carry 4 bags on each trip
 - ◆ So you still have to make 25 trips



QoS-only: Pros & Cons

- High-priority applications get priority BW access
- QoS is a zero-sum mechanism
 - ◆ Only allows you to pick winners and losers
 - ◆ Some apps get better performance: others suffer
 - ◆ Doesn't deliver additional bandwidth
- TCP and application-level chattiness still a problem

Solution 4: Caching/Cloning & Pre-Positioning

- Anticipate guests' luggage requirements
- Pre-purchase/pre-position suitcases with anticipated contents (e.g., garments, toiletries, etc...) using information from guests' prior visits
- For guests that bring identical suitcases from previous visit, you don't have to fetch them from the airport

Caching/Pre-Positioning Pros & Cons

- Potential to reduce round-trips!
- Not all guests always bring the same suitcase and contents on every visit!
 - ◆ Potential for data coherency issues (“I got the wrong suitcase!”)
- Stores application-specific objects
 - ◆ File/object processing overhead
 - ◆ File/object renamed
 - ◆ No deduplication of data
 - ◆ What about other applications?

- Get a bigger car – more data with each trip
- Don't send whole suitcases
 - ◆ Deconstruct the suitcases: open them up and send the contents
 - ◆ Don't care about the type of suitcase (application type doesn't matter)
- Don't look at just 4 suitcases at a time
 - ◆ Examine the contents of all 100 suitcases and transfer them all at once
 - ◆ Application-level read-aheads



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



**Check out SNIA Tutorial:
Understanding Data
Deduplication**

Disk-based Data Reduction

60 to 90 percent data reduction




Fixing Bottleneck #2: TCP Chattiness

- Use larger TCP windows
 - ◆ WAN acceleration solution should use larger TCP buffers
 - ◆ Send more data in each round-trip
- Send “virtual” data per TCP window/round-trip
 - ◆ Send symbols in each TCP window
 - ◆ Each symbol represents “virtual” amounts of data
 - ◆ Fewer round-trips necessary


Bottleneck #2: TCP “Chattiness”

Send a 40 MB file across the WAN



Each TCP window contain symbols that virtually represent even larger amounts of data

Larger 512KB TCP windows send even greater amounts of “virtual” data




Potentially several GB of virtual data transferred using symbols in each TCP window

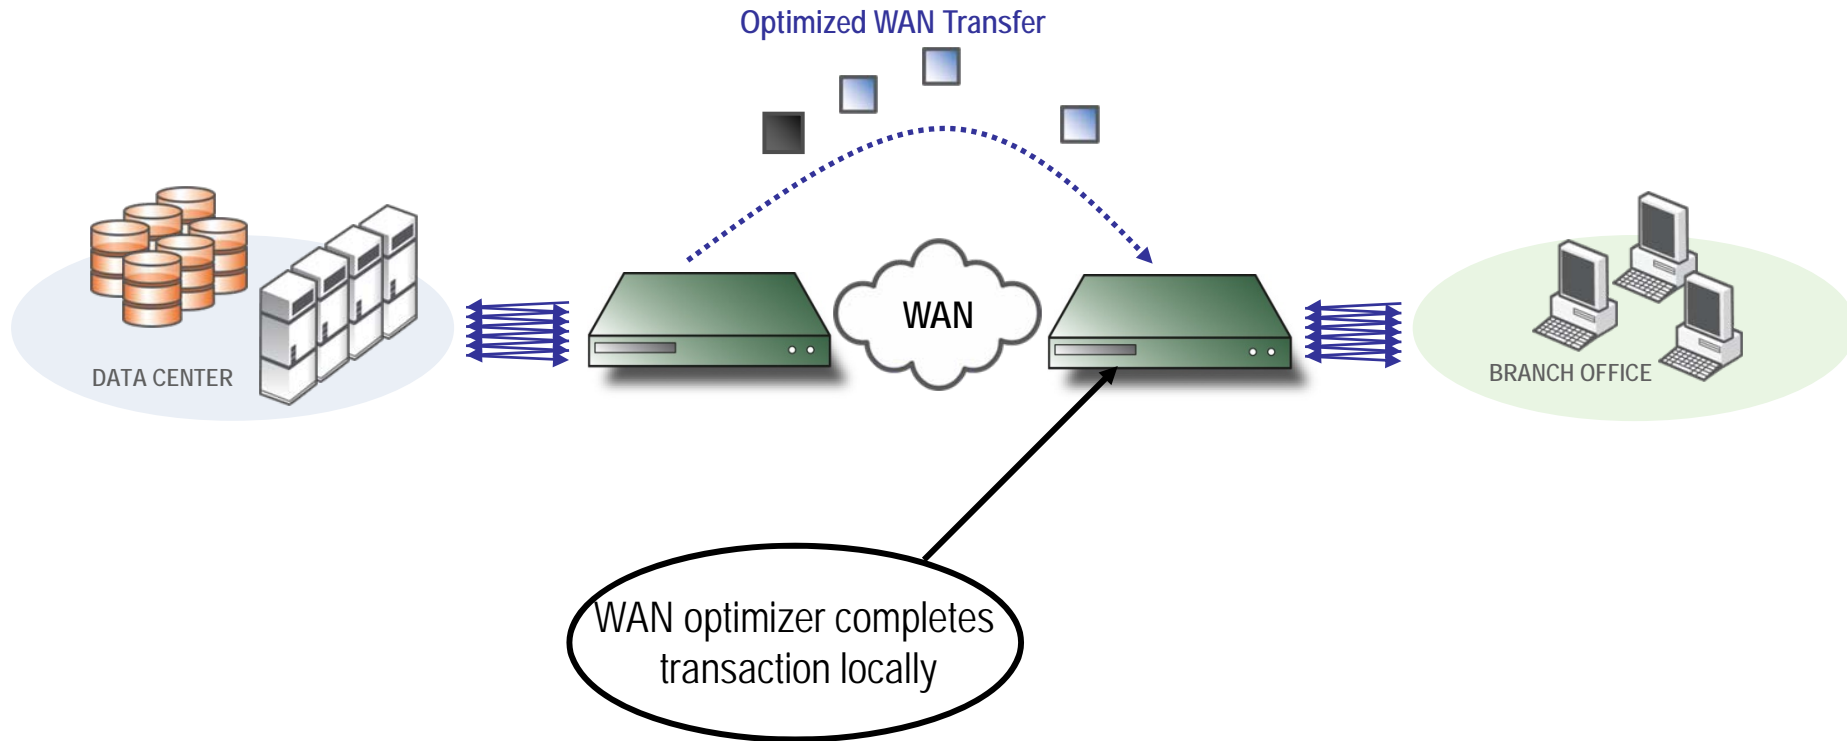
Fixing Bottleneck #3: Application-Level Chattiness

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

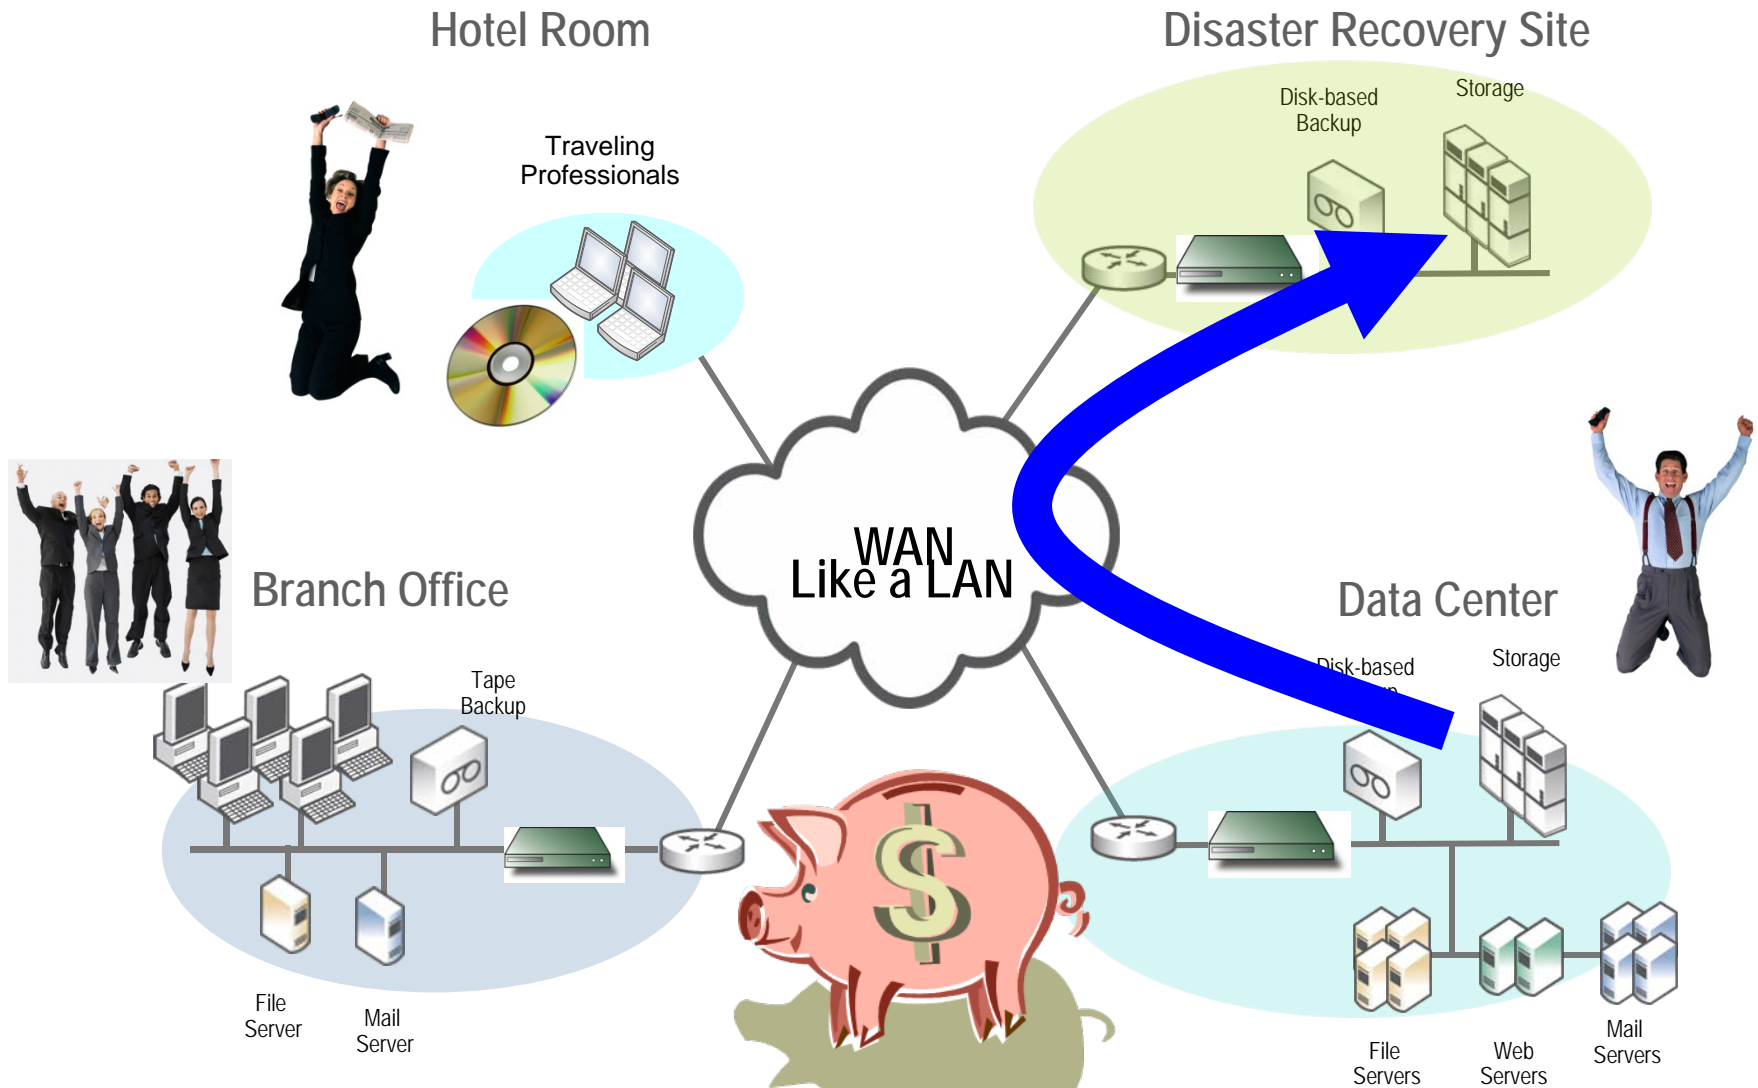
Addressing Application-Level Chattiness



Addressing Application-Level Chattiness



Solving the WAN Performance Problem



WDS Solution Requirements

- Not just adding bandwidth
 - ◆ Expensive, doesn't address latency issues
- Not just packet compression
 - ◆ Packet-level compression doesn't address latency issues
- Not just QoS
 - ◆ QoS doesn't address latency issues or bandwidth constraints
- No caching
 - ◆ Caching stores data in original application/object format with no deduplication
 - ◆ Data coherency issues
 - ◆ Scaling Limitations

How Wide-Area Data Services (WDS) Addresses Distributed Enterprise Challenges

- **Branch Office Users**
 - ◆ Can access data at LAN-like speeds
- **Traveling Users**
 - ◆ Fast data access from any location
- **Server/Storage Consolidation**
 - ◆ Consolidation saves costs and makes backup easier
 - ◆ Centralized data is more secure
- **Disaster Recovery**
 - ◆ Backup windows reduced significantly to manageable timeframes

Conclusion

- WAN performance key to IT efficiency gains
- Legacy approaches don't address all three core WAN performance issues
- Wide-Area Data Services (WDS) solutions are providing measurable benefits today
 - ◆ Productivity gains
 - ◆ Reduced infrastructure costs
 - ◆ Data protection and security
 - ◆ Strongly positive ROI

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

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

- SNIA Education Committee

**Apurva Dave
Mark Day
Kim Kaputska
Rob Peglar**