

BoF Discussion Topic

CIFS Acceleration

Led by Mark Rabinovich
Architect, Visuality Systems, Israel

Discussion Agenda

- Introduction
 - Some CIFS drawbacks
 - When CIFS acceleration comes on
- How can we improve CIFS traffic?
 - What to improve?
 - How to improve?
- CAX – CIFS acceleration example
 - Positioning
 - Platforms

Introduction

Who are we?

- ❑ Our company – Visuality Systems Israel Ltd (VS).
 - ❑ 10 years experience developing and selling embedded/mobile CIFS.
 - ❑ More than 80 world wide mainly Blue Chip customers in just about every market space possible from Aerospace to Home Devices.
 - ❑ New product introduction in November 2008
 - ❑ CAX – CIFS traffic accelerator
- ❑ Your speaker – Mark Rabinovich
 - ❑ 28 year in software development
 - ❑ From 2003: Architect and Project Manager in VS

Some CIFS Issues

- ❑ Satellite networks introduce up to one second of latency
- ❑ International network might introduce 500 milliseconds of latency

Synchronous protocols are ineffective when run over networks with latency.

Bandwidth Limitation

- ❑ Satellite networks usually have limited bandwidth
- ❑ Regular DSL/Cable end users accessing remote CIFS servers may experience significantly low bandwidth

Redundant traffic over a limited bandwidth link takes too long.

- ❑ SMB protocol has synchronous nature. CIFS is mostly synchronous. Multiplexing is limited.
- ❑ Currently Microsoft CIFS client is redundant - many commands are repeatedly sent

- ❑ Opening 1.5 MB Word document involves about 630 SMB packets and requires more than 5 minutes open time (with 0.8 sec round trip).
- ❑ Saving the same document involves about 2400 SMB packets and takes more than 13 minutes.
- ❑ Opening a remote folder in Windows File Explorer involves hundreds of sequential request/response round trips

The above results make CIFS virtually unusable

Examples (cont)

- ❑ Windows File Explorer uses a number of concurrent threads each one issuing the same or similar generic requests (SRVSVC, QueryFSInfo, etc.) when browsing folders or opening files.
- ❑ Same information requests (QueryPathInfo) are often repeated at least three times
- ❑ One request uses low info level while subsequent requests asks for more details of the same file (FindFirst, QueryPathInfo)

The Problem

- ❑ Global corporate networks must provide reasonable file sharing connectivity over satellite/international links
- ❑ Regular home users should be able to benefit from CIFS over DSL/Cable connections
- ❑ This is virtually impossible with the current implementation of SMB
- ❑ SMB2 addresses some of the issues, but not all of them

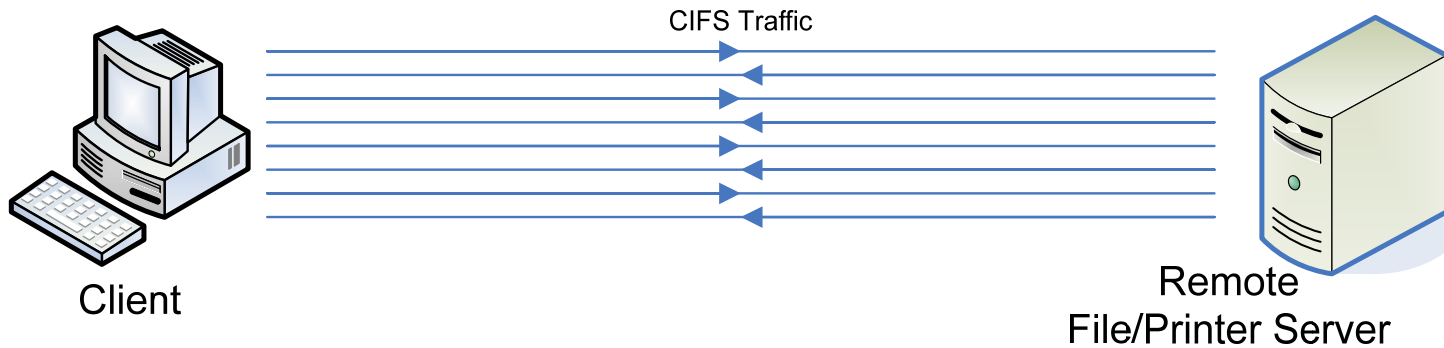
Discussion

Your opinion ...?


Accelerating CIFS Traffic

Saving on Round Trips

Non accelerated CIFS model



Accelerated CIFS model

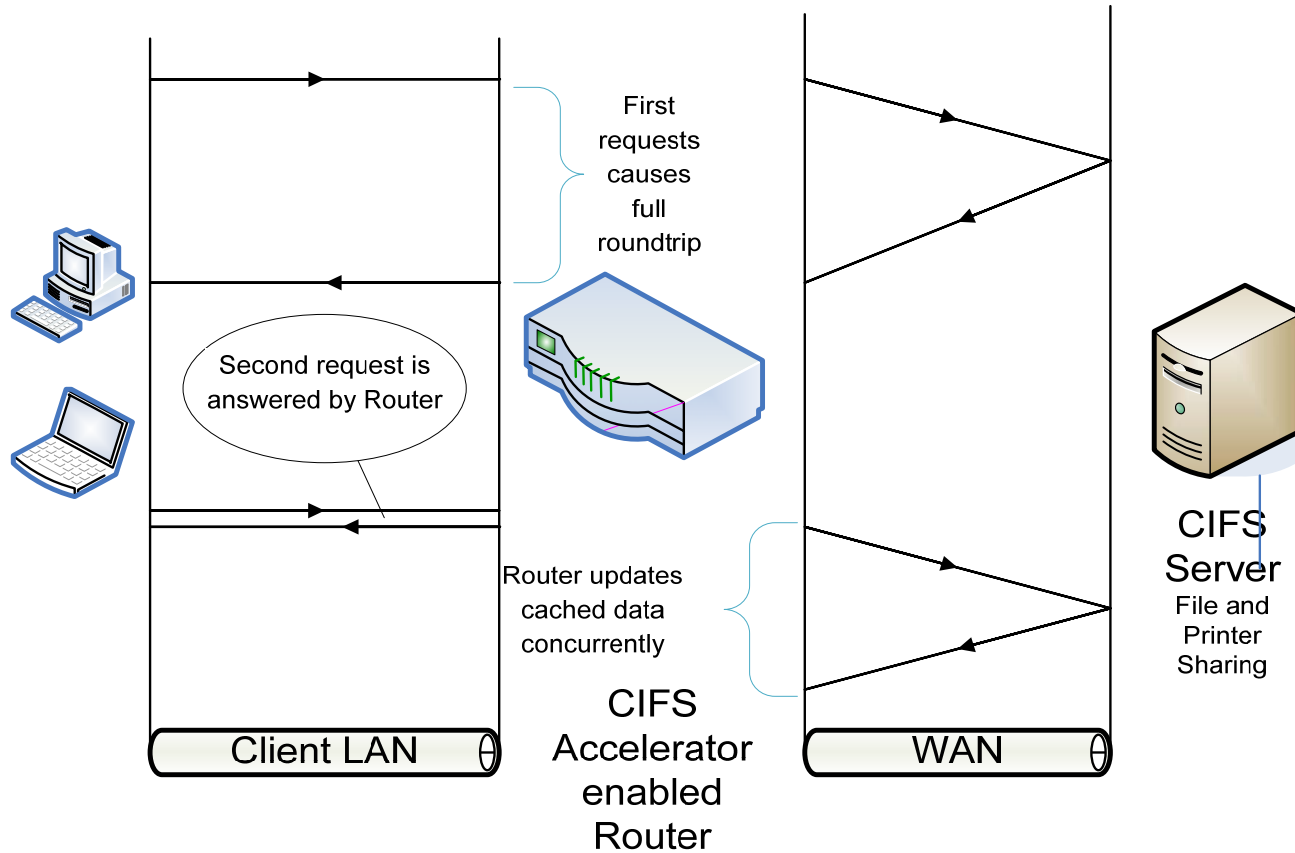


How to Save on Round Trips?

- ❑ Send less requests to server
- ❑ Send requests to server concurrently
- ❑ Utilize the entire MPX count

1. Caching
2. Predicting
3. Accumulating

Caching

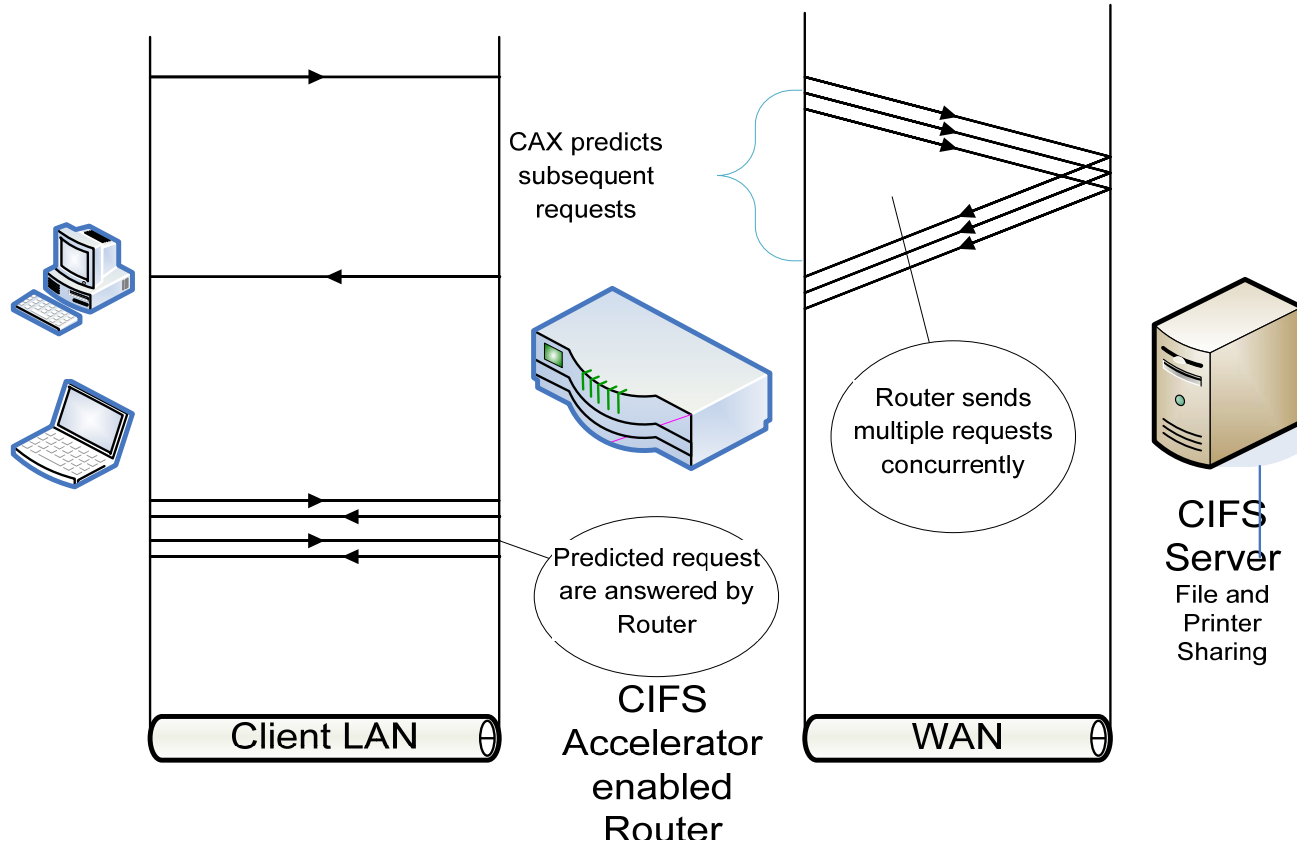


Time without acceleration – two roundtrips, with acceleration – one roundtrip

Caching What?

- ❑ File Information
 - ❑ FindFirst/Next/QueryPathInfo
 - ❑ NtCreateAndX
- ❑ Share Information
 - ❑ SrvSvc
- ❑ Server Information
 - ❑ SrvSvc
 - ❑ WksSvc
- ❑ File data
 - ❑ Write/WriteAndX
 - ❑ Read/ReadAndX
- ❑ File nonexistence
- ❑ Stream-full/stream-less file

Predicting

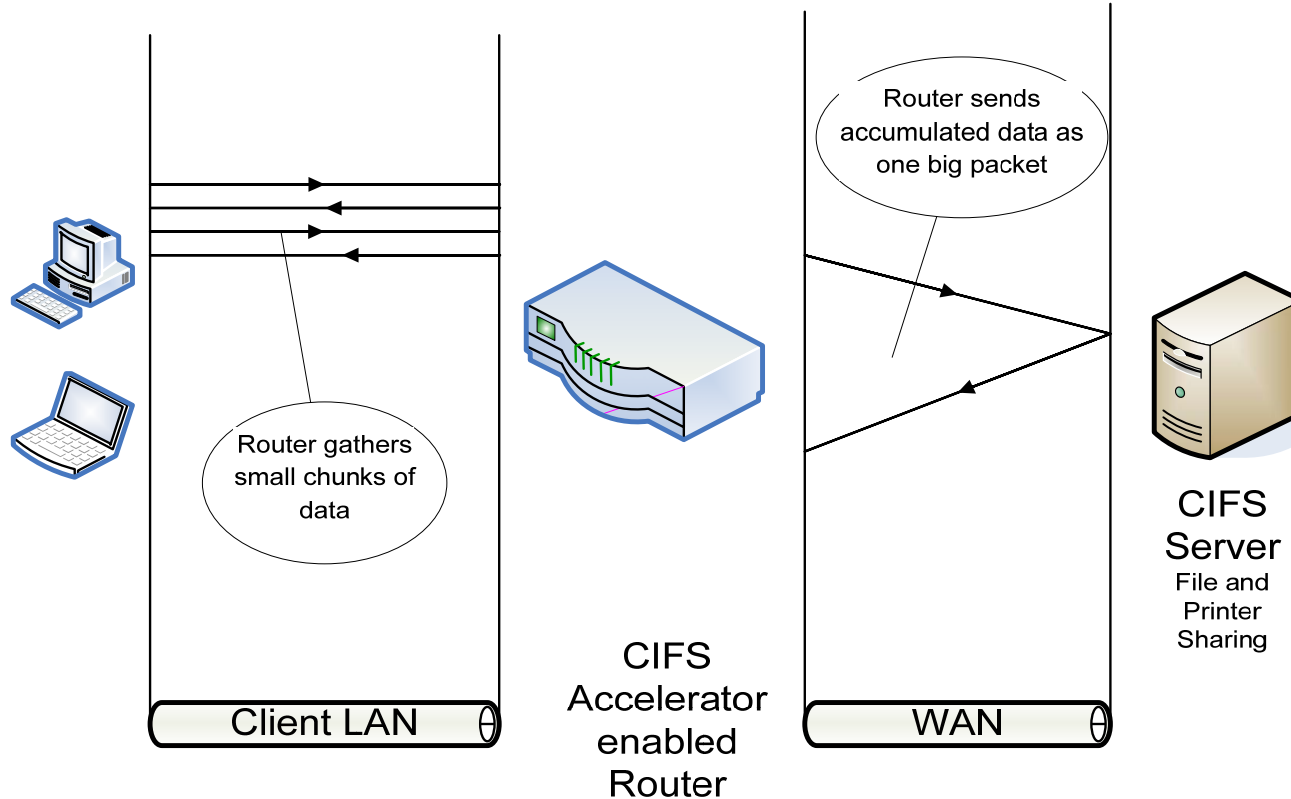


Time without acceleration – three roundtrips, with acceleration – one roundtrip

Predicting What?

- ❑ Almost always queried after TreeConnectAndX
 - ❑ FS Information
 - ❑ Root folder
- ❑ Force requests for complex scenarios
 - ❑ NTCreatAndX + IOCTL(function32) + Close
- ❑ Almost always succeed
 - ❑ Close
 - ❑ TreeDisconnect
 - ❑ Logoff
 - ❑ Subsequent Write/SetFileInfo

Accumulating data



Time without acceleration – two roundtrips, with acceleration – less than one roundtrip

Accumulating When?

- ❑ Sequential file read
- ❑ Always issue queries with the most comprehensive info level

Impact Analysis

- ❑ CIFS Acceleration introduces some inaccuracy
 - ❑ Cache goes out of sync
 - ❑ Late (error) responses
- ❑ This inaccuracy is comparable with native Windows inaccuracy
- ❑ Windows Client always knows to recover

How to reduce inaccuracy?

- ❑ Analyzing NOTIFIES
- ❑ Sending preventive requests for cached info

CAX – CIFS Accelerator Example

- ❑ For any network with latency and/or limited bandwidth
from satellite and long-range networks...
... to DSL/Cable links
- ❑ Embedded into communication environment – communication equipment
/ gateway servers / client workstations
- ❑ Transparent to user – preserves file sharing namespace semantics
- ❑ Client-side – accelerator works on a CIFS clients' side only

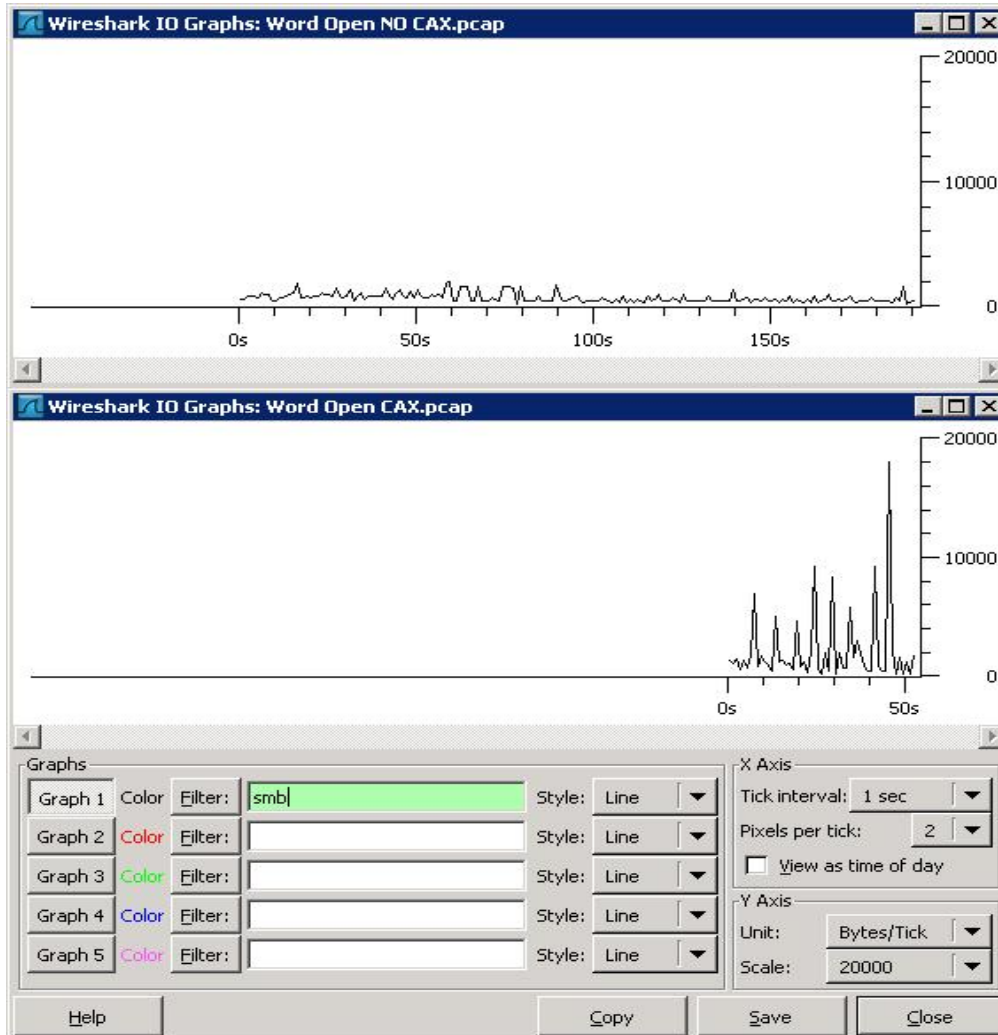
How does it work?

1. **Caches information/data**
 - ❑ Answers repeating (redundant) requests
 - ❑ File information
 - ❑ File data
 - ❑ File System/Server/Share information
2. **Predicts the answer**
 - ❑ We know which command will be next
 - ❑ We can send it concurrently
3. **Accumulates data in bigger messages**
 - ❑ We always use 60K buffers for reading

Accelerator improves the most common scenarios

- ❑ Folder browsing
- ❑ Big file upload/download
- ❑ MS Office document open/edit/save

MS Word File Open Acceleration



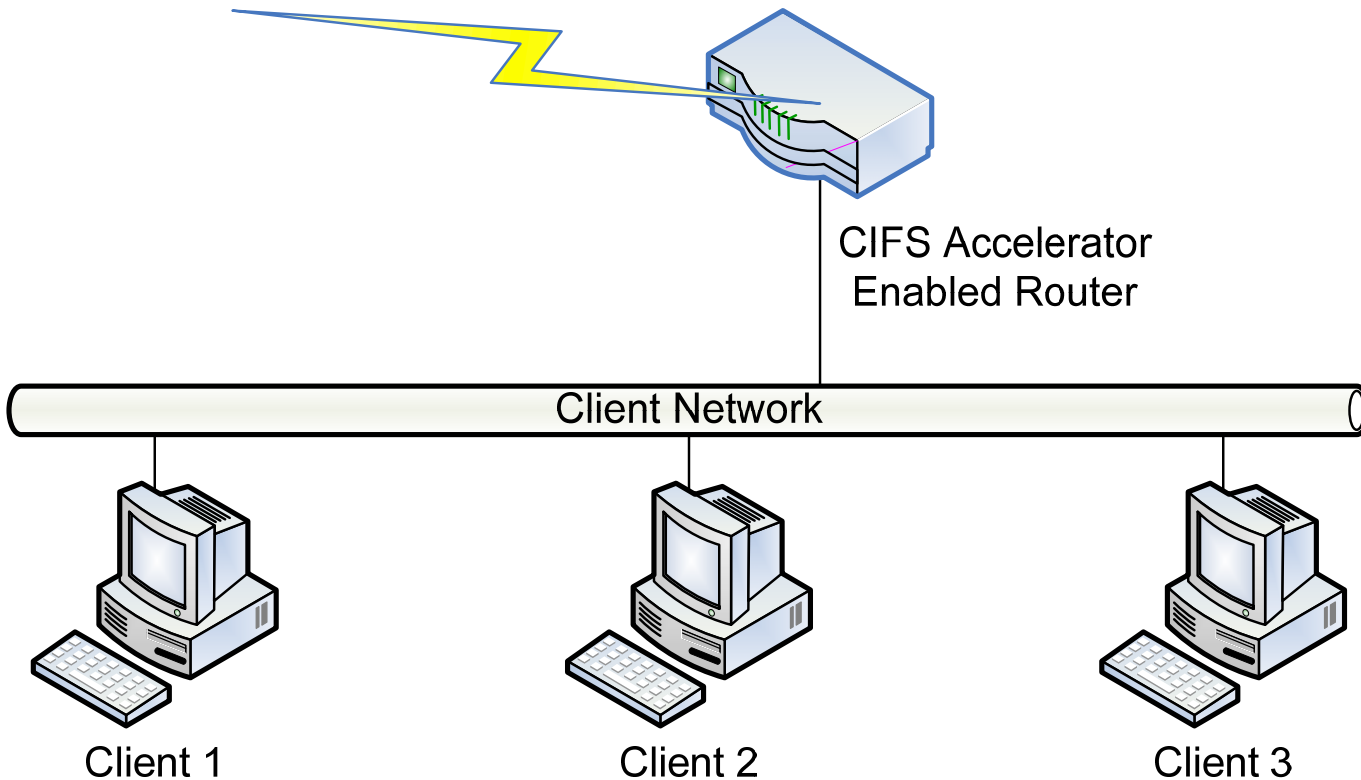
Client-side capture taken on satellite link simulator (RTT: 800 msec):

Upper graph: ~200 sec without acceleration

Lower graph: ~50 sec with acceleration

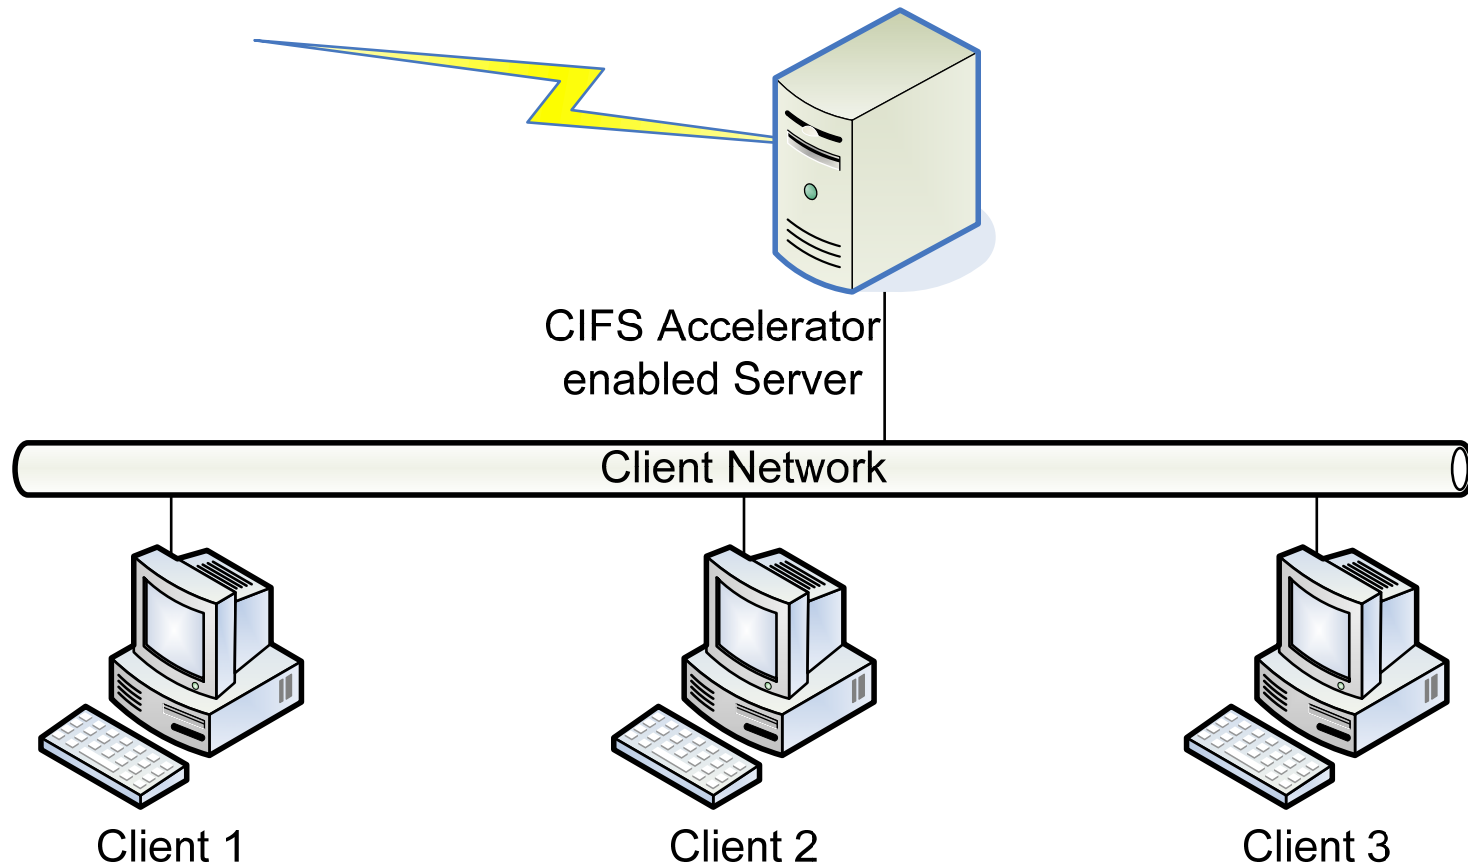
CIFS Accelerator Platforms

- Embedded into communication equipment



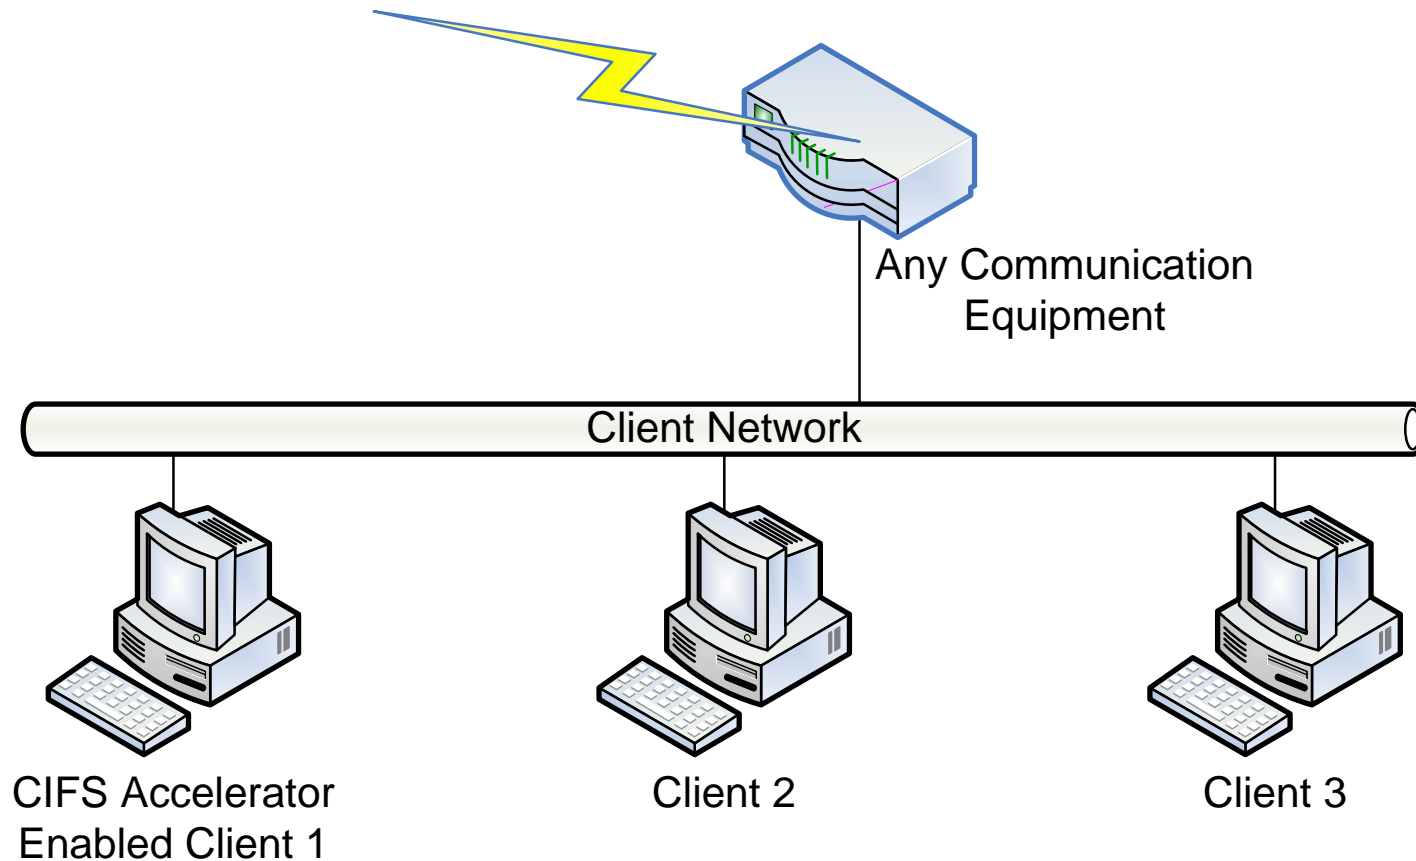
CIFS Accelerator Platforms

- ❑ Embedded into a gateway server



CIFS Accelerator Platforms

- Embedded into a client



Thank you

markr@visualitynq.com

www.visualitynq.com