LAPPEENRANTA UNIVERSITY OF TECHNOLOGY
Faculty of Technology Management
Department of Information Technology

Wide Area Network Acceleration in Corporate Networks

Examiners: Professor Jari Porras
M.Sc. Vincent Dethier

Supervisor: M.Sc. Vincent Dethier

ABSTRACT

Lappeenranta University of Technology
Faculty of Technology Management
Department of Information Technology

Maiju Kansanen

Wide Area Network Acceleration in Corporate Networks

Master's Thesis

2009

79 pages, 26 figures and 1 appendix.

Examiners: Professor Jari Porras
M.Sc Vincent Dethier

Keywords: Wide Area Network, Acceleration, performance, TCP

There are several factors affecting network performance. Some of these can be controlled whereas the others are more fixed. These factors are studied in this thesis from the wide area network (WAN) perspective and the focus is on corporate networks. Another area of interest is the behavior of application protocols when used through WAN. The aim is to study the performance of commonly used application protocols in corporate networks.

After identifying the performance problems in corporate WANs the thesis concentrates on methods for improving WAN performance. WAN acceleration is presented as a possible solution. The different acceleration methods are discussed in order to give the reader a theoretical view on how the accelerators can improve WAN performance. Guidelines on the installation of accelerators into a network are also discussed. After a general overview on accelerators is given, one accelerator vendor currently on market is selected for a further analysis.

The work is also a case study where two accelerators are installed into a target company network for testing purposes. The tests are performed with three different application protocols that have been identified as critical applications for the target corporation. The aim of the tests is to serve as a proof of concept for WAN acceleration in the target network.

TIIVISTELMÄ

Lappeenrannan teknillinen yliopisto
Teknistaloudellinen tiedekunta
Tietotekniikan osasto

Maiju Kansanen

Alueverkon kiihdytys yritysverkoissa

Diplomityö

2009

79 sivua, 26 kuvaa ja 1 liite.

Tarkastajat: Professori Jari Porras
M. Sc Vincent Dethier

Hakusanat: Alueverkko, kiihdytys, suorituskyky, TCP

Verkon suorituskykyyn vaikuttavat monet tekijät. Osa näistä tekijöistä on hallittavissa, kun taas osaan ei voida vaikuttaa. Tässä työssä näitä suorituskykytekijöitä tarkastellaan alueverkkojen (WAN) näkökulmasta keskittyen yritysverkkoihin. Toisena kiinnostuksen kohteena on sovellusten käyttö WAN-verkossa. Tavoitteena on tutkia yritysverkossa käytettävien sovellusprotokollien suorituskykyä, sekä havaita niiden mahdolliset heikkoudet.

Työssä keskitytään keinoihin, joilla voidaan parantaa WAN-suorituskykyä. Näitä tutkitaan havaittujen suorituskykyongelmien perusteella. WAN-kiihdytys esitellään yhtenä potentiaalisena ratkaisuna. Lukijalle pyritään antamaan teoreettinen näkökulma eri kiihdytysmenetelmiin ja siihen, millainen vaikutus niillä on WAN-suorituskykyyn. Työ antaa myös suosituksia kiihdyttimien asennukseen. Kun kiihdytys on esitelty yleisellä tasolla, valitaan yhden markkinoilla olevan kiihdytinvalmistajan laite syvällisempään tarkasteluun.

Työn empiirinen osuus on tehty kohdeyritykselle. Tässä osuudessa kaksi kiihdytintä asennetaan kohdeyrityksen verkkoon testausta varten. Testit suoritetaan kolmella eri sovelluksella, jotka ovat kohdeyritykselle kriittisiä sovelluksia. Testien tarkoituksena on todistaa WAN-kiihdytyksen toimivuus kohdeyrityksen suorituskykyongelmien ratkaisijana.

PREFACE

I want to thank my work supervisor Vincent Dethier for the valuable discussions and comments on the empirical part of this work. I also want to thank my work examiner Professor Jari Porras for his guidance and improvement ideas throughout the project. I am also grateful to my current employer for this opportunity, especially to Marcia Saecker for suggesting this interesting topic for the thesis. Furthermore I want to thank my family, especially Antti, for all the support they have given me during my studies and this project.

7th of May 2009

Maiju Kansanen

CONTENTS

SYMBOLS

Network (Wide Area Network) Cloud

Router

Server

Firewall

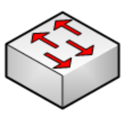

Switch

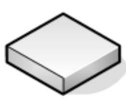

Accelerator

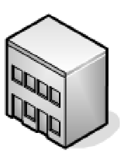

Office

Laptop

Workstation

ABBREVIATIONS

ACK	Acknowledgement
ATM	Asynchronous Transfer Mode
CAD/CAM	Computer-aided Design and Computer-aided Manufacturing
CDN	Content Delivery Network
CIFS	Common Internet File System
CoS	Class of Service
CPU	Central Processing Unit
CRM	Customer Relationship Management
CSS	Cascading Style Sheet
CWND	Congestion Window
DifServ	Differentiated Services
DSL	Digital Subscriber Line
FDDI	Fiber Distributed Data Interface
FTP	File Transfer Protocol
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
IETF	Internet Engineering Task Force
IEEE	Institute of Electrical and Electronics Engineers
IntServ	Integrated Services
IP	Internet Protocol
IPFIX	IP Flow Information eXport
ITU	International Telecommunication Union
ISO	International Standard Organization
LAN	Local Area Network
LFN	Long Fat Network
MAPI	Messaging Application Programming Interface
MPLS	Multi Protocol Label Switching
MSR	Molecular Sequence Reduction
NBAR	Network Based Application Recognition

NFS	Network File System
OSI	Open System Interconnection
PBR	Policy Based Routing
QoS	Quality of Service
RFC	Request for Comments
RWND	Receive Window
RPC	Remote Procedure Call
RTT	Round Trip Time
SDH	Synchronous Digital Hierarchy
SLB	Server Load Balancer
SMB	Server Message Block
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SONET	Synchronous Optical Network
SSL	Secure Sockets Layer
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
URL	Uniform Resource Locator
VoIP	Voice over IP
VPN	Virtual Private Network
WCCP	Web Cache Control Protocol
WAFS	Wide Area File Services
WAN	Wide Area Network
WLAN	Wireless Local Area Network

TERMS

Ethernet - Ethernet is the most commonly used local area network (LAN) technology in today's networks. The most recent one of the Ethernet standards is the 10 Gigabit Ethernet over Unshielded Twisted Pair copper cabling (IEEE 802.3an). [1]

Local Area Network – LAN is a computer network in a limited geographical area, e.g. a building or a group of buildings. LANs typically have high data-transfer rates. The network consists of cables, network devices, workstations and servers. Typical LAN devices are hubs, repeaters, bridges, LAN switches, routers and LAN extenders. [2]

Open System Interconnection Reference model – OSI is a reference model for describing how information is transferred between two applications in two different computers through a telecommunication network. OSI reference model has 7 layers: physical, data link, network, transport, session, presentation and application layer. [3]

Packet – In packet switched networks the data is transferred in the network as packets. The packets contain of control information, like destination and source address, and the actual data from the user. The control information is referred as the packet header. Packet is the term used with TCP connections, with UDP the packets are referred as datagrams. [4]

Transmission Control Protocol – TCP is a protocol which is used in network communication for sending data between two nodes. In TCP is a connection-oriented protocol. It uses sequence numbers and acknowledgements to make the data transmission reliable. [5]

User Datagram Protocol – UDP is used for sending data between two nodes as TCP. Compared to TCP, UDP is a simpler connectionless protocol. [3]

Wide Area Network – WAN is a computer network that is spread on a broad geographical area e.g. between two cities, countries or even between continents. WANs are used to connect LANs together. WANs can be either point-to-point links, also known as leased lines, circuit switched networks or packet switched networks, also referred to as clouds. [6]

1 INTRODUCTION

Network is one of the most important tools for a corporation to run their operations. Many business critical functions rely on the corporate network. During recent years there has been a switch from local area networking to wide area networking. Because most of the applications used today are designed to be used in LAN, the transition to WAN will have an impact on application performance. Also the amount of data transferred through networks is growing at the same time. Due to these reasons network performance has received a lot of attention.

This thesis studies the network performance in corporate wide area networks. The methods for improving WAN performance are discussed, focusing on WAN acceleration. The work is done for a target company, which has a need for improving WAN performance in their network. The company has done a lot of background research on WAN accelerators (more details in [7]) and selected one of them to be tested in their network. The empirical part of the thesis focuses on these tests.

This section presents the outline of the thesis also the goals and the framework of the thesis is discussed. In addition, the most important research questions are presented.

1.2 Outline

In chapter 2 network performance is discussed from the wide area network (WAN) perspective concentrating on corporate networks. Factors affecting network performance are studied. The chapter concentrates also on the performance problems caused by Transmission Control Protocol (TCP).

In chapter 3 the focus is on the performance problems that the applications and application protocols used in corporate networks are facing. In this chapter the applications suffering most from the transition to WAN are recognized and discussed.

Chapter 4 is an introduction to methods improving WAN performance. It studies WAN acceleration techniques and also looks into quality of service (QoS). Also network monitoring is discussed.

Chapter 5 continues the discussion on WAN accelerators started in chapter 4. One WAN accelerator device currently on marked is studied in detail. Chapter also focuses on how the accelerators can be implemented to the network.

Chapter 6 is the empirical part of the work. It is a case study for implementing accelerators into a network for testing purposes. First the critical applications of the target network are defined. After that the test plan and the results of the tests are discussed.

1.3 Goals and Framework

The goal of this thesis is to provide a clear view on the performance problems of the corporate networks and to offer a possible solution for overcoming these problems. The main focus is on WAN acceleration and its capabilities as a solution for improving WAN performance. WAN accelerators are first studied in theory but the main focus is on the implementing and testing these accelerators in a live environment. There are currently around 10 different WAN acceleration technologies in the market but this thesis focuses on one vendor. The vendor is selected based on the research presented in [39].

As WAN acceleration is applied to the applications using the network, one goal of the thesis is find a way to recognize the critical applications for corporate networks. This is done by studying the qualities of corporate applications in theory and then reflecting the results on the applications used by the target company.

1.4 Research Questions

The aim of these first research questions is to help in defining the problem that will be discussed in this thesis. These questions are:

- *Why network performance is important for corporations? (chapter 2)*
- *What are the factors affecting network performance? (chapter 2)*

This thesis has 2 main research questions. The first main research question in this thesis is related to the applications used in the network. This is studied first in theory in chapter 3 and then in practice in chapter 6. The research questions related to this are:

- *What are the affects of the transition from LAN to WAN for corporate applications? (chapter 3)*
    - *What are the critical applications for the target network? (chapter 3)*

An important part of this thesis is to introduce solutions for improving WAN performance. As already mentioned WAN acceleration is the main solution studied in this thesis. This can be studied by answering the following questions:

- *How is it possible to overcome the performance problems in WAN? (chapter 4)*
- *What kind of WAN acceleration methods are there? (chapter 4)*
- *What needs to be considered when selecting and accelerator for a specific network? (chapter 5)*
- *How can wan accelerators be implemented to the network? (chapter 5)*

The second main research question is related to the empirical part, chapter 6. It is important to find out the potential of WAN acceleration. This can be done by answering to the following questions:

- *How does the WAN acceleration improve the performance of these applications?(chapter 6)*
    - *How this can be tested? (chapter 6)*

2 NETWORK PERFORMANCE

Often the term *performance in networks* is used to describe the performance of applications and to how the end users experience it. In this thesis the network performance is studied from the application performance point of view. The focus is on the network performance in WAN, although many of the characteristics can be also found in LAN. First the reason why application performance is important for corporations is given. Then the different factors affecting to network performance are presented. Cole and Ramaswamy state in [8] that the factors affecting network performance are bandwidth, latency, throughput, congestion and packet loss. These are the factors discussed also in this thesis. This chapter concentrates also on the problems caused by TCP.

2.1 Corporate Networks and Performance

From the network perspective global corporations have at least 3 different kinds of offices inside the network: branch offices, regional offices and datacenters. Many companies prefer to have the employees as close to the customer as possible and therefore the number of small branch offices within one corporation can be high. The different offices are connected to each other with WAN links. Branch offices or remote locations are connected to a bigger regional office or a datacenter usually with low bandwidth high latency WAN links. Therefore branch offices are the ones suffering most from a poor performance. [9]

Figure 1 shows an example of how a corporate network can be constructed from the WAN and the LAN part. Corporate WAN consist of the carrier backbone and the point-to-point links between the carrier cloud and different office building LANs. All of these are separated from each other by routers. The point-to-point links can be Digital Subscriber Line (xDSL) or Multi Protocol Label Switching (MPLS) links, for example.

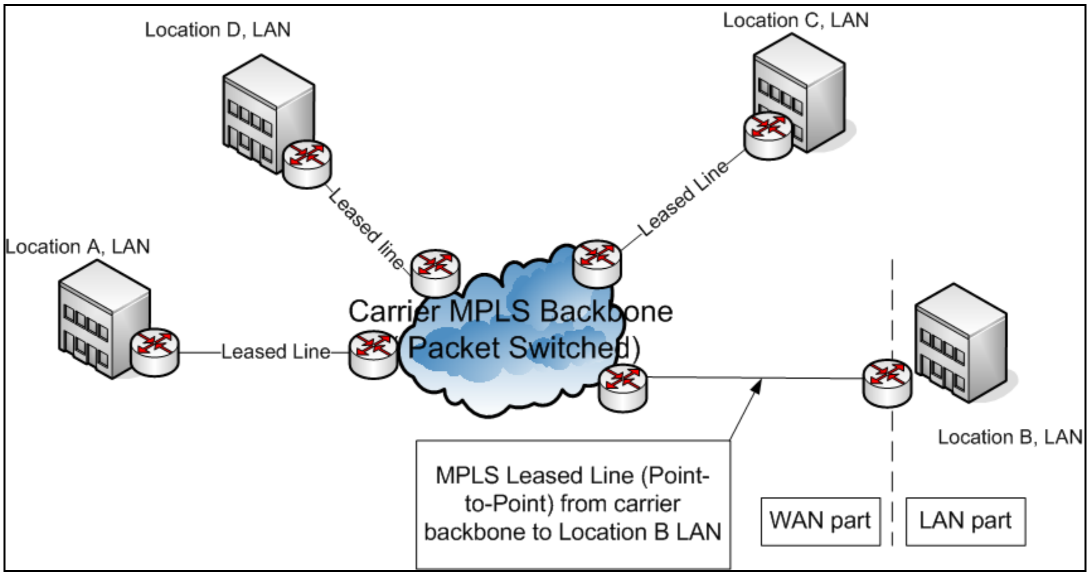

Figure 1: An example of a corporate network (WAN + LAN). The WAN consists of the carrier network cloud and the point to point links between the different offices and the carrier cloud. [1]

Another important factor affecting corporate network performance are the applications used through the network. In [1] Grevers and Christner state that largest part of the network costs in a global corporation comes from the servers, deployed infrastructure and the management of the servers in these branch offices. Only from the performance perspective it is ideal for the branch office workforce to have a local server infrastructure for applications. To reduce the costs and ease the management many companies have moved these local services into global datacenters. Due to this consolidation of services the branch office users have experienced lower application performance. The large amount of small offices and at the same time the centralization of servers into datacenters has made the network a critical point for the corporations.

2.2 Bandwidth

Bandwidth describes the capacity of the communication link and it is usually measured are bits per second (bps) [8]. As an example, corporate LANs can have bandwidths from 100Mbp to 1Gbps, where as WAN links can be 2-50 Mbps. The small branch offices can have WAN connections of 1Mbps or even less. Due to the slower transmission rate, bandwidth usually causes performance problems only in WANs. The difference in bandwidth creates a bottleneck between the LAN and the WAN. These problems, oversubscription and utilization, and are explained in the sections 2.2.1 and 2.2.2.

2.2.1 Oversubscription

The performance related problem caused by different bandwidths used inside network is called oversubscription. When data is coming from a higher rate network to a lower rate network the lower rate part might get oversubscribed. Oversubscription causes queues in the devices, usually routers, handling the change of bandwidth and thus slows down the transmission. [8]

Figure 2 illustrates the possible points of oversubscription in a WAN. As an example let us assume that the LAN part of the connection is 100 Mbps Fast Ethernet and the WAN part is constructed with a point-to-point T1 (1,544 Mbps) links. These two parts are connected with routers to each other. The first point of oversubscription can happen already inside the LAN when the three client workstations (each having a100 Mbps link to the switch) are there joined to one circuit of only 100 Mbps. Connecting three lines into one in the LAN causes 3:1 oversubscription as can be seen from the figure. When moving from LAN to WAN the difference in the network speed means an oversubscription of 67:1 (100 Mbps / 1,544 Mbps). From these two the oversubscription presented in LAN is not causing as much performance problems as the oversubscription when entering the WAN. The LAN oversubscription might not even be present at all times since the three workstations might not be in use at the same time.

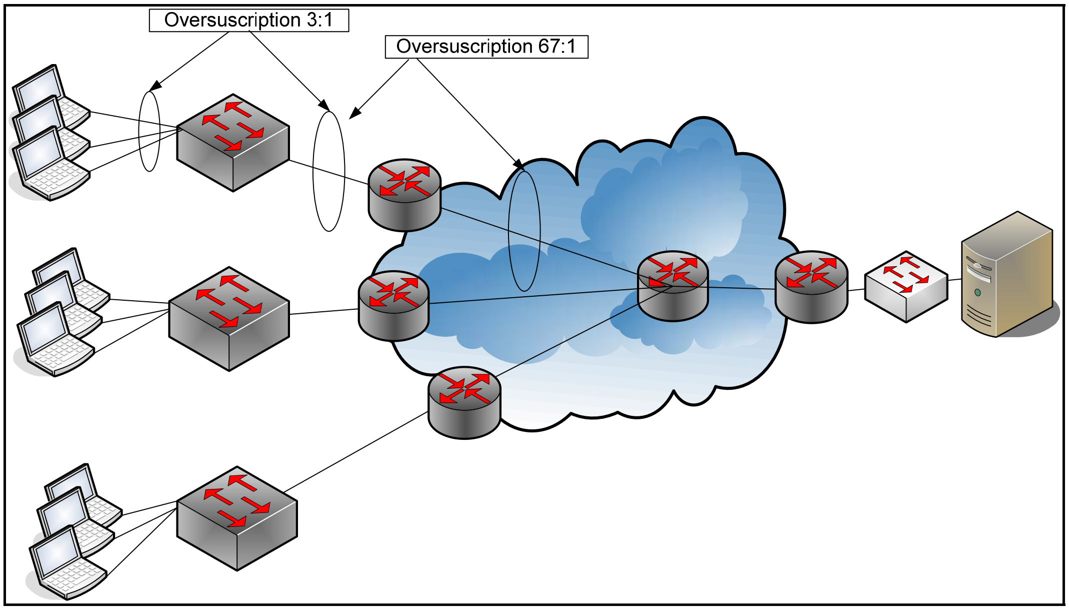

Figure 2: Network Oversubscription. Change from LAN to WAN causes oversubscription due to the change in the transfer rate. In this example the move is from 100Mbps to 1,5Mbps. [1]

2.2.2 Utilization

Utilization is a way to measure the performance of the network. Utilization percentage describes how much of the total network bandwidth, network capacity, is in use at a certain point of time.

The best way to detect over utilization in the network is to monitor it constantly. Measurements can be taken for instance from 8 working hours, the busiest hour of the day and from the busiest 15 minutes of the day. Charles Spurgeon in [10] gives an example for these reference values: for the 8-hour working day utilization should be 20% or less and for the busiest hour the average utilization should not exceed 30%. In addition the 15 minute average should be less than 50%. In [6] Bartlett, Sevcik and Moore state that overall utilization should be most of the time less than 35%. It is hard to define when the performance is good - these reference values might change based on the network and how it is used. One important factor in defining whether the network is over utilized is to get the opinions of end users. In general the reference values are based on the examinations of the statistical properties of the traffic. In this thesis these statistical properties are not discussed in detail. More information on the topic can be found in [11] and [12].

2.3 Latency and delay

Latency is used for describing the time that data travels in the network. It can be expressed as one way latency or as roundtrip latency. One way latency, also known as delay, simply means the time it takes for data to travel from the transmitting node to the receiving node. Roundtrip latency, also known as Round Trip Time (RTT), measures the time data travels from transmitting node to the receiver plus the time that it takes for the transmitting node to get a response (acknowledgement) from the receiving node. When studying application performance, the most commonly used form of latency is the RTT. [8] [3]

Latency can be divided into smaller delay components that together generate the over all network latency. These components are propagation delay, serialization delay, processing delay and forwarding delay. The following sections discuss the different components. It is important to make the distinction between the different types of delay, since some of them are fixed and some are at least partly controllable. Understanding what parts of the delay can be controlled becomes important when talking about improving the network performance. Figure 3 below has an example of a network showing how the different parts of the delay affect the over all network latency.

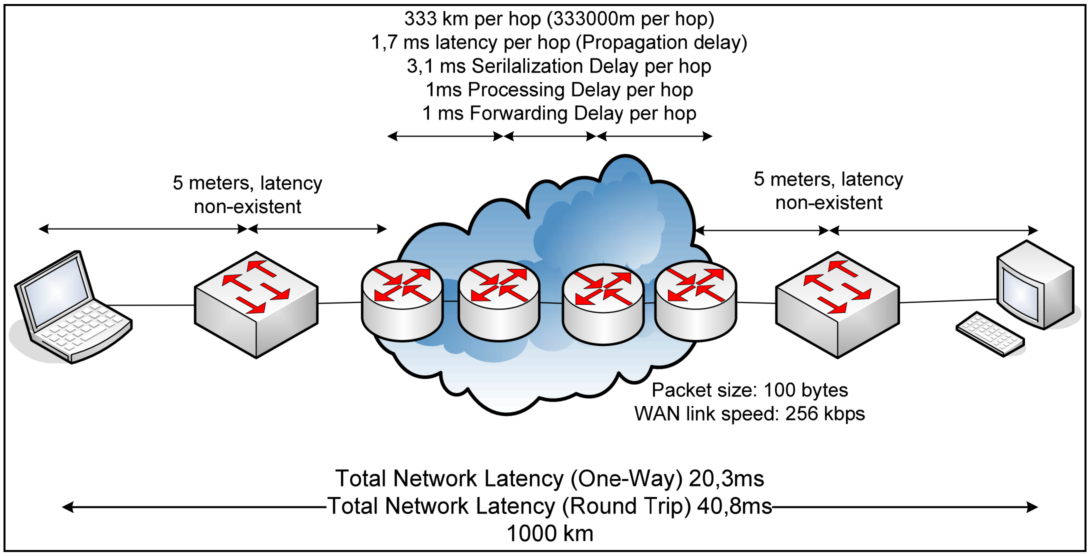

Figure 3: Example of how the over all network latency is constructed from different. The different types of delay are presented per hop and the network consists of 3 hops inside the WAN. The total latency per hop is approximately 6,7 ms which created a 20,3 ms one-way latency. [1]

2.3.1 Propagation Delay

Propagation delay is a form of delay that is caused by the distance between nodes and physics in terms of how fast data can be transferred in the network. Propagation delay is one of the fixed factors affecting the overall network latency. Propagation delay is measured as the time the data packet spends to go through the network. The speed packets can be transferred is called *propagation velocity* and it is normally around $2/3$ of the speed or light ($3x10^8 m/s$). Propagation delay becomes significant in long distances, which is usually the case in WANs. In figure 3 for example a connection of 999 km of distance between the sending and receiving end, including three 333 km hops, would have the propagation delay of approximately 5ms and RTT of 10ms. Propagation delay T_p is defined as

$$T_p = d / v,$$

where d is physical distance (m) and v is propagation velocity $(\sim 2x10^8)$. [8] [3]

2.3.2 Serialization Delay

Serialization delay is measured as the time it takes to move bits of a packet into the line. It consists of the size of the packet, network medium and speed of the interface. Usually serialization delay is more significant in lower-speed networks. In figure 3 above we assume the line speed to be 256kbps and the packet size 100 bytes. This creates a serialization delay of 3,1ms per hop. As an example of a high-speed network in a 1Gbps link the same delay would be 100 ns. Serialization delay T_s is defined as

$$T_s = s / r \, ,$$

where s is size of packets (bits) and r is transmission rate (bps). [3]

2.3.3 Processing Delay

Time that it requires for a network node like a router, switch or for example accelerator to perform required actions on the packet is called processing delay. In a router processing delay means the comparison of a piece of data to the access list. The forwarding architecture has to also be counted in processing delay: the node can either wait until the entire packet is received before it makes any decisions what to do with it (store and forward) or the forwarding of the packet can start as soon as the header is received. For instance, in a router the processing delay can vary between less than 1ms to even 10ms when the router is congested. In the figure 3 the processing delay is estimated to be 1ms for each hop in the network, making a total of 3ms end to end processing delay. [3][8]

2.3.4 Forwarding Delay

In routers or switches the time spent on deciding where to forward the packet. For example in a router the packet could go through in 1 ms and under a load it could take 3 to 5 ms for the same job. In figure 3 the forwarding delay is estimated to be 1ms per hop creating a total of 3ms forwarding delay for the end to end connection. [3]

2.4 Throughput

Throughput is the rate of successful data transfer in the network. In [3] Grevers and Christner define throughput as a sum of 3 different parameters: network capacity, latency and packet loss. Capacity means the maximum amount of information that can be transferred between two network nodes. The throughput of a network is never more than the capacity of the slowest hop within that network. For example if the network connection inside a branch office is 1Gbps and it is connected through a 1,5Mbps WAN link to the datacenter router and again the datacenter devices are connected to each other with a 1Gbps connection, the throughput of the connection would never exceed the 1,5Mbps. For throughput latency means the time it takes to transfer data between two nodes and also the distance of the nodes. Packet loss adds the element of lost data to the sum, how many packets are dropped for example due to congestion. In addition to these 3 parameters, when talking about throughput for an application, the transport protocol itself can limit the throughput. In cases like this adding more capacity to the network would not necessarily improve the throughput at all.

2.5 TCP as a Performance Factor

This chapter focuses on TCP as an application performance factor. The possible performance problems caused by TCP are related to the amount of round trips the protocol uses in communication and also to the way the protocol is constructed. Because TCP was developed to use in LANs, these performance issues exist mainly in WAN [3].

As TCP is a connection oriented protocol, it has to establish a connection between the sending and the receiving node before any data can be sent. This takes 1,5 round trips (RTT). After the connection establishment a number of protocol and application messages are exchanged between the sender and receiver and all these messages add latency usually with one round trip. The amount of roundtrips that a protocol uses for the communication might have an effect on the over all network latency when distances are long. For example in figure X in the previous chapter this would mean every time a 32,1ms (one RTT) delay for each sent message. This is not the biggest factor in application performance but it adds some overhead. [1]

In addition TCP has 2 mechanisms that affect application performance - slow start and congestion avoidance. TCP uses these two mechanisms to determine the capacity of the network and to adapt to possible changes in the network. Slow start is a mechanism that recognizes the network capacity. TCP enters the slow start phase when initializing the connection. In slow start phase a variable called congestion window (CWND) defines how many packets can be sent through the connection on every RTT. The congestion window is maintained by the sending node. The data transmission starts with a CWND of one packet and it is increased with one packet *for every received acknowledgement*. This makes the CWND to grow exponentially per RTT:

$$CWND_{i+1} = 2xCWND_i \ .$$

The connection stays in the slow start phase until congestion (packet loss) occurs. When a packet loss has been encountered for the first time it is assumed that the maximum capacity of the network has been reached and TCP enters the congestion avoidance phase. At the same time the CWND size is dropped to

$$CWND_{i+1} = CWND_i \ / 2 \ .$$

In the congestion avoidance phase the CWND is increased by 1/CWND *for each received acknowledgement*. This means that the CWND is increased with one packet per RTT:

$$CWND_{i+1} = CWND_i + 1 \ .$$

Also in this phase, whenever congestion occurs, the congestion window size is dropped to 50%. The transition from slow start phase to the congestion control phase is illustrated in figure 4. [3] [9] [13] [14] [15]

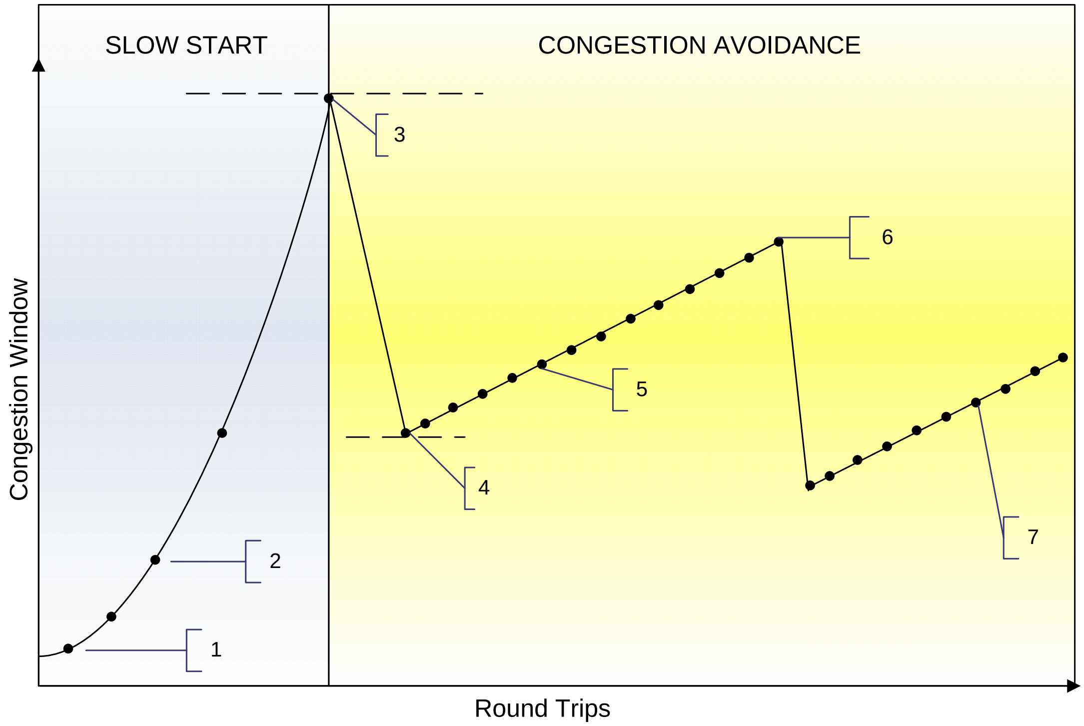

Figure 4: TCP slow start and Congestion avoidance. 1) Small CWND / low throughput during this period, 2) Exponential slow start 3) Packet loss identifies network capacity, 4) When capacity is identified throughput (CWND) drops by 50% and TCP enters congestion avoidance, 5) In congestion avoidance (+1/CWND per ACK), 6) Packet loss occurs -> congestion window is dropped 50%, 7) Returning to maximum throughput can take long time [3]

Also the receiver buffer size can be a limitation to the TCP performance because there cannot be more data in flight than the receiver buffer can handle. The receiver uses a receive window (RWND) to tell the sending node how much data the buffer can handle. The RWND size is reduced based on how full the buffer is. [13]

Slow start and congestion avoidance are useful in many ways but can cause problems in networks with high bandwidth and high latency, where it can take a long time to get the window size to the maximum so that most of the available bandwidth could be used. These types of networks are commonly known as long fat networks, LFNs. The term high speed wide area network is also used to describe the networks that can suffer from the poor performance due to these TCP characteristics. Wide area high speed networks are characterized with speed higher than 100Mbps and RTT over 50ms. [3] [13] [9]

To get an idea of how large influence the two TCP mechanisms have on the throughput of a connection we can calculate the average congestion window. Average congestion window (packets) is defined as:

$$W = BR/(8D)$$

where B (bps) is throughput, R (s) is RTT and D (bytes) is packet size [16].

As an example let us assume that a connection has a RTT of 100ms and a packet size of 1500 bytes. For a 1Gbps bandwidth the average congestion window would be 8333 packets. In case of congestion the CWND drops to about 4170 packets which corresponds a transmission rate of 500Mbps. In the congestion avoidance it would require about 4170 round trip times to achieve again the transmission rate of 1Gbps, and it would take 7 minutes to do that. To improve the recovery time from congestion in WANs, modifications to the slow start and congestion avoidance mechanisms like scalable TCP, by Kelly [13], and high-speed TCP, by Floyd [16], have been suggested.

3 APPLICATION PERFORMANCE

This chapter discusses the WAN performance problems studied in chapter 2 from the application perspective. First the section looks into performance requirements the applications have. Also some common protocols and applications used in the corporate networks and the performance problems they are facing in WANs are examined.

Applications use their own protocols, like Hypertext Transfer Protocol (HTTP), on top of TCP. These application protocols usually add extra messages to the data transfer since they normally define how the data exchange is done. This means that there are multiple messages sent over the network before any actual data is sent. In [1] Grevers and Christner use term *chatty* to describe this kind of application protocols. In short distances this chattiness of an application does not cause problems due to small latencies, but as distances between the sending and receiving node grows also the latency increases and it takes more time for messages to go through the network. Most commonly used chatty application protocols are HTTP, Common Internet File System (CIFS), Network File System (NFS) and Remote Procedure call (RPC). These protocols are covered later in this chapter.

.

3.1 Performance Requirements

Applications have different kind of requirements for the network. An application can for example demand low latency or then it can require a faultless transfer to work well in the network. It is important to find out what kind of bandwidth the application needs and what kind of latency and error rate are acceptable. This section discusses the performance requirements of applications.

3.1.1 Bandwidth and Applications

All WANs, when compared to LANs, can be said to be limited in bandwidth. If there is not enough bandwidth for the application, the affects on performance can be drastic. For applications that send large amounts of data in the network the bandwidth easily becomes an issue. In [11] Siegel has collected a set of questions that can help in identifying these bandwidth related problems that applications have:

- Amount of data transferred in both directions (download / upload)?
- How much repetitive patterns data contains?
- How much of the data will be sent only once through the network?
- Is there a variation in the rate data is sent to the network?

Answering to these questions helps in determining the best possible solution to overcome the performance problems. One of the options is to use compression to reduce the size of data sent across the network. These solutions, including compression, are discussed in the chapter 4.

3.1.2 Application Data

Studying what kind of data the application is sending is important when it comes to improving application performance. A large part of the data going through a WAN link is repetitive. Repetitive data is usual in the Internet, intranet and some client-server applications. In a Web-page, for example, this would mean a photo or instructions on how the page is formatted. In a client-server application repetitive data can be for instance a particular file that is sent over the network many times. There can be a lot of repetitive data traffic in corporate networks. Multiple users can use same files during one business day or an email can be sent to whole department at the same time. [17]

3.1.3 Latency and Applications

Latency affects some applications more than others. In [11] Siegel mentions Voice over IP (VoIP) traffic as an example of an application that is affected easily by latency. With VoIP as users expect to hear the person at the other end of the line in real time. The International Telecommunication Union (ITU) has published a recommendation ITU-T

G.114 [18] on one-way transmission times. For VoIP the G.144 sets recommendation 150ms or less as a good latency level and 300ms as an acceptable level for really long distance calls.

Latency can be also a problem for protocols that use many round-trips for simple functions like initializing the connection. For example if an application needs 20 round-trips to perform certain function with a one-way delay (latency) of 100ms, the total delay would be 4 seconds. As users are used to delays of only milliseconds they might think that the application has failed to perform the task. Due to poor protocol design it is not unusual for an application to take hundreds of even thousands of round-trips for relatively simple tasks. [17]

3.1.4 Error Rates for Applications

Many application protocols require low packet loss to function well. TCP has a way to recover from packet loss or incorrect packets but as discussed in the previous chapter 2 this slows down the network. VoIP and multimedia applications can handle a small amount of errors but a high error rate in a VoIP conversation can make the talking impossible to understand. Due to the real-time nature of VoIP the resend function of TCP is not suitable for it. Error rate analysis can be done to estimate the error sensitivity of an application. [17]

3.2 Applications and Application Protocols

In this chapter common application protocols and applications used in corporate networks are presented. As applications work in different manners it is essential to know how the application protocols are constructed. TCP and its affects to the performance of applications were discussed in the chapter 2. This section concentrates on the protocols that work on top of TCP. Main characteristics of each protocol and their capability to work in WAN are discussed to get an understanding of the possible affects on the application performance.

3.2.1 Messaging Application Programming Interface - MAPI

Messaging Application Programming Interface (MAPI) is a protocol that is usually used together with some other protocol like Remote Procedure Call (RPC). MAPI is used for directing messages but RPC is actually the protocol used for the message transfer. Microsoft Exchange is an example of an application that functions in this way. MAPI breaks down the email messages into small data blocks for sending across the network. Each block has to be acknowledged before sending the next block. Therefore sending one email through the network requires thousands of RTTs. The way MAPI works is different within different versions of the Microsoft Exchange. The recent versions have improvements to reduce the performance effects of the protocol when used in WANs. Microsoft Exchange 2003 for example uses local cached copies of data in order to reduce unnecessary traffic over WANs. There are also many other applications that use MAPI in similar ways. [3][19]

3.2.2 Common Internet File Service – CIFS

Common Internet File Service (CIFS) is a protocol used for instance for accessing, reading and writing remote files from a server. It is based on Server Message Block (SMB), a protocol developed by Microsoft. It also has a method for controlling file permissions. CIFS was designed for LANs for controlling local file shares that users can access. CIFS requires a lot of message exchange in order to ensure correct usage of the files. As an example Grevers and Christner state in [3] that, opening a 1,5 MB file requires over 1000 message exchanges. The amount of messages is related to the multiple tasks CIFS is handling. These are for example: user authentication, fining correct disk share, checking user permissions, asking which file to open and handling file locks. The file is also broken into multiple data blocks, which all need to be send one at the time through the network. Server only sends the next block after it has received the acknowledgement form the client. Due to this considerable amount of messages that are also sent one at the time, CIFS is said to be a chatty protocol. Therefore, the performance of CIFS suffers when used in WAN, and increased latency reduces performance even more. CIFS contains some optimization methods that reduce the amount of messages during the use of files [3]. Figure 5 presents an example of the messages that can be exchanged between the client and the server while a user opens a file from the file server.

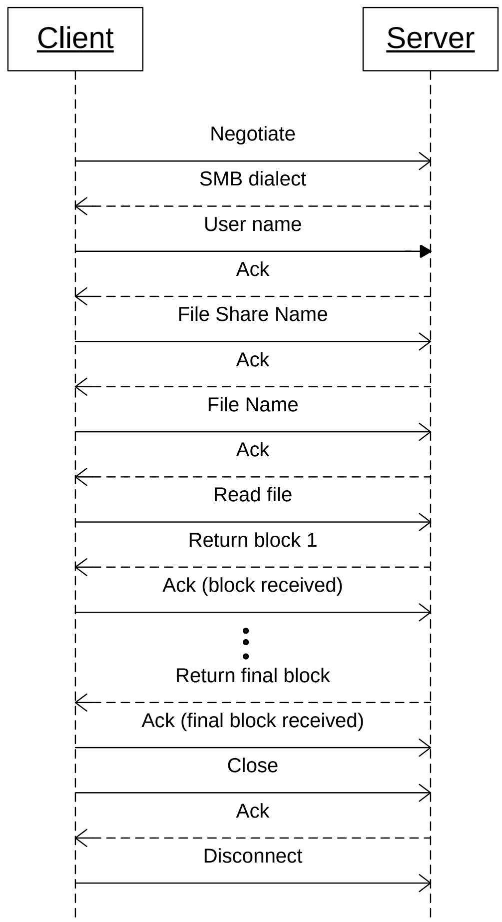

Figure 5: CIFS – a rough illustration of the messages that can be changed during opening a file from a server. First steps are: user authentication, fining correct disk share, checking user permissions, asking which file to open. After that the data can be exchanged, the number of messages related to the data exchanged is directly linked to the file size. [20] [21]

3.2.3 Hypertext Transfer Protocol – HTTP

Hypertext Transfer Protocol (HTTP) is the protocol used for Internet traffic. Web pages contain typically many small embedded objects like pictures, JavaScript codes or Cascading Style Sheets (CSS). In order to transfer the objects the protocol might have to open additional TCP connections. Both the objects and the TCP connections needed for the transfer add additional round trips for opening the web-page. When a web-page is opened for the first time, an HTML file is retrieved first. The HTML file contains links to all other objects needed for displaying the web page completely. Usually the objects are downloaded one at a time and the next object will be downloaded only after the current object has been received completely. Sending and receiving requests one at a time has an effect on the application performance in networks where latency is high. HTTP uses local caching and the second time the same HTML file is downloaded HTTP gets the objects from the local cache. This has some affect to the performance although HTTP still needs to check from the server that the objects in the cache have not been changed. This results in the same amount of RTTs as fetching the objects. HTTP has become a popular transport method within corporate applications. For example, applications that area actually client-server based can use web-browser and HTTP to deliver the application to remote workstations. [3][22]

3.2.4 Secure Sockets Layer - SSL

Encrypting critical business processes has a significant role in corporate networking. More and more applications are used through the web and require encryption. There are many possibilities for improving the security of these applications. Secure Sockets Layer (SSL) has become the most used. SSL can be used with protocols like HTTP, File Transfer Protocol (FTP) and Simple Mail Transfer Protocol (SMTP). SSL was developed first by Netscape. IETF has published Transport Layer Security (TLS) standard, RFC 2246 [23] and RFC 5246 [24], which is based on the latest version of SSL (3.0). In [25] and [26] Chou discusses the two main phases of SSL: handshake and data transmission. Handshake is done between client and server and the purpose of the handshake is to determine the secret-key parameters. These secret keys are then used during the data transfer to encrypt and decrypt the data sent over the network. The encrypted data cannot be seen by network devices like routers or accelerators. This makes analyzing or for example prioritization of the SSL encrypted data difficult. SSL also demands a lot of processing power, especially during the handshake phase, which also affects the performance. BlueCoat mentions in [27] that one of the downsides of SSL is that also spyware and peer-to-peer applications, like instant messaging, exploit the secure SSL tunnel.

3.2.5 SAP

SAP is one of the most used business applications. There are two graphical user interfaces (GUI), SAP-GUI (also known as SAP R/3) and SAP-Web (also known as MySAP). Currently SAP-GUI, which is also the original GUI, is the more popular one. In general SAP is said to have a better performance over network than many other applications. Sevcik and Wetzel discuss, in [28], SAP and the performance problems it is facing in WAN. SAP performance can become a problem due to congestion, latency or denial-of-service attacks. Latency, for example, can become a problem for SAP due to the fact that it normally is a centralized service, served from the data centers to the users.

3.2.5 Measurements on Application Performance

Figure 6 below summarizes, what has been discussed in this chapter. The figure presents measurements on the performance of different applications in WAN conducted by Sevcik and Wetzel in [21] and [28]. The figure illustrates the chattiness of CIFS, MAPI, HTTP and SAP. As basis for the evaluation, Sevcik and Wetzel have done measurements on actual data on typical user tasks performed with these applications. For example, for CIFS, the tests included opening of Microsoft PowerPoint files ranging in size from 10KB to 1MB. For MAPI, loading emails 10KB-1MB from Microsoft Exchange Server. From the X-axis “Payload per Task (bytes)” it is possible to see the amount of bytes required to transfer through the network per task per application. From the Y-axis “Turns per Task” it is possible to see the number of client-server interactions (one turn consists of 2 messages sent over the network) needed for performing the task e.g. opening of a file. The results of the measurements back also up the statement discussed in chapter 3.2.2 about CIFS requiring more than 1000 messages for opening a 1,5MB file. The measurements here show that for opening a 1MB file with CIFS it requires approximately 800 turns (1600 messages) over the network. From the figure it is also possible to see that CIFS suffers most and SAP GUI least from the centralization of services into data centers. There seems to be also some difference in the way different protocols behave when the payload per task grows, MAPI seems to handle the growth in the task size better than CIFS or Web traffic.

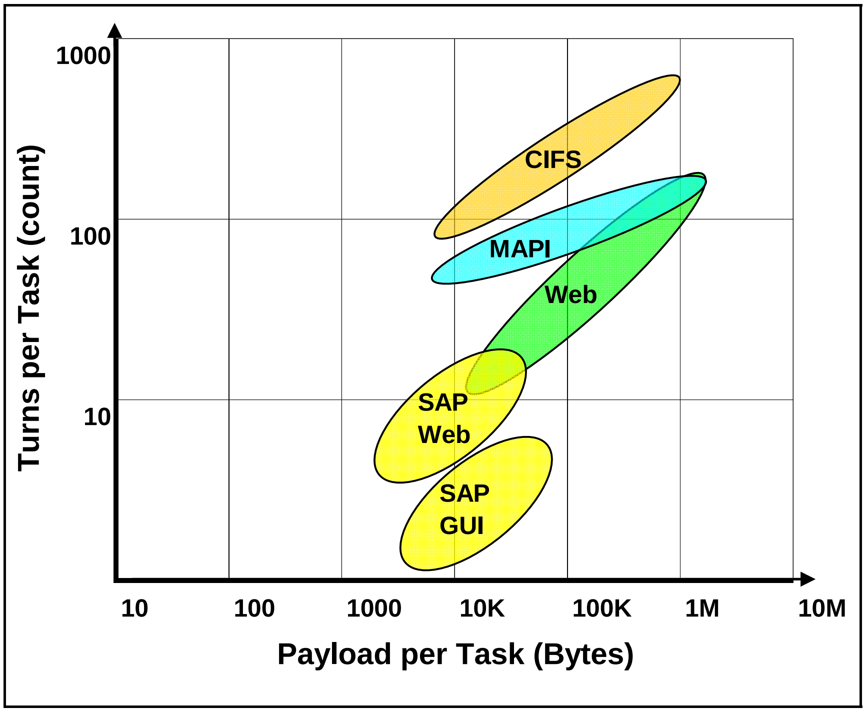

Figure 6: Chattiness of Applications. CIFS traffic is the chattiest of the protocols. For example opening a 1M file requires around 800 client-server interactions. In long distances, like WAN, this can cause declaration in the application performance. NOTE: scale of the graph is logarithmic.

4 IMPROVING WAN PERFORMANCE

Chapter 2 discussed the general WAN performance problems, for instance problems caused by latency. Chapter 3 concentrated on the performance problems that different application protocols have when used in WANs. This chapter presents *WAN acceleration* as a solution for improving WAN performance. Improving WAN performance can be divided into two categories: Quality of Service (QoS) and WAN acceleration. The main focus in this thesis is on WAN acceleration. The term *WAN accelerator* is used to describe a device designed to accelerate the applications used in the network. Also terms *WAN optimization* (chapter 4.2.1) and *application acceleration* (chapter 4.2.2) are used to refer to the services provided by the WAN accelerators. Because many of the WAN accelerators are capable of performing QoS, it is discussed in this chapter. Figure 7 below illustrates how the chapter is constructed and also how the different terms relate to each other.

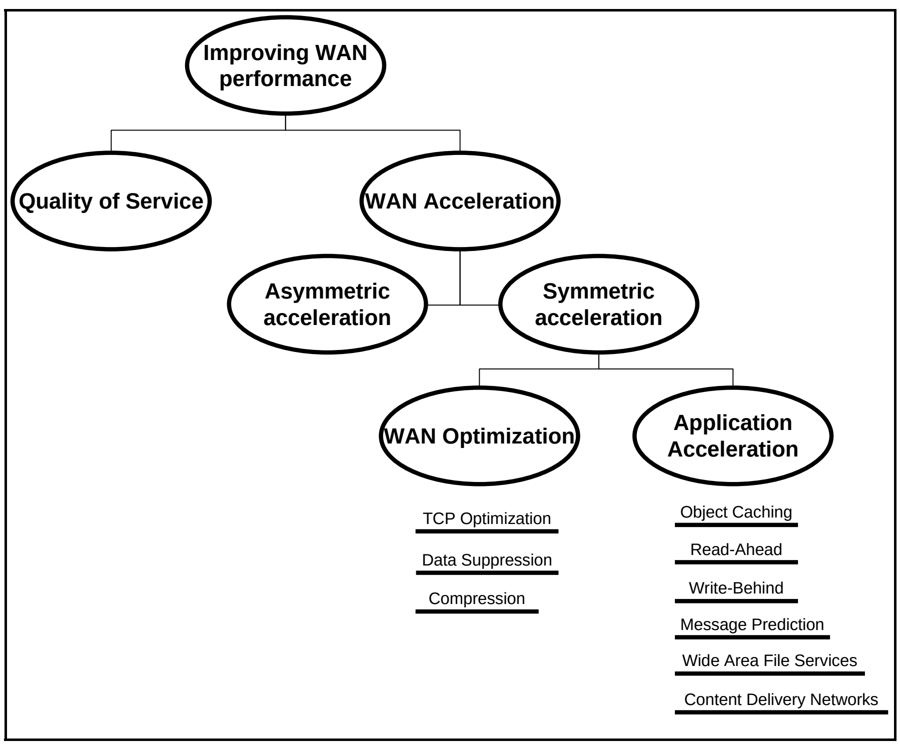

Figure 7: Improving WAN Performance - Deviation of WAN acceleration into asymmetric and symmetric acceleration and further the deviation of symmetric acceleration into WAN optimization and application acceleration.

Rolfe and Skorupa discuss the history of WAN acceleration in their article in [29]. They state that the first step in WAN acceleration were Wide Area File Services (WAFS) that focused on optimizing file access over the WAN. First vendors in WAN application acceleration market concentrated either on traffic shaping and QoS or on caching and compression. Especially the development of compression techniques and reduction in disk storage prices made it possible to create cost effective solutions for WAN. Compression had been around before but not really used in WAN. Fast CPU together with WAN normally being slow makes it worth compressing the data. Recently vendors of accelerators have implemented both caching and compression in their devices. WAN accelerator can be either software integrated in clients and servers or equipment that is installed into the network.

4.1 Symmetric and Asymmetric Acceleration

Based on Sevcik and Wetzel [30], WAN Acceleration services are divided into two separate approaches called centralized or asymmetric acceleration and distributed or symmetric acceleration. Asymmetric acceleration uses a device in a datacenter which is designed to accelerate the traffic from and to the servers in the datacenter. Currently these asymmetric solutions can accelerate only Web-based applications. Symmetric solutions on the other hand accelerate a wide range of applications between point-to-point links between a datacenter and a branch office. In this thesis the emphasis is on symmetric acceleration solutions. Figure 8 below presents an example of a symmetric solution between a branch office and a datacenter. Accelerators are installed on both sides of the network just after the WAN edge routers.

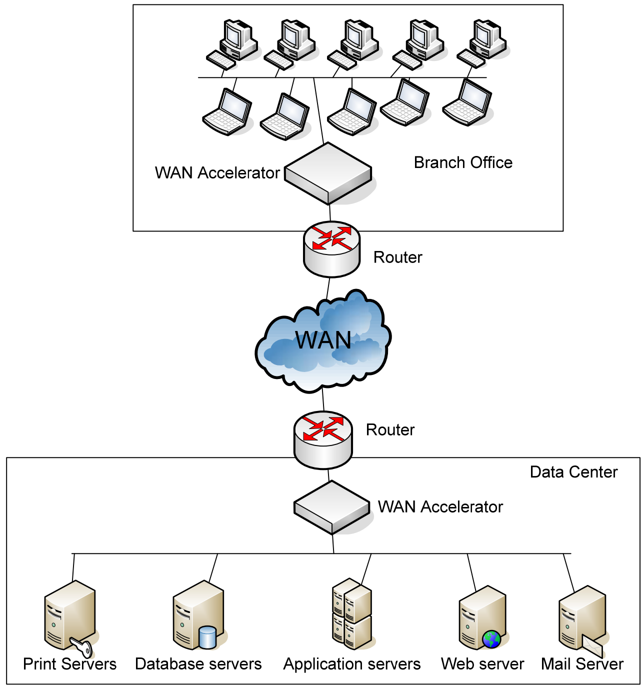

Figure 8: Symmetric acceleration – In symmetric acceleration an accelerator is installed on both sides of a point-to-point link (usually between a datacenter and a branch office) [30]

4.2 WAN Acceleration Techniques

WAN acceleration techniques try to minimize the effect of performance factors such as congestion, packet loss or latency. The goal of WAN acceleration is to keep the application performance in WAN close to what it is in LAN and address the problems that occur when switching from LAN to WAN. In this thesis the acceleration techniques are divided into two categories: WAN optimization and application acceleration. Most accelerators contain both WAN optimization and application acceleration techniques.

4.2.1 WAN Optimization

Fist type of WAN acceleration technique is WAN optimization which offers transport protocol optimization and compression. WAN optimization is further separated into sub techniques: TCP optimization, data suppression and compression. All of these try to solve the problems caused by transport protocols and network utilization. These techniques provide solutions for the performance problems described in the chapter 2.

4.2.1.1 TCP Optimization

As explained in chapter 2, TCP might not work efficiently when used in the WAN. TCP optimization is often done by implementing a TCP proxy in an accelerator. Accelerator terminates the TCP connection between the client and the server and creates a new optimized TCP connection between itself and another accelerator in the network. With TCP proxy the connection between the two nodes is split into 3 sections, see figure 9. One advantage of a TCP proxy is that possible packet loss happening in WAN is handled by the accelerators, not the end nodes. Another advantage is that the accelerators handle locally the acknowledgements coming from the end nodes, enabling a higher throughput because RTT is reduced. [3]

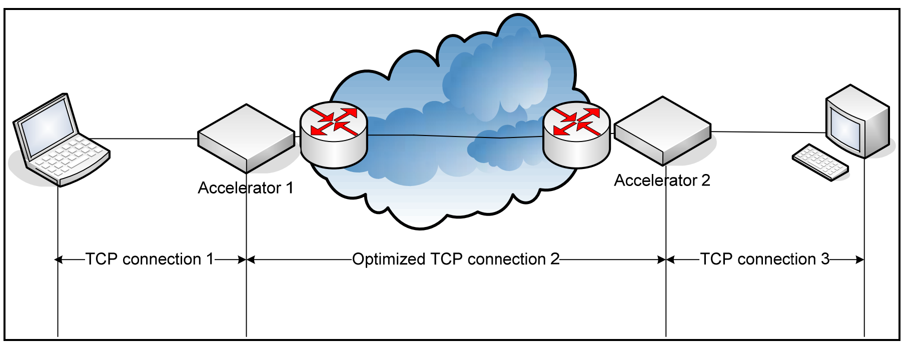

Figure 9: TCP proxy – In TCP proxy the real TCP connection between the client and the server is split by the accelerators into 3 separate TCP connections. The connection between the accelerators will be optimized. For the client and the server the connection looks as it was still the original connection.

Other TCP optimization possibilities include virtual window scaling, loss mitigation and advanced congestion avoidance. Virtual window scaling aims for better network utilization by increasing the window size. Loss mitigation tries to minimize the packet loss with more intelligent retransmission and error-correction. Advanced congestion avoidance tries to improve the transport protocol to have a better packet-loss recovery and bandwidth scalability. [3][30]

4.2.1.2 Data Suppression

The principle of data suppression, based on Grevers and Christner [3], is to reduce the amount of data by detecting data patterns sent over the network. Patterns of previously sent data are stored in accelerator in both the receiving and the sending end. These patterns get a unique token that is stored with the data pattern to the sending accelerators storage. Whenever the sending accelerator notices the same kind of data pattern coming from the sending node, it sends the token of that data instead of the data to the receiving accelerator. The receiving accelerator can then find the same data from its memory or disk storage with the token. Sending tokens instead of actual data gives significant bandwidth savings and reduces throughput. Data suppression is also known as dictionary compression [30] or byte caching [31].

4.2.1.3 Compression

Compression is a similar technique to data suppression. In compression data is modified with algorithms to reduce the size of the sent data. Popular algorithms are Lempel-Ziv and DEFLATE [3]. Compression is applied to raw data (within one window inside an object) before sending it across the network, whereas data suppression only replaces familiar data patterns with tokens. Data suppression only reduces the bandwidth usage with redundant data. Thus every new pattern of data that is sent is not reduced in any way by data suppression. Compression can reduce already the first set of data and therefore it is often used together with data suppression. In addition, the size of tokens can be reduced with compression. Compression and data suppression help to overcome the problems caused by oversubscription. Changing the amount of data sent over WAN minimizes the effects of bandwidth discrepancy between LAN and WAN. Oversubscription is discussed in chapter 2.2.1. [3] [17]

4.2.2 Application Acceleration

Application acceleration tries to improve the application layer operation in WAN. Application acceleration can be further categorized into object caching, read-ahead, write-behind, message prediction and Wide Area File Services (WAFS). These techniques are discussed in this chapter. Also the term Content Delivery Network (CDN) is defined in this section. In chapter 3 discussed the performance problems that the common applications used in corporate networks have. This chapter presents solutions for improving the performance of these application protocols. Because different application protocols work in different ways, it is not necessarily possible to apply all these techniques to every protocol. [3]

4.2.2.1 Object Caching

In object caching accelerators work as file caches for objects. This means that the accelerator at the other end of the network works as an origin server and the accelerator on the other end as a storage for local copies of objects that have been requested from the server. If an object is requested and the user requesting the object has a right to use it, a check is made that the object in the cache is identical to the one in the origin server and the local copy from the cache is sent to the user. [3]

Object caching is a similar method to data suppression with the difference that the object (e.g. a file) is not moved in any form across the network. Data suppression has to move at least the tokens, even if there are no changes in the data. In object caching the data is handled as whole objects where in data suppression the objects are handled in byte level. Object caching and data suppression work well together. Opening the object is done with object caching from the local cache and whenever the object is changed, it is sent back to the origin server using data suppression. Hence the changed bytes need to be sent instead of the whole file, because the bytes that have not been sent can be provided to the origin server as tokens. [3]

4.2.2.2 Read-Ahead

Many accelerators use application specific read-ahead algorithms for accelerating the requested objects. With read-ahead operation accelerator can investigate requests coming from an application. When an application requests a segment of an object from the origin server the read-ahead in the accelerator tries to predict the next possible application requests that the user might have and sends the predicted segments as well to the accelerator at the user end. If the prediction was correct, the next segments that the user requests are already at that end of the network. Application protocols that can benefit most from read-ahead are the protocols which send files, even large ones, in multiple small blocks of data. Figure 10 below shows an example of how read ahead functions with CIFS. Client requests first sequence of data from the server. Instead of just requesting the firs the sequence of data the accelerator will request 3 more sequences. [3]

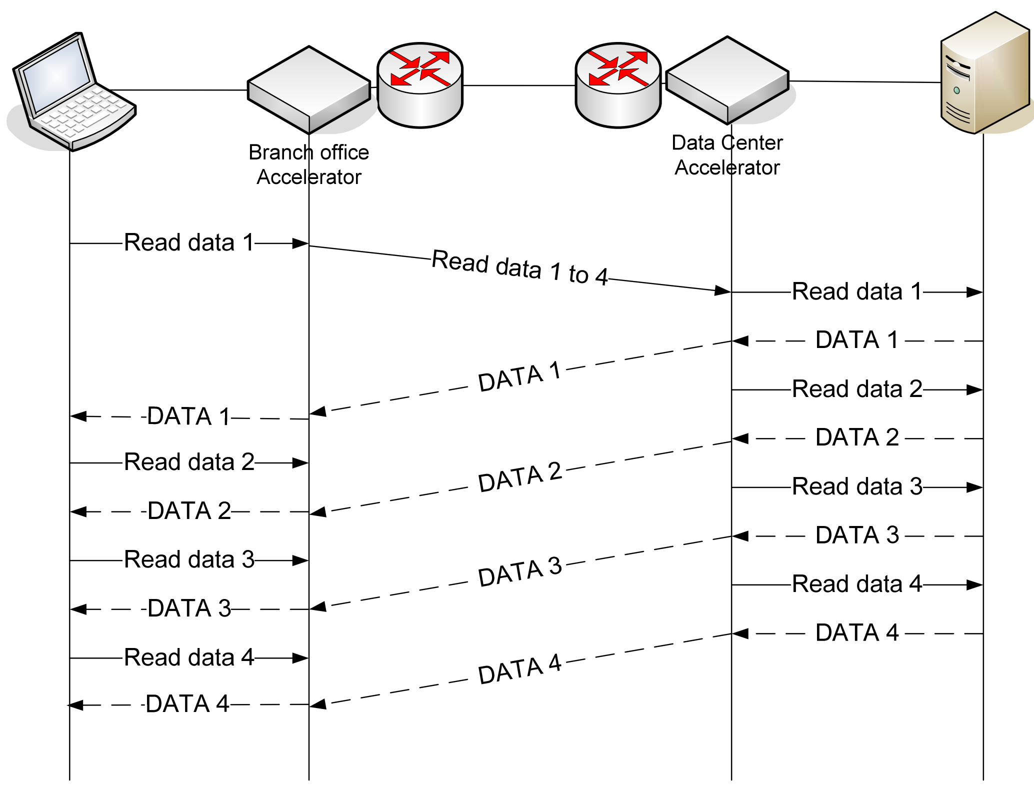

Figure 10: Read - Ahead for CIFS. When a client request comes to the branch office accelerator to read data 1, the accelerator will request also data 3 to 4 from the server. When client requests the

remaining 3 data blocks, these are already at the branch office accelerator and can be served to the client quickly.

Read-Ahead is an acceleration method that is beneficial especially, when used with the object cache. Read-ahead works well together with object caching since read-ahead can quickly deliver the objects that are requested for the first time to the accelerators cache. Using read-ahead without object caching can slow down the origin server from where the objects are requested from. Without object cache, every time an object is requested the request has to be done over the network and every request contains a set of predicting requests. This adds the load in the origin server. [3]

4.2.2.3 Write-Behind

With write-behind an accelerator can locally handle write requests coming from the client. Accelerator can in this way make the client believe that the data has been received. The accelerator sends the data to the origin server through the network but at that point the client believes that it has already been done. This reduces latency because the client gets a notification about receiving the data faster than it normally would. It is important for the accelerator to make sure that all the client data is also written to the origin server. Usually accelerators use write-behind when it is certain that it is safe and when it is possible to recover from a problem like loss of network connectivity. Many application protocols like Common Internet File System (CIFS) have built in write-behind. [3]

4.2.2.4 Message Prediction

Message prediction can be done either in a static or in a dynamic way. In both prediction strategies the accelerator sends predicted messages on behalf of a user or server. The goal is to reduce latency of client requests by predicting the actions client will do next and performing the required tasks before they are actually requested.

In static message prediction, an assumption is made that in certain situations application specific messages are always the same and in the same order. The rules how to handle a certain application protocol are saved in accelerator's logic. After a user request, accelerator can programmatically decide a certain amount of message operations on behalf of the user.

In dynamic message prediction the accelerator tries to learn from the messages that a user or a server is sending. The accelerator stores the user messages in a message history and calculates probabilities on what messages the user might send out next in different situations. In other words, the accelerator observes the previous behavior of the client and tries to operate accordingly. Figure 11 presents a probability table for an example client action. [3]

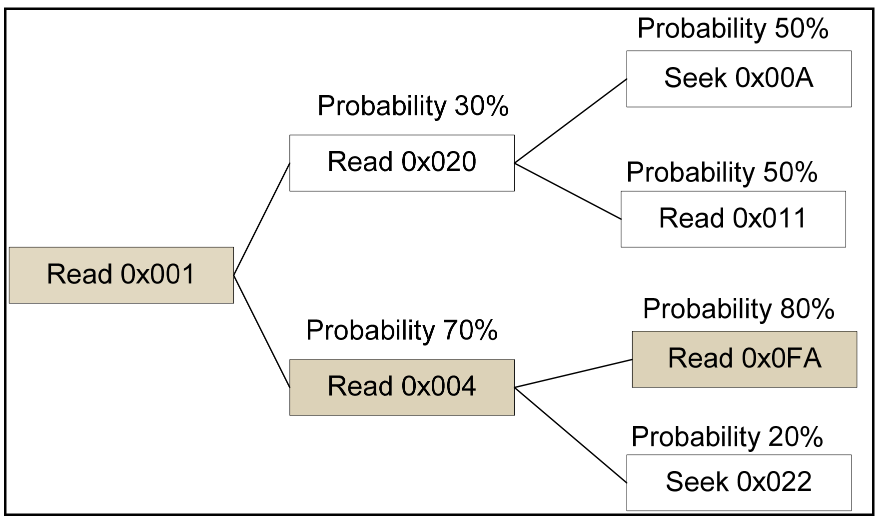

Figure 11: Message Prediction probabilities. For a read operation message prediction would select the most probable set of data reads. After 0x001 is requested by the client, the 0x004 and 0xFA would be requested next.

4.2.2.5 Wide Area File Services

Many accelerators, where caching is implemented, can also provide the servers services to the users on a disconnect mode when it is not possible to connect to the server through WAN. For a file server share, this type of services is called Wide Area File Service (WAFS). In disconnect mode the accelerator works as a proxy and offers access for the server files if the user has the right to use them. WAFS is also often referred to a broader set of services including for example printing or authentication. [3] [29]

4.2.2.6 Content Delivery Network

In the context of WAN acceleration, a content delivery network (CDN) is a set of accelerators that are managed centrally within a network. It focuses on the distribution of predefined set of objects. Usually these objects are some large files that are used by many people. WAN acceleration techniques like object caching can be used together with content delivery networks. An estimation of what objects the users might need is done and these objects are sent before hand to the accelerator working as a cache. This is beneficial with large objects like CAD/CAM files, medical images and for example with software distribution. [3]

4.4 Quality of Service

Normally IP networks use best effort delivery for handling the IP flows. Best effort delivery means that all traffic is treated equal and the flows are handled in the order they enter the network device. This might cause problems from the corporate network perspective. Treating the business critical applications with the same priority as Internet traffic for example, could have a big effect on the performance of the business critical applications.

Integrated services (IntServ) and differentiated services (DiffServ) are considered as the 2 ways to implement quality of service in corporate networks. IntServ is older and it is also more complex to implement than DiffServ. Thus, most of the corporate networks use DiffServ as the QoS method. Both use Classes of Service (CoS) to set different kind of priorities for the traffic. In DiffServ the QoS is done by assigning application traffic to different CoS queues in the network based on priority. For example a corporate network could have different CoS for voice traffic, business critical applications, email and Internet browsing. Even if Internet, set to the lowest CoS, would be used more than other applications it would not consume all the bandwidth as some of the bandwidth is reserved for other CoS levels. [32]

Quality of Service provides methods with which it is possible to define different service levels for different types of traffic. For corporate IP networks the quality of service became possible after the invention of MPLS IP VPN networks. Before this ATM networks already had QoS in them. QoS is a part of the functions improving WAN performance and many WAN accelerator appliances offer it as one of their services. Although many corporations already have implemented the router based QoS for small companies the QoS offered by the accelerators might be a desirable feature. [3] [32] [33]

4.5 Monitoring Application Performance

Network is an environment with millions of variables and monitoring tools are a way of getting an idea of how much and what kind of traffic goes through the network. Monitoring is a useful tool for detecting performance problems in the network. Monitoring tools are set up to monitor routers or switches in the network. Network devices like routers and switches support the Simple Network Management Protocol (SNMP) that enables the monitoring. NetFlow and Network Based Application Recognition (NBAR) are examples of mechanisms that are used in monitoring systems [3]. Many WAN accelerators have also implemented network monitoring in them. In this thesis a NetFlow based monitoring tool is used for evaluating the traffic in the case study in chapter 6. Therefore, the following sections discuss NetFlow based monitoring in more detail.

The term flow describes packets travelling in the network between two nodes. A flow has always a source and a destination, which are defined by OSI network layer IP-addresses and OSI transport layer source and destination port numbers. NetFlow is a service developed by Cisco which allows access to view the flow information of the network. Latest version is called NetFlow v9 and the IP Flow Information Export (IPFIX) starndard (RFC3917) is based on it. With NetFlow it is possible to view the network utilization, get detailed information on the application using the network and to see how many flows exist between the network nodes. In corporate networks NetFlow helps to determine what kind of prioritization of traffic would be most beneficial. [3][34]

NetFlow enabled routers or switches have a cache where they store all IP flows that pass through the device. These IP flows contain attributes like source IP address, destination IP address, IP protocol number and type (e.g. TCP, UDP), source port, destination port, Type of Service (ToS) identifier. If packets contain same attributes these packets are grouped to show the overall (identical) traffic. On top of these attributes, information like timestamp, subnet mask or TCP flags can be collected. [3]

The collected data can be viewed with different kind of tools either in real time or later from a server, also known as a collector, which stores the traffic history for a longer period of time. The IP flows are sent to the collector as UPD packets where one packet usually contains 30 to 50 flows. Collecting the traffic history is helpful for getting a long term view of the operation and usage of the network. Normally the traffic history is used, for example, to create reports about the top talkers (systems using the network) or top applications in the network. The real time view of the traffic can on the other hand help in detection problems in the network, for example if something is overloading a link the cause could be detected by viewing the traffic. [3]

5 WAN ACCELERATOR DEVICES

This chapter studies the WAN acceleration devices currently on market. Chapter introduces an accelerator from Blue Coat Systems. Section 6.3 discusses the installation and implementation of accelerators into the network and section 6.4 gives examples of previous experiences from using WAN accelerators.

5.1 Selecting an Accelerator

Before selecting an accelerator for a network, it is important to investigate different WAN accelerators in the market. Rolfe and Skorupa state in [29] that there are over 10 WAN accelerator vendors currently in the market. A study [7] with 3 different accelerator manufactures, Blue Coat Systems, Riverbed and Juniper Networks, has been done to help the selection of an accelerator solution for a target company. The target company is presented in chapter 6.1. The comparison of the accelerators was done based on 3 criteria: *scalability, transparency* and *device characteristics*. Scalability evaluates how well the acceleration solution suits certain kind of network, how well it works for example in a large network. Transparency evaluates the way the accelerators are implemented into the network (see chapter 6.3). With device characteristics the aim is to evaluate the WAN optimization and application acceleration capabilities of the accelerators based on the information provided in chapter 4. The results of the comparison were collected by answering to detailed questions under each criterion. The complete results can be seen in appendix 1.

The comparison in [7] showed that there are many similarities between different accelerators, especially when comparing the device characteristics and what kind of protocols the accelerators are able to accelerate. Biggest difference in device characteristics was that Blue Coat is the only one of the 3 solutions using object caching. Some more differences were found in the other two criteria: scalability and transparency. Riverbed for example was said to work well in large implementations and Juniper was the only one of the accelerators that was not offering a software version of the accelerator for remote users.

Testing accelerators in the target network is also an important part of the acceleration selection. It is important especially if the target company is aiming to accelerate only few selected applications. From the 3 compared accelerators Blue Coat Systems was selected to be implemented in the target company network for testing purposes. Chapter 6 is a case study of implementing Blue Coat Systems accelerators into a corporate network. Blue Coat accelerators are covered in the chapter 5.2.

5.2 Blue Coat Systems

Blue Coat Systems offers a group of technologies called MACH5 (Multiprotocol Acceleration Caching Hierarchy) that includes bandwidth management, object caching, data suppression (byte caching), protocol optimization and compression as a WAN acceleration solution. Blue Coat applications have a proxy based approach and MACH5 is provided as software to be installed to the Blue Coat proxy (Blue Coat SG) devices. On their Web-site, Blue Coat mentions file services, email, Web applications, video and audio and secure Web applications (SLL encrypted) as their key acceleration targets. Blue Coat is said to support the fully transparent installation mode (see chapter 5.4). The products range from small branch office devices into big accelerators that can support a data center. Blue Coat devices can also be installed as a software version to remote users' laptops. [27]

5.2.1 Data suppression

Data suppression is one of the techniques in Blue Coat devices that can be applied to many different application protocols, e.g. HTTP, CIFS and MAPI. Blue Coat uses the term *byte caching* for data suppression. Chapter 4.1.3.2 discusses data suppression. The Blue Coat accelerators are able to recognize different types of traffic and, based on the traffic, apply the most appropriate acceleration techniques to each application. Applying wrong type of acceleration to an application can cause performance problems. It is not beneficial for instance to use byte caching with streaming media data which does not contain any repetitive data patterns. Using data suppression with streaming media can cause the cache to fill up fast with data that can almost never be reutilized. [31]

5.2.2 Compression

Blue Coat uses a gzip/deflate algorithm for data compression. HTTP compression used by the Blue Coat devices is defined in HTTP 1.1 specification. The client defines whether it wants to use compression or not, and if compression is used, the client also tells which algorithm to use. The Blue Coat device then forwards the information about the used algorithm (gzip or deflate) to the server. After this, the server can send the client compressed data whenever possible. The data sent by the server is also stored in a cache in the accelerator in order to minimize the bandwidth usage even further. Blue Coat accelerators can also compress other data when the accelerators are installed to the network to as a content delivery network (see chapter 4.2.2.6). [35]

5.2.3 Acceleration for MAPI

Blue Coat MAPI acceleration is called *batching*. As mentioned in chapter 3 MAPI is a chatty protocol that uses many RTTs to deliver messages. MAPI also has only one block of data in flight at any time. Client is waiting for an acknowledgement of the previous data block before sending the next block to the network. The basic principle of batching is to acknowledge all data blocks sent by the client already at the client side accelerator. When all data has been received by the accelerator, it sends the data as one big block of data. This minimizes the RTT in the data transmission. Write-behind (chapter 4) is the general acceleration technique behind Blue Coat batching. [36]

In addition, Blue Coat also provides specific acceleration for the Microsoft Outlook email client. This technique is targeted for small branch office users behind slow WAN links who experience performance problems in receiving email when logging in to the Outlook in the morning. This technique is called keep-alive. In keep-alive the server and client side accelerators keep the connection for delivering email open even if the users log out of their email. The arriving emails are stored into the client side accelerator cache. When a user logs in next time, new emails are delivered from the client accelerator instead of the email server on the other side of the WAN. From different accelerator techniques explained in the chapter 4 the closest general technique that fits the principles of keep-alive is the WAFS service which makes it possible to use file services when the connection to the server in the datacenter is lost. The basic idea of the keep-alive is the same as in WAFS but it works to the opposite direction. [36]

5.2.4 Acceleration for SSL

Blue Coat offers a SSL proxy for WAN optimization. The SSL proxy is used to encrypt web applications with SSL. The applications can either be internally (server in side corporate network) or externally (Internet) hosted. The end-to-end SSL connection is cut and separate SSL tunnels are created between the client and the accelerator device. This is done without breaking the secure mode of SSL. The accelerators have the client and server public keys in them. The principle is similar as for the Juniper SSL acceleration (see chapter 6.4.1.6). [27]

5.2.5 Acceleration Results

Different accelerator companies have published test results of expected acceleration results with their own products. Tests are done in a test environment, and therefore the results cannot be taken literary. However they give some idea about the accelerator performance. Figures 12, 13 and 14 give reference values for Blue Coat acceleration with CIFS. The figures show three different measures for operations with CIFS object (MS Word document). Measurement (1) is done without Blue Coat device, (2) as a cold test -the file passing through the accelerator for the first time and (3) as a warm test - the file is already in the cache. Figure 12 measures the bandwidth gain that can be reached when using acceleration for opening a 2,1MB file. Figures 13 and 14 measure the response time savings when a 2,1MB file is opened/edited/saved. The results also vary based on the link size. It is possible to see bigger improvement in response time for smaller throughput or higher latency can links (compare the results of figures 13 and 14).

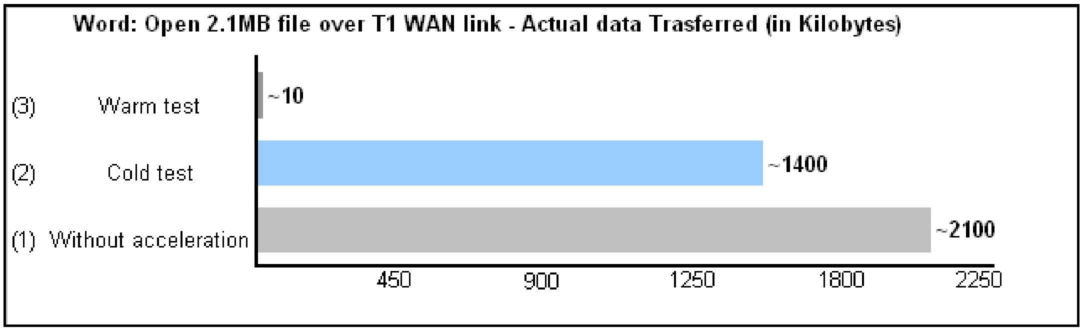

Figure 12: Blue Coat acceleration for CIFS – Actual data transferred (Kbps) (I,5Mbps WAN link, opening a 2,1MB file) [37]

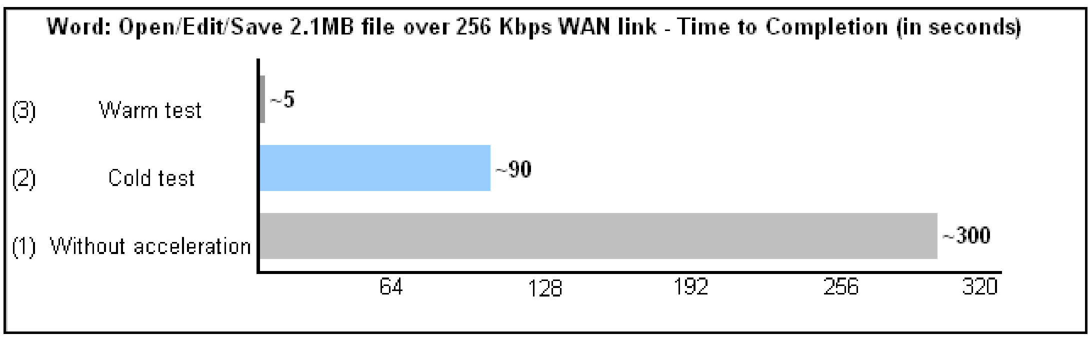

Figure 13: Blue Coat acceleration for CIFS - Time to completion (s) (256 Kbps WAN link, open/edit/save of a 2,1MB file) [37]

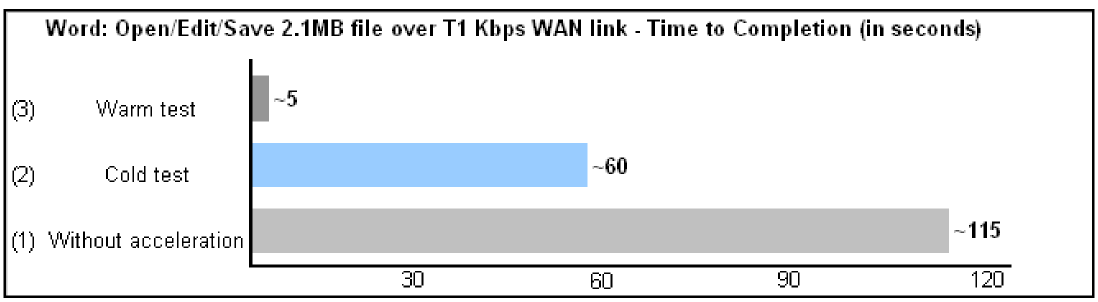

Figure 14: Blue Coat acceleration for CIFS - Time to completion (s) (1,5Mbps WAN link WAN link, open/edit/save of a 2,1MB file) [37]

5.3 Previous Experiences with WAN Accelerators

WAN acceleration devices have raised interest among non-commercial organizations with complex network infrastructures or with homegrown applications. For example, The US Federal Aviation Administration (FAA) announced in [38] a plan for testing at least 3 different commercial WAN accelerators in their network. The aim was basically the same as with any other corporation, savings in the bandwidth costs and improvement in application performance. These tests were not yet conducted but the results can be interesting as the tests were planned for "homegrown" non-commercial applications. Another example of application acceleration is presented in [39] where the US Military Sealift Command (MSC) performed acceleration tests in highly demanding WAN conditions over satellite connections. As the result of the tests MSC is planning to extend the implementation into wider scope.

Other interesting test results were executed by Network World in 2007 [33]. Four different WAN accelerators, Blue Coat being one of them, were compared in a corporate network environment of four sites and the results were stored from a seven month time frame. In this thesis these result are used as a reference when estimating the results gained in the target corporation network. These results give also some base for comparing the gained results to other accelerators, keeping in mind that the results are already a year old and that the devices might have evolved from what they were then.

5.4 Installation of WAN Accelerators

In this chapter, an overview of possible ways to install accelerators into the network is given. A general way to refer to the location of an accelerator in the network is to say that the accelerator is installed between the WAN router and the LAN switch. This chapter explains what this means and what options there are for placing the accelerators in the network. To accelerate the traffic between two network nodes, an accelerator needs to be installed on both ends of the connection. Another thing to be noted is that normally only a part of the traffic is accelerated, either due to the fact that the accelerator is not able to accelerate all traffic or simply because it is not necessary or even desirable to accelerate all traffic. Corporations might not want to accelerate Internet traffic but to concentrate only on the business critical applications for example. [3]

Chapter 5.4.1 explains the 3 different transparency levels that accelerators can use when installed into the network. The installation of an accelerator can be divided into a physical and to a logical installation. These are covered in chapter 5.4.2 and 5.4.3.

5.4.1 Transparency

Transparency measures the way the accelerator handles the visibility of the real source and destination addresses between two accelerator devices. Some accelerators change the destination and source IP and port when directing the traffic to the WAN. Grevers and Christner state in [3] that this can be a problem in networks where packet header information is used, for example, for security, monitoring or packets classification. For a monitoring tool using NetFlow to create statistics this means that the flows are seen exchanged between the accelerators and not between the destination and the source. This can basically make the monitoring tool data useless since all the accelerated data appears similar to each other. [3] [40] [41]

A scenario with hidden and with non hidden source and destination addresses can be seen in the figures 15 and 16 below. When the traffic enters the accelerator the source and destination IP and port are changed, for instance the real source IP 10.10.10.10 would be changed into the accelerator device IP 30.30.30.30 and the source port from 80 to 5000 (see figure 15). As monitoring tools use the IP and port to identify flows all traffic would appear to come from the same source (30.30.30.30 the IP of the accelerator) and going into the same destination (40.40.40.40 the IP of the second accelerator). In figure 16 the accelerators do not change the source and the destination IP and port and the monitoring tools are able to identify the real traffic. [3] [41]

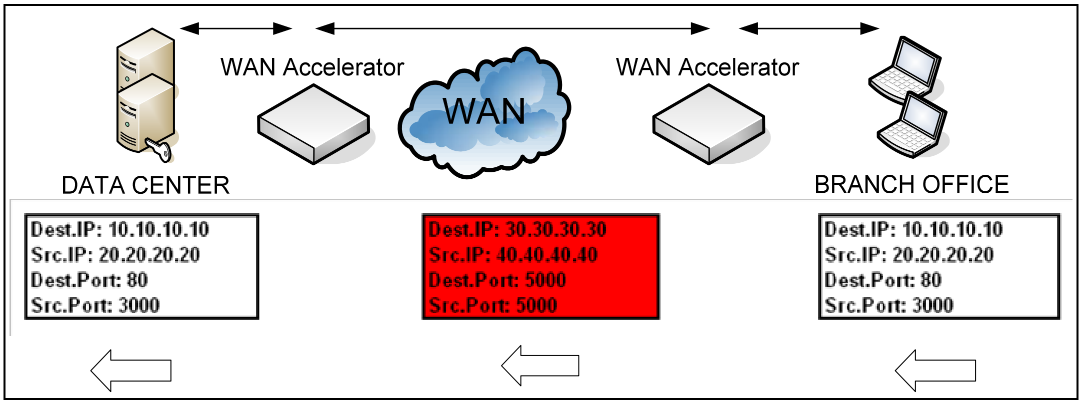

Figure 15: Nontransparent accelerators

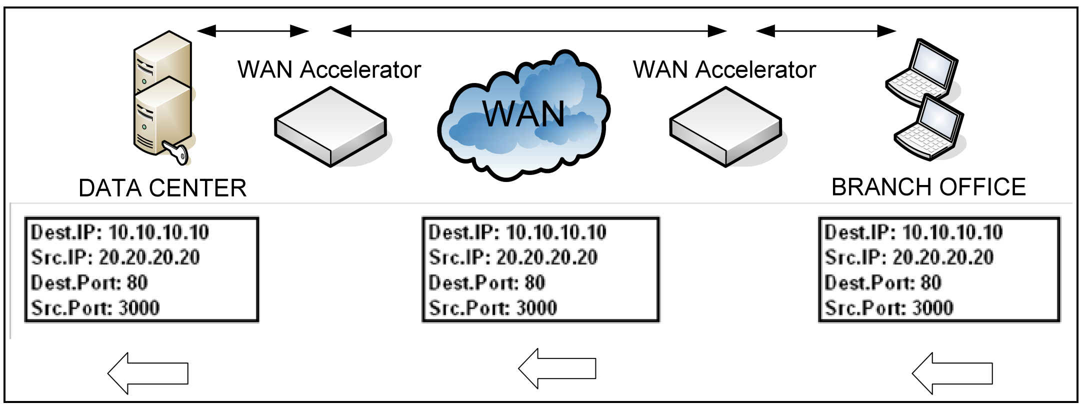

Figure 16: Full transparent accelerators (IP and Port)

5.4.2 Physical Installation

Usually the physical installation is done by installing the accelerator to an already existing network device, normally to a router. If the implementation is done in this way the router needs to support a technique which is able to route the traffic into the accelerator. Web Cache Control Protocol (WCCP) is an example of this kind of technique and it is discussed in chapter 5.4.3. [3]

Another, less common way to install an accelerator is to install it directly between the WAN router and the LAN switch, using a network card. The network card has to support fail-to-wire operation. The fail-to-wire support in the network card means that in case of a hardware failure or some other error the network card physically cuts the connection between two network cables. Installing the accelerator directly to the network path is called physical inline installation. This way the accelerator is able to see all traffic between the router and the switch. Inline installation is the easiest way to implement an accelerator to the network, but it has some downsides compared to the above mentioned other option. It is possible that in case of a failure in a device, not recognized by the network card, the card is not immediately be able to react to the failure. This can cause all network traffic to stop going though the device. [3]

5.4.3 Logical Installation

Logical installation, also known as network interception, is needed whenever an accelerator is physically installed together with the router. As stated in the previous chapter, the traffic needs to be somehow routed from the router to the accelerator. This can be done either with WCCP version 2, Policy Based Routing (PBR) or with load-balancing devices like a server load balancer (SLB). The basic principle is to use some protocol or a device, which takes care of the routing of the traffic to the accelerator. In case of PBR this means that the network administrator is able to insert new, other than destination IP based, routing policies into the router. SLB is an option for directing the traffic to the accelerator in a datacenter especially if multiple accelerators have been implemented. Load balancing means that the traffic flows are divided evenly between the accelerators based on their current load. In [3], WCCPv2 is mentioned as the most recommended way for the logical installation of an accelerator and WCCPv2 is discussed in the following chapter. [3]

5.4.4 Integration with WCCPv2

Web Cache Control Protocol version 2 is used for connecting routers and web-caches to each other. With WCCPv2 it is possible to direct traffic in the network in real time from one or more routers to one or more web-caches. An accelerator can connect to a router with WCCPv2 - the accelerator becomes a WCCP client and the router becomes a WCCP server. The client is then able to notify the network what kind of traffic can be directed to it. The WCCP server, instead of directing the traffic flows to the real destination, directs flows to the client. These WCCP servers (routers) include automatic load balancing. WCCP server keeps a list of all of the clients, called *the service group*, and divides traffic to the clients based on their load. One service group can contain 32 servers and 32 clients. If the WCCP server has an empty service group (no clients on it), the traffic is directed to its original destination. Adding new devices into a service group is dynamic and does not usually cause breakouts in the service. [3] [42]

6 WAN ACCELERATION CASE STUDY

This chapter concentrates on the implementation of accelerators into the target network for testing purposes. First the chapter presents the target company and its network (6.1) and then discusses the critical applications for the company (6.2). Section 6.3 discusses the acceleration tests and studies the results of the tests. The concentration on the tests will be on these critical applications.

6.1 Company and Network

The target company is a large-scale company in manufacturing and service business having globally offices in North-America, Asia, Europe and Africa. The variety of sites based on the amount of users is large, ranging from 600 user sites to small branch offices of 5 or less people. The standard set of offices for one country would include one or two big sites with 100-600 network users and then a large number of small branch offices. In addition to this there is a datacenter on almost every continent. As described in chapter 2, especially the branch offices are the ones affected by poor network performance due to long distances to data centers and low bandwidth connections. A typical branch office for this company would have around 10 users. The typical applications used would be standard office applications like email and web-based solutions in intranet and Internet, SAP and CAD.

The company has recently outsourced the majority of its IT operations, and at the same time switched to a more centralized IT/network infrastructure. One reason behind centralization (discussed in chapter 2) has come from the business: it is vital for the company to be as close to the customers as possible. This means a high number of small branch offices. As the first step for managing the centralization, the company invested on network monitoring tools that allow fault detection and analysis without having support on each site. As the next step, the company is considering WAN accelerators as a way of improving branch office performance. Another driving force for possible investment in the accelerators is the high cost of network bandwidth and the cost savings that accelerators might offer.

6.2 Critical Applications

This chapter introduces the selection of the applications and protocols that are considered as top priority for the company. The target company has selected CIFS, CAD/CAM, SAP, MAPI and HTTP as their primary acceleration targets. The following sections discuss these applications in more detail and explain what makes these applications the primary targets. Many acceleration companies have concentrated on developing acceleration for application protocols that suffer most from poor performance when used in WAN. Based on [33], the applications that currently benefit most from the application acceleration are the Microsoft based protocols like MAPI and CIFS. HTTP is also said to be one of the main acceleration benefiters. In addition in chapter 3 CIFS, MAPI and HTTP are also named as chatty applications that suffer most from the poor performance when used in WAN.

6.2.1 CIFS

CIFS is the protocol used for file sharing in the target network. Currently the amount of CIFS traffic in WAN is considerably low. As explained in chapter 6.1 the target company is planning to centralize file servers from the branch offices into the global datacenters. This will increase the amount of CIFS traffic in the WAN and makes also testing CIFS the number one priority for the company. Based on chapter 3 CIFS is also the worst application, when it comes to performance in WAN. Considering both of these aspects CIFS is considered as the most critical application.

6.2.2 CAD/CAM

In the target company CAD applications are mostly used by the research and development (R&D) sites. The number of these R&D sites is low but they are located far away from each other this makes the distance and latency between the sites high. These

R&D sites need to quite often use resources like AutoCAD files from the servers located on other R&D sites. This makes the latency between these sites a critical point. In addition, these AutoCAD files are usually large. Combining high latency environment with large data objects that need to be moved often across the network, makes CAD files a desirable target for acceleration.

6.2.3 SAP

SAP is one of the most important applications within corporations. This is also the case with the target company. SAP servers have already been centralized to the datacenters. The target company has SAP users in almost every site, also in the small branch offices. As discussed in chapter 3, performance of SAP is generally good even when used through WAN but problems caused by congestion or latency can slow down SAP performance. Based on the study made in [7] the different acceleration vendors are not promising high performance improvements for SAP but they say that it is possible to accelerate it in some level. Blue Coat, for example, states that they are able to improve the log in times for SAP. As a conclusion SAP is important for the target company, but the possibilities to improve the SAP performance with the acceleration are not as high as with some other applications.

6.2.4 MAPI

The target company uses Microsoft Outlook as their email software. MAPI is the protocol used by Microsoft Exchange server and Outlook client. These Microsoft Exchange servers are already centralized into global datacenters. This centralization of email servers has caused some problems in the network. Small branch offices are affected, especially during mornings, by slow email (MAPI) response as the client starts to download the new emails from the server. Another viewpoint for selecting MAPI as one of the critical protocols is that email is one of the most used applications within corporate networks. Reducing the amount of data the MAPI traffic is desirable. On the other hand the sending or receiving emails is not that time critical as actions performed by some other applications. Based on these points MAPI is set to medium priority as an acceleration target.

6.2.5 HTTP(S)

In [29] Rolfe and Skorupa state that corporations are using more and more Web-based applications for example for Customer Relationship Management (CRM) which has resulted in a need for a good HTTPS acceleration. This is also the case for the target company. The target company is using a Web-based application for CRM. Microsoft SharePoint collaboration tool is another example of a tool using HTTP. The target company sees Microsoft SharePoint as a possible replacement for normal file servers. In the future this kind of applications can come even more popular which can increase the use of HTTP(S) for not only browsing the Internet but also for business purposes. This makes HTTP one of the critical application protocols for the business.

6.3 Testing an Acceleration Solution

Before implementing any larger accelerator solution into the network, it is good to perform tests with the devices to see how well they perform in the network. In the target company accelerator provider was selected based on investigation of the capabilities of different accelerators and also based on the recommendation of one of the company's partners [7]. Blue Coat accelerators were selected as the solution to test in the network. The applications selected for testing are discussed in chapter 6.2. From these applications this thesis concentrates on tests with CIFS, MAPI and HTTP. Because the possibility to see good acceleration results with these applications is high, these three were selected. Testing the other 2 critical applications SAP and CAD is also recommendable but these tests are not covered in this thesis.

6.3.1 Test Setup

The test environment was selected inside one of the biggest countries in the company. The test country has many small branch offices and it is also using all of the most important business critical applications. The test devices were implemented between a local country head office working also as a local data center and a global data center within another country. Figure 17 shows the test implementation. The users in the local data center could benefit from the acceleration between the two sites. The users are using for example Internet, and email through the global datacenter. This makes HTTP and MAPI traffic candidates for acceleration. For more detailed information on what services the users are using from the global datacenter a traffic analysis needs to be made. This is discussed in chapter 6.3.3.

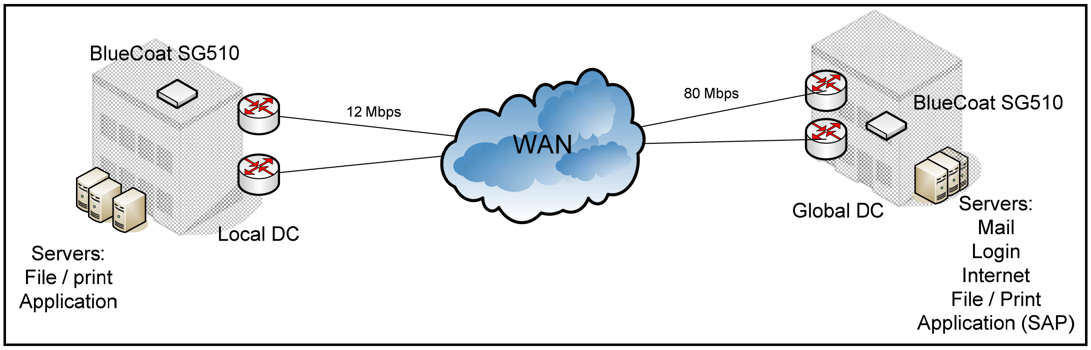

Figure 17: Test setup for the WAN accelerator installations. For example mail server is located in the Global datacenter and there fore MAPI traffic can benefit from acceleration.

6.3.2 Test Installations

For the installations the same model of Blue Coat accelerator (SG Proxy 510) was installed on both sites. The accelerators were implemented using WCCPv2 (see chapter 5) in a fully transparent mode. The fully transparent mode enables the use of NetFlow based monitoring tools for full monitoring of the accelerated data at the same time.

Blue Coat devices have a management console tool that allows web-based access for the administrators to the accelerators settings. Before starting the acceleration it is important to set up the correct acceleration for the network. With the tool it is possible to set up the ports and IPs that the device accelerates. Some of these are preconfigured for accelerating certain types of traffic like HTTP or CIFS. For the HTTP traffic, the device listens automatically to the port 80. It is possible to select, based on a known protocol which, traffic is intercepted (accelerated) and which traffic is bypassed (not accelerated). There is also a way to create priorities for accelerating different types of traffic. For example, HTTP can be set up as low priority traffic. During the main tests, the devices were set to only accelerate HTTP, MAPI and CIFS traffic, bypassing all the other traffic. After these main tests were made also some initial tests were performed on SAP and CAD, and the devices were set to also accelerate SAP and CAD.

6.3.3 Traffic Analysis in Test Network

In order to find out how much and what kind of traffic was transferred between the different connections, a WAN monitoring tool for investigating the traffic flows was used. Estimations of the traffic load and types between the test locations were made. Figure 15 shows an analysis of the traffic between the local datacenter and the global datacenter. The analysis is made based on the information from a NetFlow based monitoring tool. The estimations of the traffic were taken from both the incoming (WAN IN) and outgoing (WAN OUT) interfaces of the local datacenter WAN link. Incoming WAN traffic (the traffic coming from the local LAN and going to the WAN) is the traffic created by the users in the local datacenter e.g. email or Web-browsing. Based on the figures the most desirable targets for the acceleration would be HTTP (40% of the total traffic), MAPI (30% of the total traffic), logon (Netlogon) server traffic (12% of the total) and SAP (10% of the total traffic) traffic. Logon server is used for providing users the credentials they need for login in to the network. CIFS traffic is not shown in the figure as at the time of the tests the test site had still its own file servers.

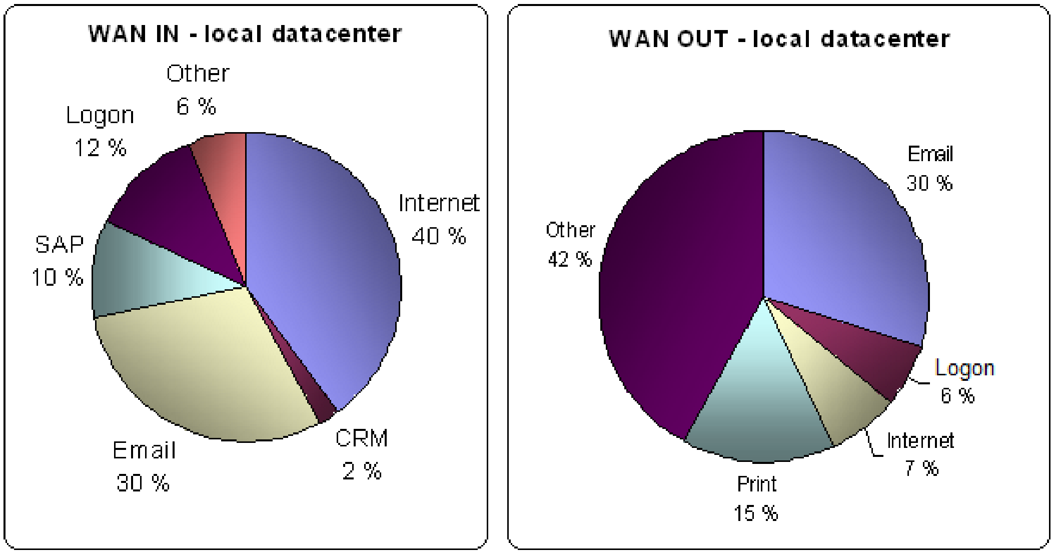

Figure 18: Local Datacenter to Global Datacenter - traffic analysis. The traffic analysis shows that the users in the local datacenter are using the WAN link mostly on web-traffic, email and getting the login credentials from the login server. SAP usage is also quite high.

6.3.4 Test Plan and Results

The tests were performed with three test users in the local DC using resources from the global DC. Tests were performed out of business hours to make sure any other traffic in the WAN link was minimal during that time. Both the measurement tool built in the accelerator and also the separate NetFlow measurement tool were used to view the changes in the bandwidth usage of the network during the tests. No other measurements like the response time were taken. The general test plan was to perform the same kind of steps with the acceleration and without. The test with the acceleration was done twice. This is known as performing a *cold test* and a *warm test* where the cold test means that the traffic is going through the accelerator for the first time and the warm test means that the traffic is going through the accelerator for the second time (the object or the bytes are already in the cache). The information helps to find out the benefits of acceleration techniques that reduce the amount of data sent through the network (e.g. data suppression or object caching). The application specific test steps are discussed for each tested protocol (HTTP, MAPI and CIFS).

6.3.4.1 Tests for HTTP

For HTTP the tests the plan was to use a website where one can download a large block of data (20,5MB) to test what kind of benefits the accelerator brings. As a second test, high definition picture (6MB) was retrieved from another website. For a third test a Microsoft Office file (24,5MB) was retrieved from a SharePoint server in the global datacenter. SharePoint is browser-based collaboration tool using HTTP for and file sharing in similar way as CIFS.

Internet traffic is generally known to use caching of its own and therefore it is important to always close the previous HTTP session before performing the test again in order to see how the caching provided by the accelerator is working. The results for the tests are presented in the figures 19, 20 and 21. For the first test with a download of 20,5 MB of data almost no improvement with the warm acceleration was seen. As can be seen from the figure 19 the amount of data transferred was reduced only with 1,5 MB when using acceleration (warm acceleration). Reason behind the poor performance improvement can be that the data could not be downloaded in a standard way for some reason. It can be that the accelerator did not recognize the traffic as HTTP traffic and therefore the benefits were not shown.

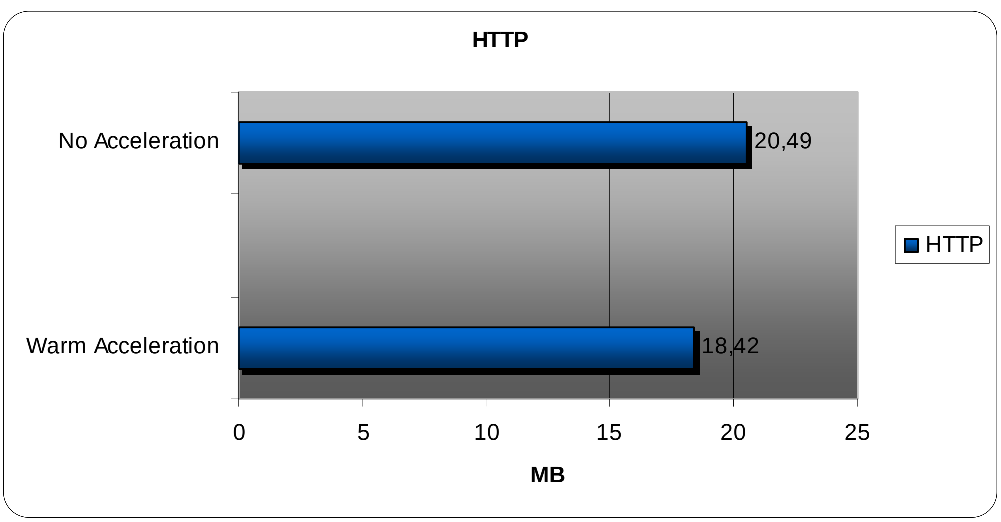

Figure 19: HTTP test 1 - download of 20,5 MB

As the first test was not as successful as expected, another website was used to perform more tests. In the second test, a 6 MB high definition photograph was downloaded. Figure 20 presents the results of the second test. In this test, the measurements were taken twice with the acceleration and once without the acceleration in order to see the difference between a cold and a warm test. The cold test should show the improvements with an empty cache and the warm test with the cache that already includes the data. For the HTTP test we did not see any improvement with the cold cache but a big improvement with the warm cache, which indicates that only object caching or data suppression was used. This means that the benefits of the acceleration appear only when an object is retrieved more than once from the Internet.

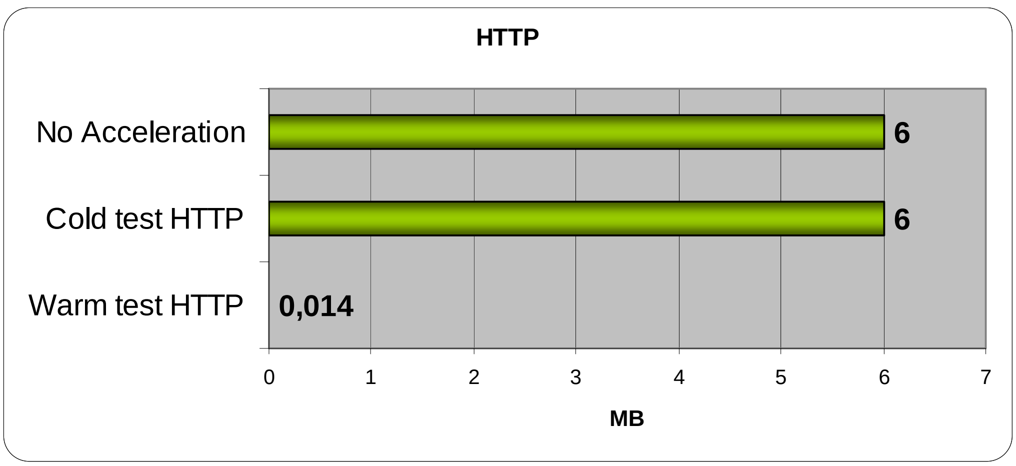

Figure 20: HTTP test 2 -- download of 6MB

In the third test a 24,5MB Microsoft office file was retrieved from a SharePoint server in the global database. The test was done once without acceleration and once with acceleration (cold test). A warm acceleration test for HTTP using SharePoint is done in later tests (see figure 23). This test was the most successful of the three tests as the acceleration showed reduction in the amount of transferred data already for the cold acceleration. Based on the HTTP test 2 (figure 20) and HTTP test 3 (figure 21) it seems that the type of the data affects the acceleration, when the data passed through the accelerator for the first time. With the photograph there was no benefit from byte caching or compression during the cold test but with the Microsoft file the accelerator was able to utilize these acceleration techniques during the cold test. Similar results can be seen in the following CIFS tests (figure 19 and 20).

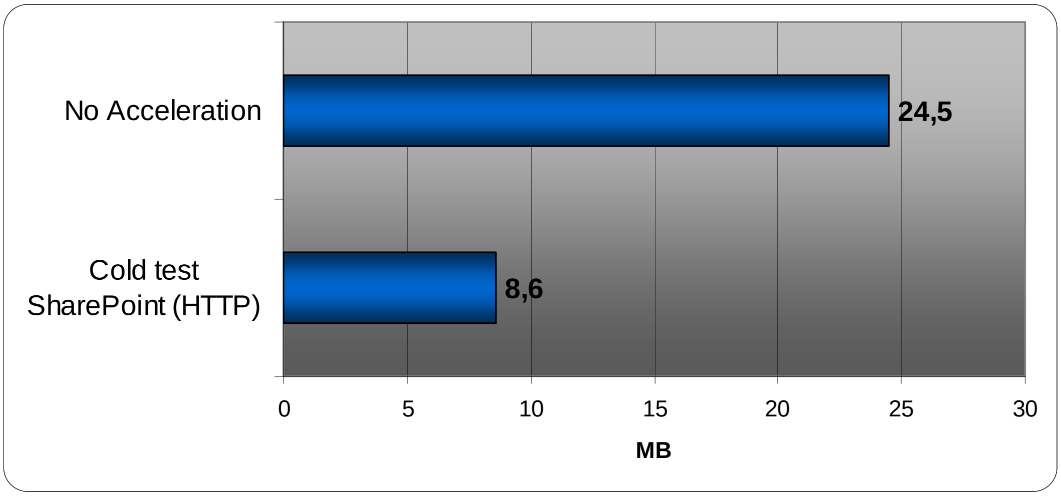

Figure 21: HTTP test 3 – download 24,5 MB from using SharePoint platform

6.3.4.2 Tests for CIFS

For CIFS, the test plan included opening a Microsoft Word document from a file server in the global datacenter. The file was retrieved twice from the file server. The results of the CIFS test can be seen in the figure 22. 40 MB file with acceleration created only 450 KB of traffic. In the first CIFS test, only the case with a cold cache was tested because the improvement seen in the measurement systems was already very significant. The big change in the amount of transferred data indicates similar results as with the last HTTP test. The accelerator is capable of performing byte caching or compression for Microsoft file already during the cold tests when the file is moved across the network for the first time.

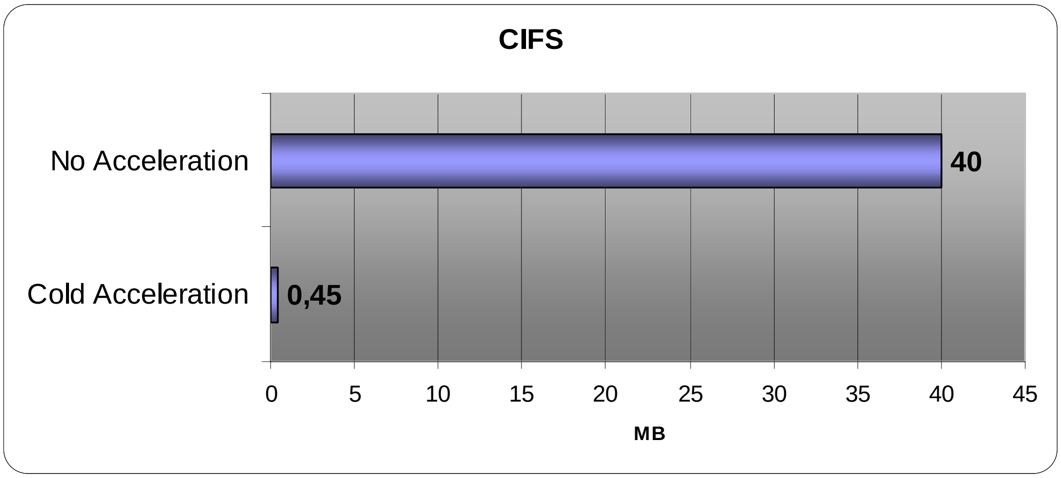

Figure 22: CIFS test 1 – retrieving a 40MB file from a server

To verify the good results from the first test, another tests was done using a 39,86 MB text document. In this test the same file was also retrieved from the SharePoint server with HTTP to enable the comparison between the CIFS and HTTP file transfer. When working with files the file name can be changed multiple times as the new changed file can be saved as a new version of the previous one. During the testing the file name was also changed in order to see, if that has any affect on the results. The results of the tests are presented in figure 23. During the test the following steps were performed: 1) retrieval without acceleration from a file server using CIFS, 2) retrieval without any changes to the file from the same server but acceleration activated (cold test), 3) retrieval with changing the file name from the same server (warm test) and 4) retrieval from the SharePoint server using HTTP, with a new file name.

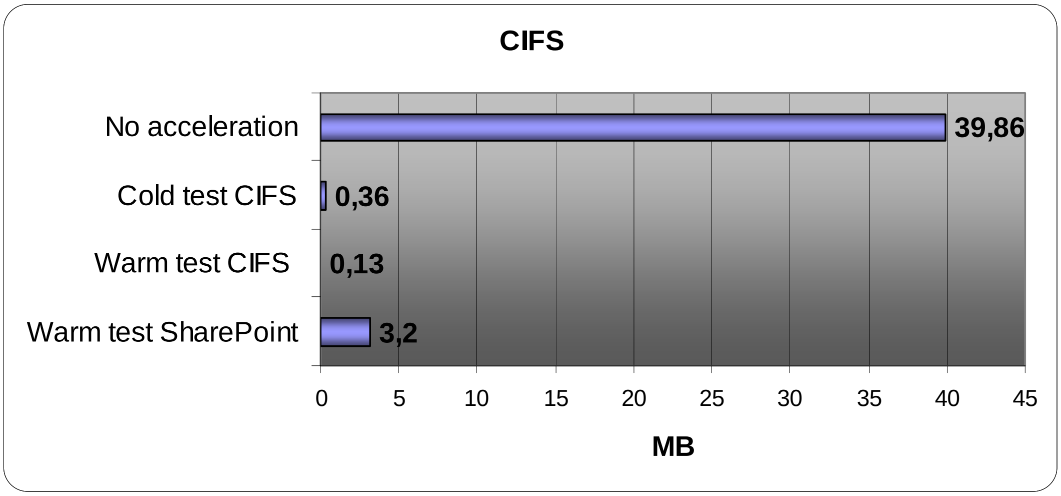

Figure 23: CIFS test 2 -- comparing CIFS performance to SharePoint (HTTP) performance

Comparing these results to the measurements presented in figure 12 (chapter 5), where similar tests were performed (by BlueCoat) with a 1,5Mbps WAN link with 2,1MB Microsoft file, the results presented in figure 22 and 23 seem better. Especially the reduction in the transferred data for the cold cache is good. One explanation for this can be that the files used in test 1 (figure 22) and test 2 (figure 23) contain a lot of repetitive patterns and that the acceleration can benefit from data suppression already during the cold cache phase, as the cache gets the repetitive patterns while transferring the file. The results (in figure 23) also prove that changes made in the file name have small or even no effect on the results of the acceleration, as the file name was changed multiple times during the testing. Final result of the tests indicates that CIFS file transfer can benefit more from the acceleration compared to the file transfer made with HTTP for the same file.

6.3.4.3 Tests for MAPI

The first tests with MAPI were performed by sending an 8,3MB email containing 3 photos. The email was sent out from the global datacenter to the user in the local datacenter. This means that the email is first sent to the exchange (email) server in the global data center without acceleration and then from the exchange server to the local datacenter user with the acceleration. Figure 24 shows the way the email was sent though the network. To test how MAPI acceleration works another email, partly similar to the first one (two photos changed into new ones), was also sent through the network.

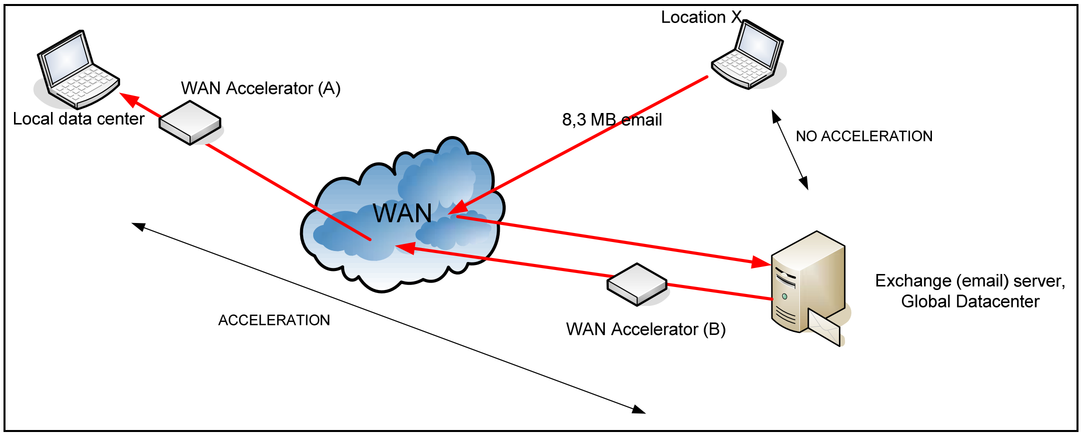

Figure 24: Test setup for MAPI – 8,3 MB email is sent from location X to the user in local datacenter. The email must travel the network first (without acceleration) from the location X to the exchange server in global datacenter and from there accelerated to the user in local datacenter.

The results of the first MAPI tests can be seen in figure 25. Four different measurements were conducted. The measurement at the bottom shows the size of the email with acceleration when the email is sent out for the first time. The size is the same as the size of the email without any acceleration (the top bar in the figure). Other expected benefits by TCP optimization for example, are not seen as any response time measurements were not taken. In the second measurement, the same email was sent again through the network. This is, as with HTTP, a warm test where the accelerator cache already contains the same data bytes. For the second test the measurements show that the data transferred through the network is 817 Bytes instead of 8,3M Bytes. The third measurement was done with an email containing one of the previously sent photos and two new photos. With changed data there is still improvement seen due to some of the bytes being already in the cache. For the third measurement, the network monitoring tools showed that two separate sessions were opened, one for the data already in the cache and one for the data going through the accelerator for the first time.

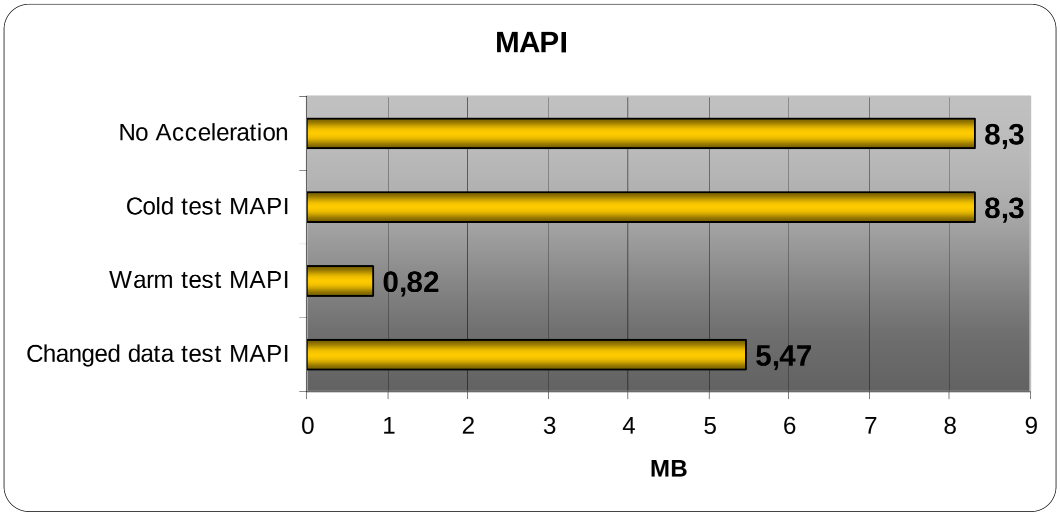

Figure 25: MAPI test 1 -- 8,3MB email

The second MAPI test with an email, size 5,97 MB, containing a Microsoft PowerPoint presentation as an attachment included three measurements. The results of these measurements can be seen in figure 26. In the first measurement an email was sent from an external site to person X in the local site without acceleration. In the second measurement the same email was sent to 2 other persons Y and Z in the same local site as a cold test with acceleration. In the third measurement the same email was sent again to the person X but no two times in a row. This third measurement represents the warm acceleration. In addition, the mailbox of user X is located in exchange server 1 and the mailboxes of users Y and Z are located in exchange server 2. As can be seen from the figure 26 the cold acceleration with Microsoft file gives better acceleration results as the cold acceleration with photographs in MAPI test 1 (figure 25). The difference between the cold and the warm acceleration is minimal. Another conclusion that can be made is that having the email sent to different servers has no impact on the acceleration results. Both of the MAPI tests indicate that the acceleration rate is not as good as for CIFS or even HTTP. It also seams that for each sent email MAPI changes the bytes in the email, since the acceleration rate for the warm test is not as high as with the other two tested protocols.

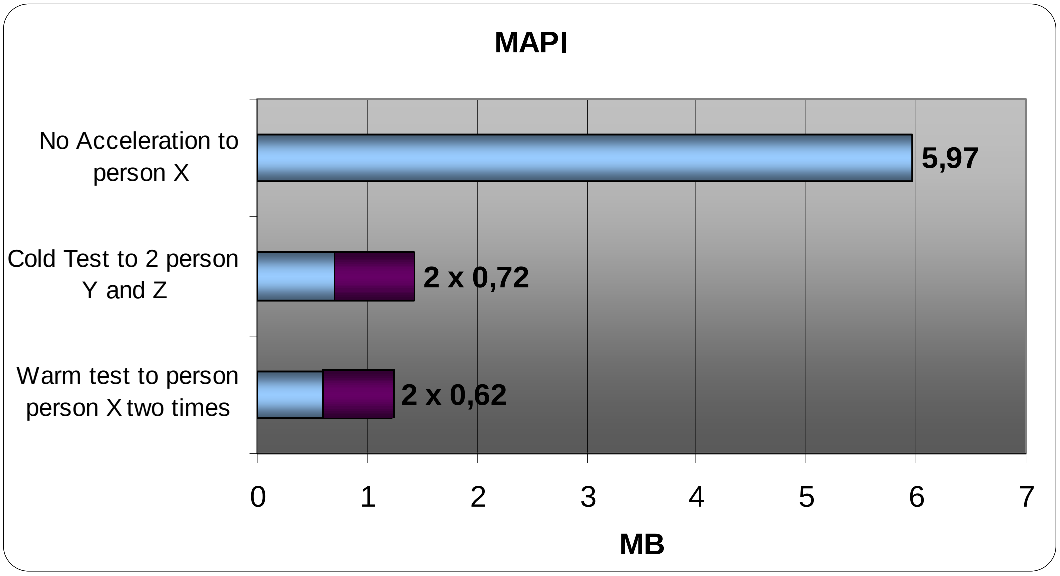

Figure 26: MAPI Test 2 -- 5,97 MB email

6.3.4.4 Problems Occurred

When enabling the acceleration it is possible to define the types of protocols or even the users that the acceleration is enabled for. During the tests it was decided that only the users participating in the tests would be accelerated. It was also decided to test the acceleration during longer periods of time after the individual tests for HTTP, CIFS and MAPI were performed. In this longer full test the acceleration would have been activated for all the users in the local site towards the global datacenter for HTTP, CIFS, MAPI, SAP and CAD. All other traffic would have been used normally without acceleration. The "full" activation of the acceleration proved to be difficult as the local site users started to experience problems with some applications using HTTP. For this reason the full acceleration had to be stopped. For the investigation of the root cause of the problem assistance from the accelerator vendor was needed. After some analysis the problem was found from the way the WCCP was implemented on the router at the datacenter. The access list of the WCCP had a typo in one of the IP ranges of the local site, and this caused the problems. As the information has to be typed by hand into the access list it is likely to suffer from human errors. This possibility for serious errors should be noted when installing accelerators by using WCCP.

6.3.4.5 Test Summary

Based on the tests some general conclusions can be made about the acceleration. First of all the acceleration seems to be working best for CIFS. This is a good thing, since based on the chapter 3 CIFS is the protocol suffering most from the poor performance. There also seemed to be a difference between HTTP and MAPI traffic. Also the CIFS test 2 (figure 23) the file download of the same file with HTTP and CIFS gave more benefits with CIFS download. This indicates that the rate the acceleration can reduce the data is dependable on the protocol.

Most variation in the test results was collected in the HTTP tests, as they were performed with 3 different types of files: first data was an installation package, second high definition pictures and third Microsoft office files. The results of these tests were different from each other, the download of the installation package showed minimal reduction in the data size (warm test) where the Microsoft file download showed data reduction already during the cold test. This indicates that the file / data type affects the acceleration results.

For the CIFS test the file name was changed during the tests. The accelerator handled it well as it did not affect the acceleration performance in any way, or if did the effects were minimal. Based on this the accelerator can function well even if the file names are changed.

These test results concentrated on the data reduction methods like data suppression, compression and object caching, although reducing the size of data transferred can also improve the response time for the user. Based on the tests acceleration seems to provide very powerful data reduction and therefore bring possibly high improvements on the WAN performance. For getting a better overview of the benefits that an accelerator can bring for a network it is recommended to perform tests in the target network. It is important to select the test environment carefully - a site can use multiple resources from various other locations and each location would then have to have an accelerator installed. It is important to investigate the way the applications are built to understand how they can be accelerated, this especially when the applications are customized for the target company and not working in they way then normally do.

Conclusion for acceleration performance and implementation:

1) Acceleration works well for CIFS, MAPI and HTTP.
2) Acceleration (reduction in data size) is dependable of the protocol used.
3) Acceleration (reduction in data size) is dependable on the type of data.
4) Some data can already benefit from the acceleration during the first time (cold cache) it is transferred through the network.
5) Acceleration (reduction in data size) is not noticeably affected by changes in the file names.
6) It is recommended to perform tests in the real environment
7) Thorough analysis on the applications planned to be accelerated is recommended

6.4 Future Work

These first tests concentrated only on the 3 main protocols CIFS, MAPI and HTTP. As a next step similar tests should be done with the other corporate applications, most importantly with SAP and AutoCAD, which were highlighted as critical applications for the target company. Measurements should also be performed within different kinds of network links, for instance with a low bandwidth high latency link, or with a high bandwidth low latency link. This helps to identify which parts of the network would benefit the most from the acceleration. As accelerators are also said to improve the response time, not only to reduce the data sent in the network, measurements on response time could be taken.

In order to get a more complete picture of about the impacts of an accelerator to the network measurements on longer time period should also be taken. The initial plan was to do this already during the case study, but the problems explained in chapter 6.3.4.5

prevented this. During the long term tests it would be good to monitor the accelerators under a heavy load, to see whether the benefits are as good as the case study tests indicate. One more aspect that could be considered during the long term tests is to do user interviews to see if they actually feel any difference in the application performance when acceleration is used.

Because the target company has a lot of small branch offices and also remote users, similar tests with the software version of the accelerator could be useful. As can be seen from the appendix 1, the performance of the software version is not as high as the performance of the accelerator device. Testing the difference in acceleration power between the two can help in deciding where a physical accelerator device is a necessity.

In general, performing similar tests with different accelerator manufacturers would be interesting. Because the case study network is mostly using the most common corporate applications, the difference in performance between different vendors might not be big. On the other hand, if the target network is used for more special data, or if the application uses for example UDP instead of TCP as a transport protocol, it can be that it is even beneficial to have tests with multiple vendors before implementation decision is made.

7 CONCLUSIONS

Due to the centralization of services and decentralization of employees to small branch offices, WAN has become a critical point for corporations. To understand how it is possible to improve WAN performance, it is important to recognize what is causing the poor performance. From the network perspective, utilization, oversubscription, latency and throughput can be considered as these factors. In addition, the transport protocol used by the network, in most cases TCP, has an impact on the performance.

The application protocols have an impact on the way applications work when used in WAN. Identifying application specific performance issues is important. This can be done by estimating how applications behave in different kind of network conditions, e.g. with high latency as different applications react differently. Current corporate application protocols are normally designed for LAN environments. Especially MAPI, CIFS and HTTP suffer from the move from LAN to WAN due to their chatty nature.

WAN acceleration is a potential solution for improving WAN performance. The techniques it provides can improve the response time of applications and minimize the amount of data travelling through the network. The acceleration techniques can be divided into two categories: WAN optimization and application acceleration.

Based on previous research and target company requirements, an accelerator of Blue Coat Systems was selected for tests in the real environment. The installation of the equipment was done by installing the accelerator behind the router using WCCP, because this was recommended based on theory. This is in conflict with experience gained in the implementation phase, as several WCCP related problems were encountered.

Finding out what should be accelerated is one of the main points to be considered when implementing acceleration into the network. It is recommended to use network monitoring for detecting how much of the bandwidth is utilized by different applications. Based on the results of the measurements combined with the theoretical results, CIFS, HTTP and MAPI are the best acceleration targets for the target company.

Based on both theory and the tests performed, WAN acceleration can improve application performance. The results of the case study indicate that WAN acceleration reduces efficiently the amount of data travelling through the network. Even if no measurements on the response time were taken, the reduction in the amount of transferred data has a positive impact on the response time. CIFS acceleration seemed to outperform the other two tested protocols. The following are the most important findings of the tests performed in this thesis:

1) The acceleration results are dependable on the protocol used.
2) The acceleration results are dependable on the type of data used.

When planning to implement acceleration, one should consider performing tests first, especially if the target applications are not using the most commonly known protocols.

REFERENCES

[1] Cisco, Internetworking Technology Handbook – Intro to LAN, [e-document] [retreived August 22, 2008] From Cisco at: http://www.cisco.com/en/US/docs/internetworking/technology/handbook/Intro-to-LAN.html

[2] Jaakohuhta, H. 2005. Lähiverkot – Ethernet. Helsinki, Edita Publishing Oy. 380 p. Chapters 1, 7. ISBN 951-826-787-1

[3] Grevers, T. Jr., Christner, J. 2007. Application Acceleration and WAN Optimization Fundamentals. Indianapolis, Cisco Press. 358 p. Chapters 1-6. ISBN 978-1-58705-316-0

[4] Comenr, D., E., TCP/IP. Jyväskylä, Gummerus Kirjapaino Oy, 2002. Original book: Internetwokring with TCP/IP principles, Protocols and Architectures, Fourth Edition, p. 674-719. ISBN: 951-826-435-X.

[5] IETF RFC 193 – Transmission Control Protocol. [e-document] 1981 [retrieved August 20, 2008] From: IETF at: http://tools.ietf.org/html/rfc793

[6] Bartlett, J. Sevick P., Moore, S., Economics of QOS on WAN Access Lines, Business Communications Review, Vol 34, Issue 10, Oct 2004. ISSN: 0162-3885 [retrieved May 20, 2008] From EBSCO at: http://search.ebscohost.com/

[7] Kansanen, M., Bachelor’s Thesis: Wide Area Network Acceleration in Corporate Networks. 2009. Will be published at LUTPub: https://oa.doria.fi/handle/10024/4000

[8] Cole, R. G., Ramaswamy, R. 2000. Wide-Area Data Network Performance Engineering. Artech House. 417 p. Chapters 1-2. ISBN 0890065691

[9] Taneja Group, Wide Area Data Services: Optimizing the Branch, [e-document], 2005, White paper [retrieved June 10, 2008] From Techworld at: http://www.techworld.com/whitepapers/index.cfm?whitepaperid=4053

[10] Spurgeon, C., 2000. Ethernet: The definite guide. Sebastopol. O'Reilly, 498 p. Chapter 19. ISBN 1-56592-660-9

[11] Bertsekas, D., Gallager, R. 1992. Data Networks. Prentice Hall, New Jersey 1992, 556 p. ISBN 0132009161

[12] Yao, L., Agapie, M., Ganbar, J., Doroslovacki, M. 2003, Long Range dependence in Internet Backbone Traffic," IEEE International Conference on Communications (ICC2(X}3). May 2003. pp. 1611-1615

[13] Kelly, T. Scalable TCP: Improving Performance in Highspeed Wide Area Networks. ACM SIGCOMM Computer Communications Review. [e-document] 2003, Volume 33, Number 2 [retrieved May 20, 2008]. From: ACM at: http://portal.acm.org/portal.cfm

[14] Jacobson, V., Braden R., Borman, D., IETF RFC 1323 – TCP Extensions for High Performance. [e-document] 1992 [retrieved August 20, 2008] From ACM at: http://portal.acm.org/portal.cfm

[15] Jacobson, V., Congestion Avoidance and Control, Proc. SIGCOMM '88, Vo118 No. 4, August 1988, [retrieved: October, 20th 2008], From ACM at: www.acm.org

[16] Floyd, S., IETF RFC 3649 - HighSpeed TCP for Large Congestion Windows [e-document] 2003, [retrieved October, 20, 2008] From IETF at: http://www.ietf.org/rfc/rfc3649.txt

[17] Siegel E., 2006. Optimizing WAN Performance: Accelerating Market Growth, [pdf-

document] Burton Group, ISSN 1048-4620 [retrieved May 20, 2008] From: http://www.zdnet.com.au/

[18] ITU-T recommendation G.114 [e-document] [retreived August 20, 2008] From ITU at: http://www.itu.int/itudoc/itu-t/aap/sg12aap/history/g.114/index.html

[19] Juniper Networks, Packet Flow Acceleration (PFA), White paper, 2005, [pdf-document] [retrieved November 15, 2008] From Juniper at: http://www.juniper.net/

[20] Hertel, R., 2003, Implementing CIFS – The Common Internet File System, Prentice Hall PTR, 641 p., Appendix D. ISBN 013047116X.

[21] Sevcik, P., Centralizing Microsoft Servers Hurts Performance, [pdf-document]Business Communications Review, Vol 36, Issue 5, May 2006, ISSN: 0162-3885 [retrieved April 18, 2009] From EBSCO at: http://search.ebscohost.com/

[22] Juniper Networks, Application Flow Acceleration for Web (HTTP) Traffic, White paper, 2005, [pdf-document][retrieved November 15, 2008] From Juniper at: http://www.juniper.net/

[23] Dierks, T., Allen, C., IETF RFC 2246 – The TLS Protocol Version 1.0 [e-document] 1999, [retrieved April, 18, 2009] From IETF at: http://www.ietf.org/rfc/rfc2246.txt

[24] Dierks, T., Rescorla, E., IETF RFC 5246 – The Trasport Layer Security (TSL) Protocol Version 1.2 [e-document] 2008, [retrieved April, 18, 2009] From IETF at: http://www.ietf.org/rfc/rfc5246.txt

[25] Chou, W., Inside SSL: The Secure Sockets Layer Protocol, IT Pro, [e-document], July/August 2002, [retrieved August 27, 2008] From IEEE at: www.ieee.org

[26] Chou, W., Inside SSL: Accelerating Secure Transactions, IT Pro, [e-document], September/October 2002, [retrieved August 27, 2008] From IEEE at: www.ieee.org

[27] Blue Coat Systems. Open, Manage and Accelerate SSL Encrypted Appilcations. [pdf-document]White paper, 2007 [retrieved August 27, 2008] From Blue Coat at: www.blucoat.com

[28] Sevcik, P., Wetzel, R., 2006, Why SAP Performance Needs Help, [e-document] NetForecast Report 5084 [retrieved April, 18, 2009] From NetForecast at: http://www.netforecast.com/ReportsFrameset.htm

[29] Rolfe, A., Skorupa, J. 2007. Magic Quadrant for WAN Optimization Controllers. [pdf-document]Gartner RAS Core Research Note G00153256. Gartner Inc. [retrieved May 20, 2008] From Gartner at: www.gartner.com

[30] Sevcik P., Wetzel, R., Pocket Guide to Application Delivery Systems, [pdf-document] Business Communications Review, Vol 34, Issue 9, Sept 2004, ISSN: 0162-3885 [retrieved May 20, 2008] From EBSCO at: http://search.ebscohost.com/

[31] Blue Coat Systems. Technology Primer: Byte Caching, [pdf-document],White paper, 2007, [retrieved December, 6, 2008] From Blue Coat at: www.bluecoat.com

[32] Willis, P., J., An introduction to quality of service, [pfd-document] BT Technology Journal, Vol 23 No2, April 2005 [retrieved December, 4, 2008] From Springer at: www.springer.com

[33] Newman, D., WAN acceleration offers huge payoff, [e-document] Network World, Aug 2007, [retrieved December 8th 2008] Available at Networkworld: www.networkworld.com

[34] Claise, B., IETF RFC 3954 – Cisco Systems NetFlow Services Export Version 9 [e-document], 2004, [retrieved April 11, 2009] From IETF at:

http://www.ietf.org/rfc/rfc3954.txt

[35] Blue Coat Systems. Technology Primer: Compression, [pdf-document], White paper, 2007, [retrieved December, 6, 2008] From Blue Coat at: www.blucoat.com

[36] Blue Coat Systems. Technology Primer: MAPI Protocol Optimization [pdf-document], White paper, 2007, [retrieved December, 6, 2008] From Blue Coat at: www.blucoat.com

[37] Blue Coat Systems. Application Performance Brief: Microsoft Office [pdf-document], White paper, 2007 [retrieved December, 6, 2008] From Blue Coat at: www.blucoat.com

[38] Atia, O., Future Trend for IP Services over FAA Telecommunications Infrastructure, [pdf-document] Integrated Communications, Navigation and Surveillance Conference, May 2008, IEEE [retrieved December, 3, 2008] From IEEE at: www.ieee.org

[39] Fowler, D., G., Application of Acceleration Technology to Military Sealift Command Afloat WAN Infrastructure, [pdf-document], 2006, Military Communications Conference MILCOM 2006. [retrieved December 4th 2008] Available at IEEE: www.ieee.org

[40] Metzler J., Optimizing Performance For Your Branch Offices, [pdf-document] Business Communications Review, Vol 37, Issue 7, July 2007, ISSN: 0162-3885 [retrieved May 20, 2008] From EBSCO at: http://search.ebscohost.com/

[41] Riverbed Technology. The Riverbed Optimization System (RiOS) 5.5, [pdf-document],White paper, 2008, [retrieved December, 6, 2008] From Riverbed at: www.riverbed.com

[42] IETF Internet-draft – Web Cache Communication Protocol V2.0 [e-document] 2001, [retrieved December, 2, 2008] From IETF at: http://tools.ietf.org/id/draft-wilson-wrec-wccp-v2-01.txt

APPENDIX 1. WAN Accelerator Comparison

The following tables present a comparison of 3 different accelerator vendors: Blue Coat Systems, Riverbed and Juniper Networks. Comparison is divided into 3 separate sections scalability, transparency and device characteristics. Each focus area is further divided into questions and the comparison is done by answering to these questions.

SCALABILITY	Blue Coat	Riverbed	Juniper
1) Are there solutions for different type of sites?	Yes, from software version to datacenter version.	Yes, from software version to datacenter version. Said to work well in large environments.	Yes multiple devices. No software version.
2) How big is the largest solution?	Unlimited amount of concurrent users, 16GB RAM	WAN capacity 310Mbps, Data store capacity 1,4TB, 40 000 concurrent TCP connections	10 Mbps to 45 Mbps, up to 140 tunnels supported, disk space 1 TB.
3) What is the software version like?	Accelerates remote applications up to 35 times. Software available free.	Accelerates remote applications 5 to 50, measured maximum 100. Licensing costs for the software version.	No software version.
4) Number of concurrent TCP connections (for a middle scale equipment)	19500 (*)	12200 (*)	32000-64000 (*)

(*) The figures for Blue Coat and Riverbed are taken from [33] and are comparable, Juniper results are not comparable to others

TRANSPARENCY	Blue Coat	Riverbed	Juniper
5) Can the accelerator be installed in a transparent mode?	Yes, and BlueCoat also recommends to this implementation.	Yes, possible to keep the source and destination. Riverbed says it doesn't work as well as with a non-transparent solution.	No information, that says Juniper uses transparency, was found.
6) Does the solution offer its own monitoring system?	Yes	Yes	Yes

(CONTINUES)

DEVICE CHARACTERISTICS	Blue Coat	Riverbed	Juniper
7) Which acceleration techniques are used?			
a) TCP Optimization	TCP proxy, and TCP optimization	TCP proxy, TCP window scaling, High Speed TCP, Max Speed TCP	Yes (active flow pipelining)
b) Data Suppression	Yes (called byte caching)	Yes	Yes (Molecular Sequence Reduction)
c) Compression	Yes (gzip/deflate algorithms)	Yes (Lempel-Ziv based algorithm)	Yes
d) Object Caching	Yes	No - says it is not as effective as their system	No
e) Read-Ahead	n/a	n/a	Yes (called read/receive)
f) Write-Behind	Yes (called batcing)	n/a	Yes (called write/send)
g) Message Prediction	n/a	Yes (called connection pooling)	Yes (called pre-fetching)

DEVICE CHARACTERISTICS	Blue Coat	Riverbed	Juniper
8) What kinds of protocols or applications are supported by the accelerator?			
a. MAPI	One of the key acceleration targets. Special solution for branch offices	Yes. Said to work well [33]	Yes, done with message prediction
b. CIFS	One of the key acceleration targets. (mainly object and data suppression)	Yes, said to work well in [33], done with data suppression	Yes, done with message prediction
c. HTTP	one of the key acceleration targets (compression, data suppression)	Yes (message prediction)	Yes (fast connection setup and message prediction)
d. HTTPS (SSL)	Yes, all use similar methods	Yes, all use similar methods	Yes, all use similar methods
g. UDP	No	Yes	Yes
i. SAP	Yes (for example for login times)	Yes	Yes
k. CAD	Yes	Yes	Yes (TCP optimization)

(*) The figures for Blue Coat and Riverbed are taken from [33] and are comparable, Juniper results are from www.juniper.net (not comparable to others)
Other resources are: www.bluecoat.com, www.riverbed.com and www.juniper.net