The Souls of DOCSIS®

Lehr und Kunst in Developing the World's First Standard Cable Modem and Internet Protocol Multimedia Delivery System

Volume 1:

Vendor Benchmarks beget a Magic MAC

Robert F. Cruickshank III

Big Indian, New York, USA

DOCSIS are the Data Over Cable Service Interface Specifications located at www.cablemodem.com

“Lehr und Kunst” means “Theory and Practice” in the German language.

The souls of DOCSIS rest when

no troubled home or business is left behind.

Bob—
It's been a great couple of years!
Thanks for getting me on-board & all your on-going support — I really appreciate it — even if it doesn't always show!
Best of luck in Herndon & let us know whenever we can help!
Andrew

Bob-
Look forward to hearing from you from back East - Best Wishes!
-Sarah ☺

Bob
Good Luck in your future endeavors
The E-Tech Gang

BOB-
ARE WE ALL CERTIFIABLE YET?
IT'S BEEN FUN - I LOOK FORWARD TO MORE! STAY IN TOUCH.
CHUCK A.
& THE BROADCOM GANG

Best of luck Bob
-Alex-

Best of luck!!!

Bob, Best of luck in your new job. Stay awesome.
Bill

Bob.
It was a wonderful experience to have DOCSIS activity led by you.
Thankyou & have good luck!
Hide & TOSHIBA/TURBO TEAM

Bob
Wish you well in your new position. I appreciate all the kindness, patience, and opportunity that you gave to me. Roger

Bob.
Hope this ...
you hope it ...

BOB,
GREAT WORK FOR YOU AT CABLELABS, GOOD LUCK,
Will

Wishing you Happyness
— Vam
Cisco Systems

Bob-
Every time I have lobster, I'll think of you. The nice thing is I'll still be able to give you you know what!
Dotty

Bob
It was great working with you even the DOCSIS. It has been great team (DOCSIS) because of your leadership. Hope we keep on meeting and also hope to be part of ... project
All the best
Bhef Shivale

Phasecom.

Bob, It's been a blast!
Take care and keep your ... level! Bill W.

BOB,
HAVE FUN
GOOD LUCK,
Christopher

Bob-
Have a great time & good luck!
Rima Rhone

BOB-
I'll miss you!
Who'll format the OOP 2
Take care
Kelly

Bob,
Wish you best on your new job.
you will be missed.
Majid

Bob.
Congratulations!!
We'll miss you.
Come back to see us.
Mukesh.

Bob -
Have a high-speed time!
Dan T.

Bob,
Good luck in your new adventure.
let's see those cable modems in the field

Bob, thanks for everything.
TNX ITP. lifer,
C. Clark

Bob,
Best of luck!
Stay in touch!
Laurie

Bob,
You must feel like I do. (AHHHH)
Congradulations!
Melanie

Bob,
Good luck with your new adventure.
Have fun & ...
Dolly

Bob,
THANKS FOR HELPING ME TO ESCAPE THOSE COLD CANADIAN WINTERS!
TAKE CARE BUD.
Jason

Bob,
It's been great working with you.
Your encouragement and perseverance is extraordinary and appreciated.
Charlie William ...

Good luck Bob
Kathy Motte

Good Luck!
Sandy

Acknowledgments

For reinvigorating work on this manuscript and numerous helpful connections, I thank G. Kim Dority.

For encouraging and publishing this work, I'm indebted to the capable staff of The Cable Center, including Larry Satkowiak, Jana Henthorn, and Diane Christman.

For tireless editing, I owe a special debt of gratitude to my editor, Cynthia Lang, at The Cable Center.

For picking up the DOCSIS story 20 years ago, right about where this document leaves off, I am forever thankful for the support of cable journalists Roger Brown (RIP), Paul S. Maxwell, Leslie Ellis, Stewart Schley, Jonathan Tombes, Carl Weinschenk, Jeff Baumgartner, and Alan Bresnick.

For their encouragement and gracious support in reviewing this document, I'd like to thank the following:

CableLabs: Dr. Richard R. Green, Christopher J. Lammers, Daniel J. Rice and Wayne Surdam.

SCTE: Mark L. Dzuban

YAS Capital Partners: Dr. Rouzbeh Yassini

DOCSIS Core Team: Thomas E. Moore, Brian S. Reilly, Pamela D. Anderson, William P. Kostka, Jason K. Schnitzer, Andrew W. Sundelin, Kazuyoshi Ozawa.

Superior Court Judge and Adjunct Professor at USC Gould School of Law: Hon. Daniel L. Brenner

Industry Reviewers: Louis A. Steinberg, Andrew J. Audet and any others I may have missed citing above.

I'd like to thank CableLabs and member company CEOs and CTOs for granting me the opportunity to lead, and for encouraging us along the way: John C. Malone, William T. Schleyer, Tony G. Werner, Brian L. Roberts, David M. Fellows, Mark A. Coblitz, James A. Chiddix, Dr. Mario P. Vecchi, Edward Samuel "Ted" Rogers Jr., Nick Hamilton-Piercy, George Hart, James Robert "JR" Shaw, Randy Elliot, James O. Robbins, Alex B. Best, Jay A. Rolls, William D. Bauer, and many others who believed in our efforts and ensured our success in building over 1 Billion DOCSIS modems.

This was a meaningful, memorable, and wonderful experience to have shared with you all. Thank you!

Dedication

This work is dedicated to three groups of individuals. First, the gifted and capable teachers, professors, and instructors who touched my life in such a profound way that I began to believe in my ability to do almost anything. The materials and concepts they presented, the challenges they asked be overcome, and the loving kindness they shared remain an inspiration!

I owe a debt of gratitude to the Peddie School and it's wonderful teachers: Jim Brink, who gave me a love of physics; John Sprout, who taught me functions and single-handedly convinced me that my mathematical limitations would not prevent me from being an engineer (and who sent me a birthday card every year for over a decade); Jeffrey Holcombe (RIP), who gave me a love of show business and taught me to speak, sing, and act; Charles Haney, who helped me learn to write English; and dear Mr. Everett Swift, who got me hooked on studying history with J. K. Galbraith's *The Robber Barons*.

I also owe a debt of gratitude to the fine institution that is Worcester Polytechnic Institute (WPI), where I learned that good engineering was based on Theory and Practice, or as the Germans say "Lehr und Kunst." WPI's wonderful teachers included: Harold S. Corey, who taught sketching and drawing so lovingly and so well that I took his class twice; Walter A. Kistler who graciously taught analytics, programming, and so much more; John A. Van Alstyne, who taught with great patience, introductory calculus; Michael M. Sokal, who taught the importance of studying the history of technology; Joe Gale, Bobby Taylor, and John Grzyb who provided instruction and materials to fabricate anything (even trying to cut and convert a 1969 SAAB Model 96 into a motorcycle trailer); and most of all to the one man, who like a father, selflessly gives himself over and over to his children, or in this case his students– Ryszard (Rich) J. Pryputniewicz.

As a world-recognized senior research scientist and engineer, Rich organized and produced numerous research conferences, symposia, and papers. He contributed in hundreds of professional society activities, published over 375 research papers, and gave more than 700 lectures and seminars—all while teaching full academic loads and developing the Center for Holographic Studies and Laser micro-mechaTronics (CHSLT) as well as WPI's NanoEngineering, Science, and Technology (NEST) Program. Yet, among all these accomplishments, what is most remarkable about Rich is that his research contributions were made while always putting the education and welfare of society and his students above his own needs.

I first met Rich when entering WPI's Mechanical Engineering program as a freshman in the fall of 1979. I took Rich's Engineering Experimentation class and by Christmas was employed by this wonderful researcher as his work-study student. For the next six years, I cherished time with this man and by the time I graduated with my BSME in 1983 and MSME in 1985, I was his graduate teaching and research assistant. He was also my academic advisor, major projects advisor, and master's thesis chairman.

At the time, the CHSLT was a very impressive single-room laboratory. By the time I left, we had grown the CHSLT to eight laboratories. Last I checked, it is 14 laboratories! The CHSLT has always been a place where unique “not available in the literature” experiments are conceived and built—revealing for the first time amazing phenomena to be explored, explained, and quantified.

During our years together, I came to know and love Rich very much. Of all the wonderful things he shared with me, the most memorable were his kindness, generosity and scientific method. Rich’s method was thorough, yet simple; easy to understand, yet challenging. When tackling a new research project, the first step in the laboratory was to create a real-life prototype of the system to be observed and the physical phenomena to be explored. Rich would carefully describe the experimental goals and we the lab staff would propose building a physical apparatus and data acquisition system that would enable us to recreate, capture, and quantify the environment to be studied. My favorite memory of Rich is his never-ending enthusiasm for our experimental apparatus suggestions, which were almost always met with his cheerful command: “Get it! Do it! We’ll need it!” This meant “Start shopping and generating purchase orders now! Start building now! Work around-the-clock and get the parts and subsystems you need! Fabricate what you need in the machine shop! Assemble the apparatus! Call me anytime day or night as soon as the experiment is running! Meanwhile, I’ll be staying up all hours getting the experimental apparatus and methodology fully documented!” In this way, we conceived, built and researched huge and elaborate systems; the air conditioning and power requirements alone were staggering.

Rich’s research, in my view, has always been about saving and improving lives. The very first experiment he shared with me was the holographic optical inspection of primary stage and secondary stage compressor and turbine blades of the United Technologies, Pratt & Whitney JT-9D high-bypass turbofan jet engine. At the time, the JT-9D was the largest commercial airliner engine ever mass-produced and was the first engine that made the Boeing 747 fly (thousands are still in service). I was immediately moved by the gravity of the responsibility we were undertaking as we (with Rich’s guidance) looked for tiny sub-visible-wavelength sub-surface manufacturing defects that could cause catastrophic engine failure, and in so doing, put hundreds of lives in the air and thousands of lives on the ground in harm’s way.

Another experiment for the Wyman Gordan Company involved researching and developing the theory that later enabled the development of an automated optical inspection system for very-fine-mesh powdered metal sifting and blending screens. The powdered metals sifted through these screens were later single-impact-forged in the largest press in the world to produce life-critical parts such as helicopter main rotor hubs (where all the lifting blades are attached to the main drive shaft), which simply cannot fail without certain death.

Rich had an uncanny ability to enthuse industry chieftains when he gained interest from organizations unlike one another as General Dynamics Electric Boat, Amp, and Chesebrough-Ponds. Under Rich’s guidance and leadership, we also built and researched all sorts of fixed and mobile (even battery powered) systems for making optical-based free-space and underwater measurements of mass, temperature, time, distance, vibration, translation, rotation, reflectivity, and surface roughness.

More than any other professor at WPI, Rich helped me become a research engineer. Certainly there were many other wonderful professors who touched my life, but Rich Pryputniewicz's research gifts to myself and those students around him rose well above those of all the professors I had the honor of working with during my six years at WPI and five follow-on years at the University of Colorado at Boulder. Suffice it to say that along with all the great research and research skills this man has given to his students and the world, he still made time to teach me and those around him who wanted to learn: individual responsibility, scientific responsibility, social responsibility, willingness to share responsibility, willingness to lend a hand, kindness, generosity, ethics, composure in the face of disaster, encouragement, confidence, pride of accomplishment and much, much more. The research skills, life values and "go for it" attitude I learned from Rich have served me well over the years. Thanks to Rich, I am a better researcher and am closer than ever to my pre-WPI goal of reducing global warming and fossil fuel consumption through improved efficiency of electrical energy generation, transmission, distribution and use. Thank you WPI, and most of all, thank you Rich.

I also owe a debt of gratitude to the wonderful faculty and staff at the University of Colorado at Boulder– from the Interdisciplinary Telecommunications Program and Laboratory: Professors Stan Bush, Gerald Mitchell, Frank Barnes, and many others who taught the fundamental bits and Bytes of data and telecommunications; from the Department of Civil Engineering: Professors Jan Kreider, Mike Brandemuehl, Louis Sommers, and others who taught the fundamentals of solar power, heating ventilation and air conditioning, intelligent control systems, and so much more; from the Electrical and Computer Engineering Department: Professors Michael Mozer, Renjeng Su, Ewald Fuchs, and others who taught the fundamentals of neural networks, digital design, analog to digital converters, digital to analog converters, power systems engineering, vision and pattern recognition systems.

The second group of individuals to whom this work is dedicated is the management and staff of Cable Television Laboratories (CableLabs) and the cable industry Multiple System Operators (MSOs) that support CableLabs. Relentlessly from the very outset, a core group of these people understood the societal and business importance of cable modems and encouraged us to forge ahead. Many whom I will tell you about herein were mentors who taught what we needed to know and showed us the way. As time passed, the number of believers loyal to the DOCSIS project grew exponentially. At every turn, the commitment of these individuals consistently turned impossibility into reality. Read on and I will tell you about them and the adventures they made possible!

The third group of individuals to whom this work is dedicated is the tens of thousands of men and women of the cable television companies throughout the world who, every day, continually seek to improve service levels. Read on and I will tell you about them and what they do—and what we did and still do to help these dedicated individuals provide the best service levels possible.

Last but not least, I owe a life-long debt of gratitude to my Grandparents, Aunts & Uncles, Mother, Father and Siblings who instilled the value of education and love of farming, engineering and machinery, and to my wife Laurie for inspiring me every day.

Introduction

The DOCSIS project marked the beginning of something critical that the Cable TV industry still needs to finish. As a team, we defined methods that remotely diagnose service problems in an automated fashion—sometimes before individuals or groups of subscribers experienced degraded TV, phone or Internet service. After minimal industry success implementing remote diagnostics, much of the original DOCSIS Team came back together and engaged in the furtherance of these methods by writing software to allow cable operators to proactively address issues that, if left unaddressed, would result in extended periods of degraded service to cable customers.

Together, the reassembled DOCSIS Team built software to detect and quantify degraded service causing customer discomfort. Our software also verifies the successful installation of modems and set-top boxes in customers' homes, capturing performance parameters in a birth certificate that even includes a "born on" date—just like your beer. Most importantly, we built software to proactively evaluate every management team's ability to grow the network and customer base while providing the best possible service on a market-by-market basis.

Out of absolute necessity to meet pressing operational demands of the business, others less informed in the art and science of DOCSIS also built software tools—they are still building software tools to this day. However, not having the benefit of guidance from "being there" during the birth of DOCSIS, the vast majority of homegrown tools often lack critical functionality for improving business operations and customer satisfaction. Despite setbacks, the industry has accomplished a lot, but has a long way to go.

Regrettably, many still don't understand the importance and value of all the DOCSIS "management hooks" that we created and we have, as an industry, lost the better part of two decades in getting our customer service act together.

For example, though we have the technology to account for and appropriately notify cable company staff of each and every customer with a DOCSIS connectivity issue, the vast majority of customers still have to call their cable company for assistance when suffering from degraded service.

The good news is that we have the tools necessary to improve customer experience and operations efficiency. There are an increasing number of cable companies using recent advances in database and mobile technologies to fully leverage DOCSIS management hooks to proactively ensure customers have the best service levels possible. I'd like to acknowledge the people who understand DOCSIS well enough to operationalize these management hooks as the real heroes in our industry!

Preface

Once in a great while you may come across a job that really suits you and those around you. A rare exception to the norm of everyday life, an opportunity arises and allows you to soar as a team like you have never soared before. The DOCSIS Project was just that kind of job!

I spent most of my high school years in the wood shop, the metal shop, the photography laboratory, and the automotive shop. While I enjoyed these places and activities very much, they did not help me get good grades in traditional studies. By the end of junior year in high school, I came to the realization that my academic grades were not going to get me any further than the local community college. So, I was destined to find a college preparatory school where I could repeat my junior year in the hope of getting grades that would get me admitted to a great college.

After five years of high school, I spent the customary four years in Worcester Polytechnic Institute to complete my undergraduate degree in mechanical engineering. I had met some amazingly gifted faculty and staff in prep school and college, and was given the opportunity to return for my master's degree with all tuition paid and a stipend to cover monthly expenses. In evaluating this opportunity, I decided to speak with my dearest friends who, having completed high school and college in eight years versus my nine, had by now been in the "working world" for a whole year.

"Work sucks!" and "Stay in school!" was the resounding reply from my dear friend Fran Mastoloni. I believed him and stayed on to get my master's degree in mechanical engineering. Of course, there was another huge deciding factor in staying on at Worcester Polytechnic Institute—I loved my boss and I loved my job! It was as if I was in some sort of Charles Dickens *Christmas Carol*. I had never worked for a man who so closely who resembled Ebenezer Scrooge's beloved boss, Mr. Fezziwig.

> "Yo Ho! My boys," said Fezziwig. "No more work tonight! Christmas Eve, Dick! Christmas, Ebenezer! Let's have the shutters up!" cried old Fezziwig with a sharp clap of his hands.
>
> – *Christmas at Fezziwig's Warehouse by Charles Dickens*

Happy to have stayed on to learn more engineering, I finished my Master of Science in Mechanical Engineering two years later and was thrilled to start my first job at AT&T Bell Telephone Laboratories in Denver. But a harsh reality set in quickly. Fran was right, work did suck. Here I was at one of the most prominent laboratory organizations in the world, yet was bored out of my mind. I made many great friends and we had some wonderful times together, but as a whole it was like the *Dawn of the Dead* movie. It was as if the staff (except us kids) all needed coffee, but already had four cups!

It was here that I met Mike Nienaber a young engineer who, as chance would have it, had worked on multiple product designs that for one reason or another never got out the door. Don't be mistaken, Mike designed good products. But somewhere along the line, management or marketing decided the products were not to be sold, leaving Mike with what felt like another product failure. More than anybody else, Mike instilled in me the importance of getting a product "out the door." He never spent a whole lot of time telling me why it was important, but I could see in his face the frustrating effect of having worked on so many projects for nothing.

Table of Contents

Chapter One

ISDN: Innovation Subscribers Don't Need

While at AT&T Bell Telephone Laboratories, I witnessed a product "road kill" that was bigger and sadder than anything I could imagine.

It was an incredibly exciting and promising time in the telephone industry. The 100 year-old analog telephone network was transitioning to an all-digital access network. The technology that made this possible was called ISDN, the Integrated Services Digital Network.

The promise of ISDN was enormous. Its ability to enable new types of telephone services was mind-boggling. The potential power of ISDN was its ability to allow subscribers (customers) to "signal" their intentions to the network like never before. To appreciate the potential impact of ISDN customer signaling, step back for a moment and consider the evolution of signaling methods used in telegraph and telephone networks. When Samuel B. Morse invented and productized his system for the transmission of messages in 1840, the entire signaling method consisted of only two symbols; the dot and dash. In Morse Code, the dot and dash are assembled to send messages; such as the internationally recognized distress message SOS with a dot-dot-dot, dash-dash-dash, dot-dot-dot.
(…---…)

The dot and dash served us well, and the dot alone was all we used for signaling the telephone network—for over 100 years! Those old enough to remember rotary dial telephones may recall that the dial sent out a cadence of dots (aka pulses, flashes) that sounded like clicks when you put the handset to your ear. For example, when dialing the number three you would hear three clicks. And if the dial was locked or removed to prevent unauthorized use, one could simply "click out" the desired number by flashing the switch hook (i.e., pressing and releasing the switch hook quickly). Though it was difficult to mimic the exact cadence of the rotary dial for long telephone numbers, one could simply flash the switch hook ten times (the equivalent of dialing zero), to reach the operator who would complete the call. Try it yourself. This simple signaling method still works today.

It wasn't until over 120 years after Morse's invention that the signaling method used in the telephone access network evolved. In 1963, when the Touch-Tone® telephone was introduced by AT&T, the signaling method made available to customers was expanded to include the numerals 0-9, plus * and #. What a cornucopia of functionality was now possible! The number of symbols in the subscriber signaling alphabet grew overnight from one (the dot) to thirteen.

With Touch-Tone, new services came like rain from the sky! One of the first new services was paging, where one could enter a callback number using touch-tones. Almost overnight, every doctor had a pager and soon many others in the service industry got a pager too. Along came voicemail, where touch-tones enabled customers to identify their mailbox and their password, and perform functions like play, skip, previous, next, delete, etc. Then came tele-banking, credit card authorizations, gasoline pay-at-the-pump, and much more, with just the twelve symbols 0–9, plus * and #.

To appreciate the impact of ISDN, suffice it to say that the customers signaling alphabet grew from twelve symbols in Touch-Tone to any combination of 256 symbols in the American Standard Code for Information Interchange (ASCII). Overnight it seemed that the signaling opportunities were endless. What customers could ”say to” and “hear from” the network seemed completely unbounded. Even the fundamental concept of having to make a call to find out simple things was dramatically changed. For example, with ISDN one could ask the network the current local time without calling anybody; both the query and the reply would be handled by ISDN signaling. A telephone call was no longer required to listen to a time recording, or to ask an operator the current time. With ISDN, the possibilities were endless. And for many, ISDN meant I Smell Dollars Now!

I was such a believer in ISDN. In my excitement, I built tiny network-powered ISDN sensors to be used in intelligent “smart building” monitoring and management systems that would be part of what we today call the Smart Grid. Akin to, but much more than a thermostat on the wall, these integrated sensors could measure the presence of occupants, room temperature, light levels, and sound levels—in order to provide lighting and heating/cooling needs for the occupants, and to save energy, reduce global warming, and conserve natural resources when occupants were away. These sensors could be plugged into and powered by the phone network and didn’t require any other wires.

And now for the sad ”road kill” part of the ISDN story. Perhaps it was capitalism; the companies involved in the development of ISDN for North America simply were not able to get along. They fought over which specifications should be in the standards. They made last minute changes to the standards (for example to the physical transmission layer) well after initial deployments were underway. And most devastating of all, left to their own devices, they committed the gravest of mistakes; they misinterpreted the standards and never did sufficient interoperability testing until it was too late—until after they had fielded their finished products! The net effect was that ISDN simply didn’t work. A Nortel ISDN phone would not work on an AT&T ISDN switch and vice versa. ISDN was stillborn in North America and the promise of powerful customer signaling didn’t materialize. People started mocking ISDN saying that it stood for: Interface Subscribers Don’t Need, It Still Doesn’t Network, It Still Does Nothing, and I Still Don’t Know. The moniker of I Smell Dollars Now was dead.

I was disgusted. I went to see the just released Tom Cruise movie *Top Gun* and almost immediately took a military leave of absence from AT&T to go have some fun learning to fly jets and transports in the United States Air Force. What a refreshing trip!

I returned to my job at AT&T eighteen months later, and on nights and weekends flew the Lockheed C-130 Hercules transport aircraft locally and throughout North America for the Wyoming Air National Guard. Even with the excitement of part-time flying, within another year I was again bored out of my mind, prompting me to start an interdisciplinary Ph.D. to further my energy management and “smart buildings” research at the University of Colorado at Boulder.

A year later, despite my second “flying” job and the thrill of going to school at night, I was again bored with my day job and took advantage of AT&T’s newly announced Selective Enhanced Leave of Absence (SELOA) to join a small six-person startup, CyberLYNX Gateway Corporation in Boulder, as their vice president of energy management. The AT&T SELOA program offered a safety net to departing employees who were eligible to return to AT&T at any time during the first year. Although, we went out of business at CyberLYNX twice in one year, I never returned to AT&T. Instead, I chose to pursue my energy management Ph.D. fulltime at the University of Colorado, while flying part time for the Wyoming Air National Guard.

Chapter Two

My first week and indoctrination to Cable TV

By 1994, I had been working for several years in the Laboratory of the Interdisciplinary Telecommunications Program (ITP) at the University of Colorado at Boulder. We had an extensive well-equipped laboratory facility and the instructors were fantastic.

I remember exactly where I was standing in the ITP lab looking, for the first time, at a demonstration of the very first Mosaic web browser from the National Center for Supercomputing Applications at the University of Illinois. Up until this point we, and the rest of the Internet world which was comprised mostly of the scientific community, had been using very crude programs for transferring files among computers. FTP was one such program that lives on today; another program was called Gopher. Suffice it to say, that the Mosaic browser enabled the Internet and was an indescribably huge advancement that is still being felt today. Mosaic was the basis of all web browsers such as Microsoft Internet Explorer, Netscape Explorer, Mozilla Firefox, Apple Safari, Google Chrome and any other browsers that may come along. Mosaic was the cornerstone of the World Wide Web.

I was actively testing Asynchronous Transfer Mode (ATM) switches for Network World Magazine, and had been published in their February and August 1994 issues. What I found most interesting about ATM at the time, was that it seemed the telephone companies were screwing up the deployment of this new technology similar to how they had screwed up ISDN. In ATM, there is the concept of Switched Virtual Circuits which are very smart and efficient, and can be set up and torn down as needed just like a telephone circuit that is connected (set up) as needed, and then terminated (torn down) when no longer needed. However, ATM also has the concept of Permanent Virtual Circuits that are very dumb point-to-point connections, which remain "up" and consume network resources forever—whether or not there is any data to be sent! The analogy would be like calling Grandma and both of you leaving the phone off the hook forever, just in case you might ever want to say something to each other again. Under this scenario, in order to call anyone else, you would need a second phone line and a second telephone. From a network resource efficiency standpoint, I couldn't figure out why in this new world of ATM all the telephone companies were deploying permanent virtual circuits. It seemed so wasteful to me.

About this time, I met two outstanding gentlemen engineers, Tom Moore and Brian Reilly, who helped me with the Network World magazine ATM Switch Testing program. We quickly became close friends and I treasured the time with them. Little did I know then how much they would help me in the next year at CableLabs.

One day in the summer of 1994 the telephone rang. It was Anna Maria Larson, a dear friend and executive search consultant from Larson Consulting International who had been retained by Scott Bachman of Cable Television Laboratories to search for a project manager, telecommunications projects. I knew Anna Maria from my CyberLYNX days, which made it very easy for us to discuss this new position at CableLabs. I explained that I was very busy with my studies and she explained that I was "made for this opportunity." She sent me the job description. We talked some more about all the telecommunications and data communications acronyms in the job description, most of which I happened to know very well. With nothing to lose, I decided to go for an interview. I liked everybody at CableLabs a lot, especially Richard (Dick) Green the CEO and Scott Bachman, Vice President of Operations Technology Projects who was the hiring manager. During the interview just south of Boulder, Dick and I found common ground and talked about thermal imaging of the moon. CableLabs was cool and fun! Scott soon made me an offer that as a student with a part-time job and a young family, I simply could not refuse.

My first week in Cable TV commenced in September 1994 when I reported to CableLabs to work for Scott. I had asked Scott for, and he gave me, a beautiful window office with a grand view looking out over the Flatirons foothills of the Rocky Mountains. My initial responsibility was to help fellow employee, Craig Owen (across the hall), evaluate the soon to arrive responses to the first—and already famous—Telecommunications RFP (Request For Proposal) issued by the cable television industry. The issuance of the Telecommunications RFP had been surrounded by a lot of hoopla on Wall Street, with personal appearances and a press conference by the CableLabs Executive Committee. This committee included the chairman of CableLabs, Dr. John C. Malone, who was also chairman, president and CEO of TCI (Tele-Communications Incorporated), the largest cable company in the United States at that time.

But first, I needed to learn about Cable TV and Scott Bachman made sure of that! Within the first couple of days, I had met and dined with senior engineering executives including Tom Staniec of Time Warner Cable who, along with the others, opened my eyes to the challenges faced in this wild and fast growing industry.

The foundation of my "Cable 101" practical education came in the first week after two days riding in trucks with the local Scripps-Howard cable company in Louisville, Colorado, performing cable TV service installs, upgrades, disconnects, and outside plant maintenance. Many times during those two days my eyes were opened so wide they almost popped out of my head. Was the whole cable television industry really doing business like this? These cable technicians in the field had no network monitoring tools to help them.

On our first job that day, we did a new service install in a lady's apartment. The TV was on one side of the fireplace and the existing cable TV outlet was on the other side of the fireplace, and the installer technician simply ran the coax cable right across the floor in front of the fireplace. Job done in record time! I was shocked that a cable across and on top of the hearthstone was satisfactory customer service.

On a second job we were adding an additional outlet for cable TV in another room. The house was pre-wired for cable TV and there were cable outlets in most all the rooms. The only thing we had to do was to determine which of the wires downstairs fed the specific room that was to now be turned on. The installer took a standard 9-volt battery from his pocket that had a crude cable TV F-connector attached to it and stuffed it into the mating F-connector on the wall—without knowing if that particular room's cable was hooked up to anything else on the cable TV network! Then, we went downstairs to where all the cable TV wires from the whole house came together. One by one he put the ends of each cable on his tongue and in so doing, detected which cable ran to the room with the 9-volt battery attached. In no time, he felt that familiar 9-volt tingle telling him which cable he needed to terminate. He put an F-connector on that cable, screwed it into an existing coax signal splitter (9-volts and all) and after returning upstairs to disconnect the battery and hook up the TV, we were done. Another job completed in record time!

A third job that also stands out was detecting and then finding a split in the cable TV network trunk cable buried underground in the neighborhood near a (cable TV) distribution pedestal. This was a scene of several technicians with shovels, a pick axe, giant cable cutting shears, two large mating cable connectors, a surgically sealed piece of large and incredibly gooey and sticky heat shrink tubing, and a blowtorch. But, the real magical instrument here was a Geiger counter-like device that would wail louder and louder the closer we got to the actual split in the buried cable. Another highlight of the day was putting on a pair of gaffes and a safety belt, and unassisted, climbing up then back down a telephone poll.

A very important take away from my first few weeks in cable TV was that the cable industry staff—who build and maintain the network and interact with customers every day—need as many network intelligence tools as possible to guide their daily efforts, especially while working on outside plant and onsite in customers' homes. I have never forgotten this need and to this day, continue to work on putting cable TV network monitoring and service-level intelligence into the hands of our industry's technicians, customer care staff, engineers, supervisors, managers, directors, vice presidents and general managers.

Chapter Three

What were these vendors thinking and just how well would these modems perform?

Responses to the first CableLabs Telecommunications RFP came from "all the usual suspects"[1], such as Nortel, AT&T, Siemens, and Alcatel. In addition, responses also came from many companies that we had never heard of before. Each vendor provided approximately fifteen or so copies of their response. Some responses were in multiple three-ring binders that were each two inches thick. With nothing practical to grasp in our hands except volumes of paper, and with our minds grasping at many new and old thoughts and concepts described therein, evaluating the Telecommunications RFP was an enormous and challenging theoretical exercise.

Most of the 43 responses to the Telecommunications RFP described a technology known as Next Generation Digital Loop Carrier (NGDLC), a simple and not very novel approach to timesharing bandwidth among customers. I very quickly came to think of NGDLC as being an absolutely brain-dead approach for cable TV systems and here's why: in NGDLC, a defined time slot is given to each user to send his/her data, whether and not that user has any need to send data at that point in time! NGDLC allots the same amount of bandwidth to active users and non-active users. Merely being a potential participant on the network means that you consume network resources. If you were recently active in sending data but don't have anything to send right now, then a packet is still shipped from you; it is just full of null data (i.e., all zeros). Even worse, in this approach, if you have lots of packets to send then you must wait your turn between sending each packet. You only send one packet at a time while everyone else on the network who is inactive (i.e. has nothing to send) sends their packet of nulls, and those who are active send their packet; only then do you get to send just one more packet. Then the process starts over again, synchronously transporting small portions of bandwidth for every user who might or might not have something to send. It is a terrible use of shared network resources! And to make matters even worse –in order to have more bandwidth available for any particular user, that user is assigned multiple time slots, which, when that user is inactive, further ties up network resources that could be more efficiently used by others who actually have something to send.

Luckily, there were a few respondents such as LANcity and Com21 who provided brilliant suggestions for allowing users who needed lots of bandwidth to get it without wasting network resources and without slowing down other users or the network as a whole. It was these brilliant responses to the Telecommunications RFP that really got us thinking about the inner workings of cable modems and would soon spur us to create and issue the High Speed Cable Data Service RFP.

During this time, several of CableLabs members already had cable modem trials and even deployments under way. A group of cable operators banded together to form the CableLabs High Speed Data

[1] "All the usual suspects" was a favorite phrase of Milo Medin who became the Chief Technology Officer of @Home. Another favorite phrase Mr. Medin attributed to NASA was "With enough thrust, anything will fly."

Working Group under the leadership of Bill Bauer of WinDBreak Cable. David Fellows, then CTO of Continental Cablevision, advised us how his company was using LANcity modems in Cambridge Massachusetts; Steve Craddock of Comcast in Philadelphia advised they were using Hybrid Networks modems in Lower Merion, Pennsylvania; Marty Weiss of Cox Communications in Phoenix advised they were also using LANcity modems; Mark Millet of Cox Communications in San Diego advised they were using Zenith modems; George Hart of Rogers Communications in Toronto advised they were using Zenith modems in Newmarket, Ontario; and Doug Semon of Viacom was using Hybrid/Intel/GI modems in Castro Valley, California. I studied and visited these trials and deployments of cable modems and asked each CableLabs member company how CableLabs might be of service to them in their cable modem efforts. The resounding reply heard universally from everybody I asked was something like this: "We are so busy rolling out a particular vendor's modems, [such as Zenith] that we have no time to work with or otherwise understand other vendors' modem solutions [such as LANcity]. Will you [CableLabs] create a sort of 'Consumers Report' showing the strengths and weaknesses, benefits, pros and cons of all the modem systems that are available out there?"

Bang. In that instant, I had a clear enough mission to get started. The way I saw it, the CableLabs members were asking me for expert unbiased product ratings and reviews. This called for a great deal of educating myself on modem performance testing very quickly. This was the kind of theory and practice—theoretical research and practical laboratory work—that I loved to do!

I got things started by asking the three existing modem vendors if I could attend the formal classes in their facilities where they trained their customers on how to deploy and maintain their cable modem system. Rouzbeh Yassini agreed and invited me to a LANcity class in Andover, Massachusetts. Tim Frahm also agreed and invited me to a class at Zenith in Chicago. The folks from Hybrid Networks in California had me visit them as well. Within a couple of weeks, I had a fair understanding about these three modem systems. Soon after, I asked each manufacturer to loan me a complete modem system that I could set up in the laboratory at CableLabs. They all graciously agreed to loan CableLabs a complete modem system of their own manufacture. I had to pinch myself; this was going very well.

There were many burning questions already on my mind that I asked each of the three manufacturers. How did each manufacturer know how well their modem system performed under the best of conditions? How did they know when their modem system was working OK? How did they know when it was sick? How did they quantify its performance? How well did the modem system perform (move traffic) in the face of a little traffic or a lot of traffic? What happened when you presented too much traffic to their modem system? Did it stop working and fall over? And if so, did it recover swiftly and gracefully or was the recovery prolonged and ugly? And how does one debug the system when it stops working?

I couldn't shake these questions. They gripped my mind, would not let go, and sent me full speed ahead into a deep inquisition of modem behavior, performance, and abilities[2] (capability, supportability,

[2] Howard Pfeffer gave me a great appreciation for this lovely encompassing term "abilities" much later. Thank you Howard!

deploy-ability, survivability, etc.). I needed to understand how these products differentiated themselves from each other, and how well each worked in the cable TV operating environment—a real challenge as I knew little about cable TV Hybrid Fiber Coax (HFC)[3] operating environments.

[3] See Walter S. Ciciora's overview of *Cable Television in the United States* where he states that in the early 1990s, the vision of what fiber in the cable plant could accomplish inspired the cable industry to dramatically upgrade its physical facilities.

Chapter Four

The first little thing that changed everything: Theory and Practice on Steroids

On my first visit to LANcity I had the incredible fortune of meeting (for the first time) their founder and CEO, Rouzbeh Yassini. As we dined over breakfast, having known each other only a few minutes, I could already sense the intensity and kindness of this special man. For starters, he was dressed impeccably. He spoke quickly and with great precision. He was kind and put me at ease. I immediately got the sense that he wanted to help me. "What do you want to know?" he asked. I blurted out "I want to know how well your modem system works. And I want to know how you test performance of your modem system. And I want to know how you, your staff and your customers know if your modem system is healthy or not." Just like that, I spilled my guts in an instant, as if talking to the great "Wizard of Oz."

Rouzbeh assured me I would learn these things during my visit, and after the formal classroom training with Gene O'Neil, I would then meet with LANcity's chief technical officer, Gerry White. The day started quickly and the training Gene delivered was full of useful information. Gene was a wonderful instructor and his enthusiasm was evident throughout the full day of training. As thrilled and excited as I was in learning so much about the LANcity modem, I had no idea that in just a few moments I would be graciously given a gift of great value—that I didn't even know existed—yet was hoping with all my heart to find.

It was the spring of 1995 when I first met Gerry White on the fifth floor of 100 Brickstone Square, Andover, Massachusetts. I can take you to the very spot where, looking out the window to the south, one can see the beautiful leafy trees of New England. I couldn't help but feel privileged to be in this revived textile mill and learning so much about new technology. What Gerry had and shared freely with me immediately blew my mind.

Gerry demonstrated for me a very simple, yet powerful cable modem performance test that was already implemented, highly repeatable, and able to run at the push of a single button! I get goose bumps just writing about it. I was in heaven realizing so much extremely relevant and immediately useful theoretical and practical work had already been completed and embodied into this tester. The performance test was based on research by the Internet Engineering Task Force (IETF), Benchmark Methodology Working Group (BMWG), under the leadership of Scott Bradner at Harvard University. The test implemented much of the IETF Request For Comments document 1242[1] (IETF RFC 1242) and ran on a "luggable" customized PC manufactured by Wandel and Goltermann. The moment Gerry demonstrated the first little bit of what the tester could do, I knew it was exactly what I needed.

The Wandel and Goltermann DA-30C was a high performance Data Network Analyzer, used for acceptance and benchmarking testing. It was perfect for my needs. It was 25 pounds, was a no-brainer

---

[1] https://www.rfc-editor.org/rfc/rfc1242.txt

to use and even had a carrying handle! I simply had to have one. As I somewhat expected at the time, this little tester put "cable modem testing" on the map and helped change the course of history for cable modems.

I was so excited that I could hardly sleep. Where was I going to get one of these testers? What would the DA-30C tests reveal about the performance of modem systems from various manufacturers? I was pretty sure it would reveal a lot—and it did!

The man who helped me get a DA-30C at CableLabs was Gary Archuletta, our local sales representative from Wandel and Goltermann. I didn't have any money to buy one or rent one, and to get money I needed to put together a project plan to be approved by the CableLabs project committee. That was going to take some time. But, Gary understood and made sure I had this critical piece of test equipment right away, before I had a purchase order or could make any commitment. It was with this loaned machine that we were able to run the tests that surprised everybody!

The Wandel and Goltermann DA-30C had two Ethernet interfaces on the back of the unit. Internet traffic was sent out of one interface and received in the other. The simplest test setup was to attach a short Ethernet cable from the one interface to the other to make sure that the tester was working properly.

Oh, the things this little box could do! To invoke the test, one would type RTBENCH at the DOS command prompt. The tester would then automatically start sending traffic out the first port and look to see if that traffic appeared on the second port. The traffic was in the form of variable length Ethernet frames. Computer programmers refer to the functionality embodied in this test as a "nested loop." Here is how it worked: To begin with, the tester sent short Ethernet frames, initially very slowly, and then faster and faster until the maximum frame rate had been achieved; next, having sent all possible frame rates for small frames, the tester began to send medium sized frames, again slowly at first. After all possible frame rates for medium-sized frames had been sent, the tester began to send large-sized Ethernet frames. And it did all this in practically no time, with just the push of one button!

This box was a practical tester's dream. The theory behind how the tests were constructed and how to explain the test results was straightforward and easy to grasp.

In addition to the practical laboratory testing that this little box enabled us to perform, there were other completely different theoretical exercises that were follow-on activities to the Telecommunications RFP, namely the Multimedia Communications Network Request For Information and Integrated Development Proposals (MCN RFI/IDP), and the High Speed Cable Data Services (HSCDS) RFP.

Wandel and Goltermann DA-30C Internetwork Analyzer

Chapter Five

The High Speed Cable Data Service RFP and Hell Weeks 1 and 2

A wonderfully resourceful woman, Dr. Masuma Ahmed, wrote the High Speed Cable Data Service RFP (HSCDS). Behind mostly closed doors, she worked relentlessly on the document and delivered, as the French would say, a tour de force. When she shared the document for the first time, it was by far the most complete draft RFP I had ever seen from a single author. It was very clear that Masuma knew what she wanted, and though a few of us gave her detailed comments, she went on to finalize this incredibly comprehensive document mostly on her own.

As of the publication date, no other RFPs had ever documented the Cable TV HFC and RF requirements in such great detail (especially providing the ranges for each RF parameter based on several CATV network requirements). The HSCDS RFP was the first industry document to do so and it took a lot of initiative and collaboration from the MSOs to develop the requirements. The requirements became the foundation of the RF standards initiative at the IEEE and later became the baseline for the MCNS and DOCSIS RF specifications for HFC networks. Masuma and Mario Vecchi from Time Warner Cable, in collaboration with vendors General Instrument, Scientific Atlanta, LANcity, Motorola, and others, created the first DOCSIS RF MIB (Management Information Base). When starting the Internet Protocol Cable Data Network (IPCDN) Working Group at the IETF, as the chair of the group, Masuma submitted the MIB and other documentation that were included as items for consideration and discussion for IETF standards initiatives.

The RFP was a collaborative effort of CableLabs and North American MSOs, and Masuma led the effort. MSOs participating in the RFP included: Tele-Communications Incorporated, Time Warner Cable, Rogers Cablesystems, Jones Intercable, Continental Cablevision, Comcast, Shaw, and Videotron (please see RFP for compete list). The RFP was not specific to one MSO network and included requirements for several different types of HFC networks. In addition, several of the MSOs provided various infrastructure architectures such as; 1,000 vs. 300 homes passed per Fiber Node, Distribution Hub vs. direct Fiber Node connectivity to the headend, varying number of amplifiers per serving area, number of homes per drop, coax vs. fiber-deep, various RF requirements, OSS/BSS requirements, etc. The companies who provided extensive data included Rogers Cable, Time Warner Cable, Tele-Communications Incorporated, Continental Cablevision, Shaw, and Videotron.

Responses were due June 30, 1995. We received responses from 50 companies in all, including; 3com, ADC Telecommunications, Akzo Nobel Electronic Products, Alcatel Cable Contracting NA, Applied Digital Access, AT&T, Augat Communications, Aydin Telecom Division, Bay Networks, BBN (Bolt, Beranek, Newman), Cisco Systems, Com21, ComPath, ComStream Corporation, DataEquip, Digital Equipment Corporation, Ericsson Network Systems, First Pacific Networks, Fujitsu Network Switching of America, General DataComm, General Instrument, GTE Government Systems, Hewlett Packard, Hybrid Networks, IBM, Integrated Network Corporation, Intel, LANcity, Loral Data Systems, MessagePhone, Microunity

Systems Engineering, Motorola Multimedia Group, Net Labs, Northern Telecom, Nuco Information Systems, Objective Systems Integrators, Perot Systems, Phillips Broadband Networks, Probita, Scientific Atlanta, Siemens Stromberg-Carlson, Stratus Computer, TAG International, Telegate, Tellabs Operations, Terayon, Thompson Sun Interactive Alliance, TrackCom Systems, Unisys, and Zenith Electronics.

It looked like a moving company had arrived. There were huge and heavy boxes stacked in our offices and in the hallways … everywhere. Each reviewer had somewhere around 70 three-ring binders that each were two-inches thick. There were about fifteen copies submitted from each manufacturer—all in, there were over a thousand binders to review! Needless to say, we had our work cut out for us, and the theoretical work of digesting this information went on for many months. My set of 70 binders covering the responses from each of the 50 manufacturers were set on my office floor for easy access and went completely around the perimeter of my office.

To assist in the evaluation effort, I constructed an evaluation matrix using Microsoft Excel on a Macintosh computer. It had a column for each of the 50 respondents, and several hundred rows consisting of the review criteria.

The review criteria included a section on the architecture overview along with proposed subsystems such as whether the respondent was proposing a complete end-to-end system, whether a telephone line could/would be used for the return path, whether the modem "to and from" user "data rates" were asymmetric or symmetric, whether the modem system was completely distributed, or if instead there was a Headend controller. Here is a chuckle for those in the cable industry– at the time we referred to the headend controller as a CDMTS (Cable Data Modem Termination System), a term we later shortened to CMTS; both terms were never liked by marketers!

The Evaluation Matrix also had a short section titled "Reviewers Perspective on Critical Issues" which included these items: 1) Traffic Control (fairness and optimality, ability to enable/disable transmission of data), 2) Bandwidth Efficiency (of huge importance to me and others), 3) Simplicity/Economy of operation, 4) Time to Market, 5) Security (tamper proof abilities), 6) Provisions for Operations Support Systems (OSS) and Business Support Systems (BSS), 7) Coexistence with other services on the network, and 8) Remote Diagnostics (loopback, ability to ping a modem, see subscribers equipment, etc.).

Lastly, the Evaluation Matrix included an enormous section of the vendors' self-reported scores to Masuma's many detailed questions. There was an Overall Functional Requirements section that included service features, security features and performance and availability objectives such as packet loss ratio, packet delay, and inter-arrival time (i.e., jitter, aka packet delay variation). There was an Architectural Functional Requirements section which included end-to-end network architecture and functional requirements such as support for industry standard interfaces in wide and local area networks, compatibility with existing cable network designs that considered numbers of homes passed per node (500-200) and number of nodes per CDMTS (40-200), return path functional requirements such as resiliency against upstream ingress and impulse noise and ability to use a telephone return path, forward path functional requirements such as RF transmission parameters, and channel transmission failures. The final section of the vendors self-reported scores were around systems and interfaces functional requirements including the use of standard Internet protocol (IP), sub-netting IP, network management interfaces and protocols, the cable data modem interfaces, addressing, powering, reliability, physical security, plus the RF interfaces and protocols.

In keeping in line with the original "Consumer Reports" request, I developed a method for reviewers to visually assess this sea of information by representing all responses on a single poster-size page. An improvement to readability was a shaded-circle font of just three characters that indicated how well a vendor covered every question in the HSCDS RFP. In my scheme, one of three different shaded circles represented the compliance, coverage and/or level of treatment for each question: 1) An open circle was a rating of none/minimal/poor, 2) a half-shaded circle was a rating of some/fair, and 3) a fully shaded circle meant full compliance/coverage/level of treatment. Evaluating the High Speed Cable Data Service RFP was exciting and the shaded circles greatly assisted reviewers in assimilating all the different vendors' proposed capabilities. Now that we could see all that was proposed in a succinct manner, the challenge was determining what to do next.

In evaluating responses to the RFP, we thought long and hard about what each vendor proposed, what each already had in the way of working products, how far along each was in their development of new or next generation products, and most importantly, the "fit" of each proposed solution within what the cable industry actually needed. This was an extremely tedious, iterative and repetitive process; largely because the responses were, little by little, helping us uncover and create a better understanding in our minds eye of what it was that the cable industry actually needed for the next millennium.

A final step in evaluating the RFP was having each vendor present their solution for approximately one hour to an audience of CableLabs and MSO staff. With 50 respondents, this process went on for almost two weeks with catered breakfasts, lunches and dinners. It became known as "Hell Week".

During almost every one of the 50 or so presentations given by vendors, a rich in-depth technical dialogue ensued. Sometimes, the dialogue consisted of gentle thought-provoking suggestions from vendors and the audience, but other times the dialogue was downright combative. Sometimes big brains were graceful; sometimes they clashed.

Here is a related example: The interface from the cable modem to the personal computer was suggested by most to be Ethernet 10BaseT, though some companies, including Scientific-Atlanta, were adamant about the eventual domination of ATM25 displacing Ethernet in the home. ATM25 is an ATM (Asynchronous Transfer Mode) version wherein data is transferred at 25.6 megabits per second. Though 100 megabits per second Ethernet networks had yet to arrive, which would have made the decision a whole lot simpler, we collectively shook our heads "no" and put a bullet into the idea of ATM25 right then and there.

The tussle around ATM didn't stop at the ATM25 interface. There were vendors such as Com21 who felt that ATM should be transported (in native mode) through the entire network to/from customers' homes and cable TV headends. This started a long and rich discussion that went on for weeks, even months. The question at hand for years was whether access and transport networks throughout the world would evolve to fixed length cells, namely the 53 byte cells used in ATM, or instead would use variable length IP packets.

In any case, we knew that future modem systems should be able to support both modes, so the question became, should: 1) ATM be the underlying native transport with additional support for variable length IP packets, or 2) should IP be the native transport with support for ATM cells? The outcome of the decision would have enormous bearing on the design of a cable modem system, and as such, we gave it a lot of thought. In the end, we collectively decided to use IP as the native transport and left hooks in place to be able to support ATM. We were clearly placing a network efficiency bet and we all knew it. If IP ended up being ubiquitous, then we would win the bet. If ATM ended up being ubiquitous, we would then lose due to greater transport inefficiencies. Even if our bet was wrong, we still could transport ATM and make all services to subscribers work just fine.

There was another thought provoking area that came out of RFP evaluations and Hell Weeks 1 & 2, namely, what services were really going to be prevalent as we moved to the next millennium? We are still thinking about this question!

Evaluating the HSCDS RFP was an intense theoretical exercise. There was so much information and our hours were long. With breakfast, lunch, and dinner inside and only a few minutes between vendor presentations, we rarely saw the sun. I guess that's why we called it "Hell Week."

On the plus side, never before had we created an environment that resulted in such a rich dialogue with so many experts on cable modem systems. Our two Hell Weeks were the largest collection of prominent cable modem researchers and users in the world, and we were able to kindle long and fruitful relationships with some of the greatest telecommunications minds on the planet.

All the theoretical work was starting to hurt my head and I yearned to get started on practical laboratory testing of the existing modem systems from LANcity, Zenith, and Hybrid Networks.

Chapter Six

"Do you guys know what you're doing?"

A welcome complement to the intense theoretical work of evaluating the High Speed Cable Data Service RFP came in the form of practical cable modem testing in the laboratory. To evaluate all existing cable modem technology in a "Consumer Reports" style, I knew I was going to need serious help and thankfully, I knew just the capable engineers for the job. In April 1995 as the RFP was being readied for publishing, I reached out to two friends at the University of Colorado, Tom Moore and Brian Reilly, who immediately agreed to come on board as CableLabs interns for the summer. The very first multi-vendor proprietary cable modem testers are shown below with Bob Cruickshank at left, Brian Reilly in the middle and Tom Moore at right.

Dashing in and out of the laboratories, these guys were immediately loved by CableLabs staff. Someone even made them lapel buttons that said "I'm Tom and he's Brian" and vice versa. And the guys would switch them just for fun to keep everybody guessing. Young and capable, few knew of their prestigious backgrounds. Tom Moore holds a master of engineering degree in telecommunications (with honors) from the University of Colorado and a master of business administration degree (with honors) from the Harvard Business School. Prior to joining CableLabs, Tom spent five years as president and general manager of the Natural Fuel subsidiary of Public Service Company of Colorado. Brian Reilly spent 10 years with GTE Government Systems Corporation where he was a member of technical staff. His work there included network design of ATM/SONET systems, circuit and system design of computing hardware, and ASIC design. He received his bachelor of science degree in electrical engineering from The Johns Hopkins University and a master' degree in computer science from the University of Colorado.

I quickly passed along to Brian and Tom what I had learned during training at LANcity, Zenith, and Hybrid Networks. Having completed all their available training programs, the cable modem manufacturers were each very gracious in helping us install and configure their respective modem systems in the laboratory at CableLabs.

By mid-May, 1995, each cable modem system was configured as a standalone network and represented, albeit very simply, a tabletop version of a cable television Internet access network. There was a headend modem from which Ethernet frames were injected and then sent down the network to customers. There was a short piece of coaxial cable that acted as—though grossly simplified—the HFC network that you would find in any neighborhood. And there was a customer modem that received the Ethernet frames.

The entire tabletop modem system from headend to customer was considered a single device under test by the Wandel & Goltermann DA-30C, which would send Ethernet frames to the headend modem and then look to see if all those frames were received by the customer modem (and vice versa).

In very short order, measured in days not weeks, we had everything we needed to run and publish our first "Consumer Reports" modem tests. We had the three modem systems. We had the Wandel & Goltermann DA-30C, and we had two amazingly gifted engineers, Tom and Brian.

Needless to say, the three modem systems distinguished themselves from each other a lot. A whole lot … we were shocked!

We reran the tests again and again. Something must be wrong, we thought. We went back and studied RFC 1242. We called RFC 1242 author Scott Bradner at Harvard and started making plans to visit his lab. We called Wandel & Goltermann. We reran the tests again and again. We showed some of our results around CableLabs and were told that we better share our laboratory findings with the CableLabs member companies. But what did we know? We were brand new to cable and were not the CableLabs engineering department. It was just a couple of interns and me.

We downplayed our findings as much as we could. We referred to our work as "Unimpaired Modem Testing", because it was performed in a pristine radio frequency (RF) environment without the impairments experienced by real life cable systems that operate in the presence of sunspots, AM, FM and other transmitters such as cell phones and CB radios, electromagnetic interference from lightning, and countless other RF ingresses. The tabletop versions of cable systems that we had created in the laboratory reduced an entire cable network, which in real life could span thousands of miles, to just a few feet of cable that you could hold with just two fingers. And as such, we knew we had simplified things a lot and tried to make the simplification apparent in the description of our findings.

What did we find out? Well, of the three modem systems, the LANcity modem moved the most traffic successfully—by far. We were able to present the data using graphs generated in Microsoft Excel that showed on the x-axis the offered data load and on the y-axis the forwarded data load. This was not rocket science. The DA-30C created the offered load that was sent to the headend modem, and in turn, down the tabletop cable network and then the DA-30C measured the forwarded load that it received from the customer modem. For each frame size, small, medium, and large, the DA-30C created an output file of two columns: 1) offered load and 2) forwarded load. In an ideal world, the forwarded load would always equal the offered load and if the Device Under Test (DUT) was just an Ethernet cable, then the graph of offered load versus forwarded load would be a 45° line showing that all traffic that was offered was indeed forwarded. And sure enough, when the DUT was an Ethernet cable the test results were as expected. However, when the DUT was a cable modem system, the results were, let's say, interesting.

Management told me "You guys are going to get a call at 1:00 PM this afternoon. Be prepared to defend your results." I went down to the lab to find Tom and Brian so we could prepare for the call. What are we going to say … that the whole test took us three minutes?

I think it was Brad Dusto and Steve Craddock that called from Comcast that day. They had received the graphs we had made that showed offered load versus forwarded load for the three modem systems LANcity, Zenith, and Hybrid Networks. Previously, most of us had little if any interaction with these gentlemen and I don't think had ever spoken with them on the phone. I sensed they were as shocked as we were with the results.

"Do you guys know what you're doing?", they asked. "We think so", we said. "Are you sure you know what you are doing?", they asked again(perhaps a little frustrated). We described our methodology and said, "We stand by our test methodologies and we stand by our numbers."

Chapter Seven

Benchmarking Methodology, the IETF, and the W&G RTBench tests

The pioneering work of Scott Bradner and the Internet Engineering Task Force, Benchmark Methodology Working Group (BMWG) allowed us to forge ahead and evaluate modem systems with incredible speed and accuracy. During the summer of 1995, we were able to accomplish a great deal very quickly because back in 1991 the Network Working Group published Benchmark Terminology for Network Interconnection Devices. As I understand it, the motivation behind Scott's work was to allow anyone interested in procuring networking hardware to perform an "apples to apples" comparison of available solutions. The following quote from the Introduction section of RFC 1242 provides insight:

Vendors often engage in "specsmanship" in an attempt to give their products a better position in the marketplace. This usually involves much "smoke & mirrors" used to confuse the user. This memo and follow-up memos attempt to define a specific set of terminology and tests that vendors can use to measure and report the performance characteristics of network devices. This will provide the user comparable data from different vendors with which to evaluate these devices.

The work of Scott and the BMWG allowed us to produce results immediately.

The first test described in RFC 1242 is a Back-to-Back buffer capacity forwarding test where "Fixed length frames [are] presented at a rate such that there is the minimum legal separation for a given medium [i.e. minimum separation on the network wire] between frames over a short to medium period of time, starting from an idle state." This was important because of the growing number of locations/devices on a network that can produce bursts of back-to-back frames. Today's peer-to-peer applications fall squarely into this category and create large packets that must be fragmented to be sent via the cable modem. "Since fragment reassembly will only be attempted if all fragments have been received, the loss of even one fragment because of the failure of some intermediate network device to process enough continuous frames can cause an endless loop as the sender repetitively attempts to send its large data block."

The back-to-back capability of LANcity very closely resembled the performance of a wired Ethernet connection across the full range of frame sizes. At medium to large frame sizes all three vendors' modems performed as expected. For small frame sizes, the Hybrid system came in last place as it could only manage 30-45 back-to-back frames; the Zenith system came in second place with a limit of about 160-190 back-to-back frames. The LANcity system was able to forward the maximum number provided by the RTBench test, about 450 frames.

The second test described in RFC 1242 is Frame Loss Rate, which measures the "Percentage of frames that should have been forwarded by a network device under steady state (constant) load that were not forwarded due to lack of resources. This measurement can be used in reporting the performance of a network device in an overloaded state. This can be a useful indication of how a device would perform under pathological network conditions …"

For small (64 Byte) frames, forwarded load versus offered load was quite different among the three modem systems. The LANcity system, rated by the manufacturer at 10 megabits per second, performed most like a wired Ethernet connection forwarding all offered load up to about seven megabits per second, which is the theoretical maximum for 64 byte frames. The Zenith system, rated by the manufacturer at four megabits per second, came in second place but was only able to consistently forward about one megabit per second. The Hybrid Networks system, rated by the manufacturer at 10 megabits per second, came in last place forwarding less than one megabit per second—and exhibited degraded performance as offered load was increased to put the modem in an overloaded condition.

For medium (768 Byte) frames, Hybrid and Zenith forwarded load versus offered load improved considerably though LANcity was still the clear winner. Zenith improved to slightly over three megabits per second. Hybrid improved greatly to just over five megabits per second. LANcity achieved almost nine megabits per second.

For large (1518 Byte) frames, LANcity again came in first place forwarding slightly over nine megabits per second. Hybrid came in second place forwarding approximately seven megabits per second. Zenith came in third place and started to show overloaded behavior, falling back to a reduced forwarding rate of about two to two and a half megabits per second.

The third test described in RFC 1242 is Latency, which for "bit forwarding devices" measures "The time interval starting when the end of the first bit of the input frame reaches the input port and ending when the start of the first bit of the output frame is seen on the output port." We knew variability of latency was a real concern and could be problematic, especially in voice networks and real-time applications such as streaming video, online gaming, etc.

When measured in milliseconds, a three-foot Ethernet connection has nearly zero frame latency. LANcity performed best with latency never exceeding one and a half milliseconds across the full range of frame sizes. Zenith and Hybrid had about the same performance, ranging from one millisecond for small frame sizes to about three and a half milliseconds for large frame sizes.

The fourth test described in RFC 1242 is Throughput, defined as "The maximum rate at which none of the offered frames are dropped by the device. The throughput test allows vendors to quantify and report a single value that has proven to have use in the marketplace. Since even the loss of one frame in a data stream can cause significant delays while waiting for the higher level protocols to time out, it is useful to know the actual maximum data rate that the device can support."

The throughput test yielded dramatic results. Again, LANcity most closely followed the performance of a wired Ethernet connection. For small (64 Byte) frame sizes, LANcity was able to forward 14,000 frames per second; by comparison, Zenith and Hybrid could only successfully forward about 2000 frames per second.

We learned so much during these one-button tests and our respect and appreciation for Scott Bradner and the IETF Benchmark Methodology Working Group grew each day. Other items in RFC 1242 which helped educate us included the concept of Constant Loading, which specifies that fixed length frames be transmitted at a fixed rate. Though constant loading (i.e. Constant Bit Rate or CBR) is an extremely rare occurrence in practical networks, it does represent a steady state condition which can be easily and universally applied in many different settings (of frame size and frame rate) across many different vendors; in our case, with quite surprising results as shown above. Endeavoring to be good engineers, we continuously worked to get away from CBR loading, instead seeking to understand and use real world traffic profiles whenever possible. This was not easy, but we persisted and made a big breakthrough later when developing a virtual model that later became known as the Common Simulation Framework.

Other items in RFC 1242 of great interest to us included: link speed mismatch, inter frame gap, MTU-mismatch behavior, overhead behavior, overloaded behavior (already apparent in some of our test results), policy based filtering, restart behavior, and single frame behavior.

We called Scott Bradner, described our cable modem testing efforts, and sought his guidance. He was gracious with his time and even hosted our visit to his facility at Harvard University in Cambridge, Massachusetts.

We were very excited the day we got to meet Scott Bradner! As chance would have it, Hertz gave Tom and me a huge Lincoln Town Car so we had some fun getting to and through Cambridge in the evening rush hour.

Over the years, Scott was a mentor in describing what he thought was necessary from a user interface and “dashboard” perspective, so that would-be testers/users could readily decipher our cable modem performance measurements.

Chapter Eight

Building an Internet "Mini Me" to Characterize Network and User Behavior using World Wide Web, E-mail, and FTP

Given that our benchmarking tests took all of three minutes to run and were relatively easy to understand, we figured we'd better spend some time brushing up on other areas of testing data networks. During a quick scan of the June 1995 issue of Network Computing magazine, we found reviews of other practical network benchmarking tests such as Netbench and Servbench from Ziff Davis, WebStone from Silicon Graphics, and Netperf from Hewlett Packard. From what we could tell, our choice of Wandel and Goltermann's DA-30C and RTBench software was very appropriate for what we were trying to do.

But, something else really caught our eye! The Network General Company made a product called a Sniffer that performed detailed higher layer protocol analysis combined with strong capture and decode capabilities. We quickly realized we could use the Network General Sniffer in characterizing user and network behavior, while simultaneously having a look at traffic from the World Wide Web, E-mail, and FTP (file transfer protocol).

Though the World Wide Web was in its infancy, we knew the web was going to have incredibly widespread use in the future. As such, we decided to think about how we could/should characterize web traffic. The existing online services America Online, Prodigy, and CompuServe all were very popular and we knew it was worthwhile having a good look at the traffic from these as well. Given that the most popular method at the time for sending large files through networks was FTP, we decided to have a look at FTP traffic, too.

Our work got really fun very quickly. We created an Application Testing Setup that simulated five different networks: the three different cable modem networks (Zenith, LANcity and Hybrid Networks), a 14.4 kilobits per second[1] dial-up telephone modem-based network, and a short run of 10BaseT cable that made an Ethernet hard wire network. This is how "Mini Me" was architected[2]. We had so much fun thinking about how we should characterize user and network behavior in the presence of a realistic traffic!

How did we get real traffic and make Mini Me emulate reality? We choreographed a World Wide Web "surf" that started on the CableLabs home page, then went to the CableLabs Internet Resources Section, then went to the Rogers Cablesystems home page in Canada, then went back to the CableLabs Internet

[1] Yes, 14.4 Kbps was the highest speed telephone modem generally available and in use at the time.

[2] Back in the day, we simply called it the lab network, though Austin Powers' popularized term "Mini Me" fits well today.

Resources Section, then went to the Tele-Communications Inc. (TCI) home page, then to @Home's home page, and finally to a picture of Longs Peak, Colorado.

We then started performing this surf on the real Internet (just as one would today by clicking on various HTML links while web surfing) and relied heavily on the Network General Sniffer to capture every click, html web page, all JPEG files, TCP acknowledgements, etc. The Sniffer caught and time-stamped everything and was able to dump the results to a file so we could quantify and even recreate all the bits/Bytes/packets of data that had passed. With some much appreciated assistance from Andy White and Jerry Stachowski, we then loaded all the source files onto our own CableLabs server and performed a few Internet Protocol (IP) routing tricks needed to recreate a user's web surfing experience in a completely controlled environment. This is how "Mini Me" was born, and it worked just great! One could "surf" on Mini Me and get the exact same experience as surfing on the Net, except Mini Me performed exactly the same each time whereas the Net would perform differently every time due to a whole host of network and server performance anomalies that were beyond our control.

The World Wide Web surf took about 107 seconds and passed a cumulative total of 460 kilobytes of data in the downstream direction to the home. In the upstream, the cumulative total traffic was about 65 kilobytes. Once scripted, this surf was performed on our Mini Me network for each of the three cable modem networks, the telephone dialup network, and the Ethernet hard wire network.

As expected, the LANcity network performed the best having times very close to the hard wire Ethernet network. The time to complete the surf over the LANcity network was 108 seconds. The time to complete the serve over the Zenith network was 112 seconds. The time to complete the surf over the Hybrid Networks network was 125 seconds. And the time to complete the surf over the dial-up network was almost two and a half times as long—a whopping 282 seconds!

We immediately and over the long term, got a lot of mileage out of the Mini Me network concept. Looking to the future we knew we needed richer graphical content and were in constant pursuit of more realistic graphics-rich surfing experiences that we would capture with the Sniffer and move on to the Mini Me network.

From a research and testing perspective, the beauty of the Mini Me network was our ability to decouple from the Internet and control as many aspects of the surfing experience as possible including: the selection of the web surfing personal computer[3], it's operating system[4] and web browser; and the

[3] We were in heaven when Intel Corporation gave us two brand new PCs with Pentium 90 MHz processors and Microsoft Windows 95 (which forever removed TCP/IP PC bottlenecks). These machines had a TCP/IP throughput in excess of four megabits per second.

[4] Early versions of Microsoft Windows such as Windows 3.1 and Windows for Workgroups 3.11 could not access the Internet without a TCP/IP "shim" or "stack" from a third party such as Chameleon; tests showed that Chameleon 4.0 on a Intel 486 processor 33MHz PC had very poor TCP/IP data throughput—about 200 kilobits per second.

selection of the "far end" content server, selection of the router, DNS server, DHCP server, etc. In making these selections, we were able to control the origination, delivery, and rendering of web pages.

As much as possible, we performed similar activities in capturing and analyzing AOL user sessions and FTP user sessions—and recreating them on the Mini Me captive network.

Mini Me quickly took on a life of its own and under Andy White's capable supervision even made a trip to Washington, D.C. for the National Cable Television Association (NCTA) Technology Tour in which CableLabs presented a side-by-side demonstration illustrating the speed and power of cable television technology for data transmission. Here is a quote about Mini Me from that time:

The CableLabs demonstration was featured as part of the opening presentation for each of eight tour groups over two days. Visitors included members of Congress and their staff advisors, as well as high-level FCC and other Administration officials. A total of 300 people attended the event.

NCTA's President and CEO Decker Anstrom delivered opening remarks, and then drew the audience's attention to the CableLabs demonstration. The CableLabs' demo showed three computers connected to a remote information server. Each computer was connected by a different technology:

One computer was connected via a telephone modem running at a speed of 14.4 kbps over regular phone lines. This is widely available technology now being used by some consumers to access the Internet and online services. Most consumers still utilize 9.6 kbps or slower connections.

The second computer was connected via integrated services digital network (ISDN), a phone company-provided service running at a speed of 56 kbps over regular phone lines. ISDN is not yet widely available to residential consumers.

The third computer was connected via a cable modem, running over coaxial cable at a speed of 10 Mbps. Cable data transmission is 1,000 times faster than most existing telephone modems.

Andy White, CableLabs' senior information systems specialist, started all three computers on a retrieval of a complex, digitized photograph of the Space Shuttle cockpit, a very large file from NASA's Internet site. The speed and efficiency of the cable connection vividly drove home the advantage of cable television technology in delivering practical, high-speed data applications to homes, schools, and businesses. Guests remarked that cable's advantages had not been clear until they saw this live demonstration.[5]

Mini Me was a "star" and clearly demonstrated that the cable television networks that passed 97% of the property lines in North America had tremendous bandwidth that allowed for transmission of many two-way services in addition to (plain old one-way) video.

---

[5] SPECS News From CableLabs, September/October 1995, Volume 7, Number 7.

Chapter Nine

What We Had Learned thus Far and a Look to The Future: What Does a Cable Modem Need to Do Anyway?

In a very short period of time, people looked to us to be the resident experts on the strengths and weaknesses of the three existing vendors' cable modems when operating in an unimpaired cable TV RF environment. We had learned many insightful distinguishing characteristics among the modems from Zenith, LANcity, and Hybrid Networks and were on a quest to put these insights to good use.

A question for the reader: Can you see enough writing on the wall to tell where we have arrived and where we're headed? Here are a few more details ...

The Zenith and LANcity modem systems had no "master modem" and operated autonomously as peers on the network with minimal knowledge of each other. Both modem systems offered symmetric data rates to and from a set of users. In the LANcity system, QPSK (Quadrature Phase Shift Keying) modulation was used in a 6 MHz channel resulting in a user data rate of about ten megabits per second—just about as fast as Ethernet networks at the time. One 6 MHz channel transported data to users and one 6 MHz channel transported data from users. A frequency translator in the headend coupled the upstream and downstream frequencies so all the individual modems could hear when each other were transmitting and could see what data was destined for each. The LANcity System acted as an Ethernet bridge meaning that individual modems only forwarded Ethernet frames destined for addresses not appearing on the originating segment. Think of a mail carrier who takes an outgoing letter from your mailbox that is addressed to your next-door neighbor. Instead of bringing the letter back to the post office, the mail carrier instead immediately puts the letter in your neighbor's mailbox. Ethernet bridges act like the mail carrier; they only forward frames to other networks if the frames cannot be delivered locally. They do this by building a "bridging" table as they observe traffic on a network.

The Zenith modem system used BPSK (Binary Phase Shift Keying) modulation. Two symmetric data rates were offered: one with a 6 MHz channel that delivered four megabits per second and another with a 1 MHz channel that delivered 500 kilobits per second shared among all users. The Zenith modems acted as simple repeaters meaning that all traffic was repeated across each modem, even traffic destined for other endpoints on the originating segment. By design, more traffic was repeated on a Zenith system then was bridged on a LANcity system, an inherent inefficiency drawback of Zenith.

The Hybrid Networks modem system was quite different and had at least two distinguishing features. 1) There was a master modem referred to as the "Pop" that governed all client modems, and 2) the "to and from" user data rates were asymmetric (unequal in speed). In the forward direction (i.e., transporting data to users), 4 VSB (Vestigial Side Band) modulation was used in a 6 MHz channel resulting in a potential user data rate of 10 megabits per second, though as mentioned previously, the highest data rate we observed with W&G's RTBench was about seven megabits per second with large frame sizes and as little as a half a megabit per second at small frame sizes. In the return direction (transporting data from users), two options were supported: 1) a dial-up telephone modem operating at 9.6 kilobits per second for cable TV plants that were only one-way capable, or 2) for two-way capable cable TV plants, FSK (Frequency Shift Keying) modulation in a 200 kilohertz cable TV channel resulting in a user data rate of 19.2 kilobits per second. From what we could tell, the Hybrid Networks modem system acted as an Ethernet bridge with some limited IP Routing capabilities.

We had a lot of information to digest. The Zenith, LANcity and Hybrid Networks modem systems were so different from one another that we were often given cause to reflect on the goals of any cable modem system design. These three modem systems had many significant distinguishing characteristics. Data rates to and from users were different. Modulation types were different. The occupied cable TV spectrum, that is the RF bandwidth used, 6 MHz versus 2 MHz, versus 200 kilohertz, was different. And there were many other distinguishing characteristics that we had yet to explore or even think about.

We were almost blind to the physical layer RF operating requirements of the modem systems, meaning what each system needed from the cable TV plant in order to operate. What were acceptable values for Carrier to Noise Ratio (CNR) in the forward and return paths, transmit and receive power levels, amplifier “tilt”, group delay levels, and so forth?

What were the distance limitations? Our tabletop cable TV plant was only three feet long. Could modems be located tens or hundreds of miles apart from each other? From our industry’s perspective, what would be acceptable values for all these parameters?

And what about security, encryption and network management? We knew any one of these categories were huge areas worthy of much research in and of themselves.

What about two-way versus one-way cable TV systems? Were there either short- or long-term roles for the “telephone return” option? Could/should a wireless “PCS return” be planned for use later? How long would it take cable TV operators to upgrade all of their hundreds of thousands of miles of HFC plant to two-way, so that the cable return option could be used instead of the telephone return option?

Despite all the open questions, we had made significant progress in a record amount of time. All three existing cable modem systems and much of the necessary data communication test equipment were procured, installed, and functioning in our labs. We had attended onsite training with the modem vendors and engaged in much productive research. We had recreated an ideal Internet access cable plant (a Mini Me version of the Internet) that helped us explore user and network behavior of the World Wide Web, FTP, and online services such as AOL, Prodigy, and CompuServe. We had performed an initial IETF RFC 1242 benchmark on each of the modem systems—which showed they all outperformed telephone dial-up modems by as much as a factor of 1,000 (think megabits per second vs. kilobits per second). We had determined that cable modems delivered an enormous boost in communications power for cable TV users, and that this boost in communications power was due to the “always on” nature of cable modems and their potentially blazing speed. Whereas users had become accustomed to telephone dial-up modem speeds doubling every few years (e.g., 300, 600, 1200, 2400, 4800, 9600 bits per second), cable overnight had a 1,000 times speed advantage.

Despite the entire positive buzz around cable modems, one of the things we knew for sure was that none of the existing modems from Zenith, LANcity or Hybrid Networks could do telephony (the generic name for telephone services). Amazing as these modems were, including all the incredible stories[1] from MSO user groups, we knew that lack of support for telephony was a long-term showstopper for the industry. For the next millennium, cable modems needed to support voice, video, and data—and we didn't see any designs that did a really great job addressing this need!

[1] Quotes from early users of cable modems:
"You will have to pry my modem from my cold dead hands",
"Please start charging for my trial of high-speed data service so I will know I can keep my modem forever", and
"It [my modem] has changed my life!"
Many thanks to Doug Semon and others for these quotes!

Chapter Ten

Modeling for MSOs, Evaluating Many New Entrants and MCNS

I remember exactly where I was standing in the basement of my home in Lafayette, Colorado, on a Saturday morning when Tony Werner, then CTO of Tele-Communications Incorporated (TCI), called wanting to discuss the question, “just exactly how many modems/users can we put on an upstream or downstream?” We talked it over for a bit and it seemed that another way to frame the question was/is, “How many users can be added before there is a noticeable slowdown in service?”

Luckily, my dear friend and capable engineer, Tom Moore, had been thinking about this for a while and I ended my conversation with Tony saying we needed to circle back with Tom. Tony’s question was right on the money and soon most MSOs (Cable TV Multiple System Operators) were asking us to help with this type of capacity planning for cable modem-based high-speed data services.

Well, the answer was/is it depends on a whole lot of variables, some within, and some beyond our control and understanding. The beyond our control part we’d have to live with. The beyond our understanding part gnawed at us and we were on a mission to figure out all that we could. Certainly, it depended on exactly what users were doing such as email, web surfing, and file transfers such as FTP (and today P2P and P4P). And it depended on how often users wanted to do whatever they do and on how long user sessions lasted. It depended on how graphically-rich, or music-rich, or video-rich the content was/is. These questions and answers are timeless—and are still being asked and answered by capacity planning engineers today and will be in the future. Needless to say, we thought hard about traffic and eventually upgraded and refined our captive World Wide Web surf that we had running on Internet Mini Me. With these refinements, over time, our captive World Wide Web surf affectionately became known as the “60-second surf.”

Around us the Internet world was changing at an amazing pace[1]. As of February 1995, the Wall Street Journal reported growth of 3.8 million households with personal computers in the preceding six months (from July to December); 58% of U.S. homes had a computer two years or less and had a CD-ROM and multimedia software; and by the year 2000, 70% of homes were projected to have a personal computer. It was becoming clear that work-at-home, education, shopping, games, financial, and other services were a better match to the PC than the TV. In addition, 56% of customers were willing to pay the same or more than cable TV for electronic information services. Note this phrase “electronic information services”… few had yet to call this explosion The ‘Net or The World Wide Web! Business Week magazine picked up on the Web relatively early and on February 27, 1995 reported that the Web had 27,000 sites and was doubling every 53 days. Fifteen thousand people per day were accessing the Web

[1] Special thanks to Mario P. Vecchi, then CTO of Time Warner Cable’s Road Runner High-Speed Online Service, for bringing these numbers to our attention. I am forever grateful to Mario, Tony Werner, Dave Fellows, Alex Best, Randy Elliott, and the many other cable industry CTOs who were immediately staunch supporters of our efforts.

through Prodigy alone and a new Internet domain was registered every two minutes. It was a worldwide phenomenon—even 16% of Estonian K-12 schools were already connected to the Internet. In Mario Vecchi's opinion "It won't be long before revenues for services to the home PC will be comparable to the revenues for Basic TV." Mario certainly got that right!

Mario also warned of backfires the industry couldn't control such as bog downs (slowdowns due to excessive demand and resulting traffic) in the network cloud; bog downs in content providers' networks and content sources; and bottlenecks in endpoints such as workstations and personal computers. In our testing we had already seen bottlenecks in underperforming TCP/IP stacks in personal computers.

As you can imagine, the world around us did not stand still while we were theoretically evaluating responses to the High Speed Cable Data Services (HSCDS) RFP and practically evaluating the existing modem vendors LANcity, Zenith, and Hybrid Networks. Demand and innovation in cable modems was high and many new modem vendors fielded products in an incredibly short period of time. New cable modem systems became available from General Instrument, Scientific-Atlanta, 3Com, Hewlett Packard, Intel, Motorola, Sega[2], and many others. A typical design of the day included HFC nodes of 500 homes passed supported by 64 QAM (Quadrature Amplitude Modulation) in a 6 MHz downstream channel yielding approximately 25 megabits per second to a set of users, and QPSK (Quadrature Phase Shift Keying) in a 2 MHz upstream channel yielding approximately three megabits per second from users. The new modem systems touted network management capabilities and were expected to be available for less than $500 by the fourth quarter of 1995 or first quarter of 1996. I wanted an industry standard modem under $100, and eventually got one for about $35—but I am getting ahead of myself in this story.

The new systems had many improvements over previous designs from existing modem manufacturers who also raced to keep ahead and were busily upgrading their designs. We poured over the HSCDS RFP responses again and again looking for a modem system that at least in theory was "good enough." Among all the different modem systems from dozens of manufacturers, we just weren't seeing a design that met all the criteria we thought necessary to take the cable industry into the next millennium. And, despite all the proposed innovations that tantalized our heads, our guts where telling us that we must keep up our theoretical and practical exploration into the details of cable modem system operation and design.

In 1995, we became aware of an already established activity separate from CableLabs which was funded by several MSOs under the name MCNS (Multimedia Cable Network System). The original MCNS partners included Tele-Communications Incorporated (TCI), Time Warner Cable, Cox Cable Communications, and Comcast Cable. The MCNS partnership intended to invest in MicroUnity System Engineering Corporation of California, who was tasked with helping to define general microprocessor technology for a next generation of set-top boxes.

[2] The Sega modem was part of a gaming offering available as a tiered service known as the Sega Channel. It was used to download a Sega games and was not a general Internet access cable modem.

At the end of 1995, cable executives and equipment suppliers put out a call for standardized data connections at the Western Cable Show, and cited an agreement with key equipment manufacturers to specify some of the technical ways cable networks and data equipment talk with one another[3]. The announcement included the Executive Committee of CableLabs, the MCNS partners, Rogers Cablesystems of Toronto, and Continental Cablevision of Boston. Representing equipment manufacturers were Michael M. Ozburn, General Instrument Communications Division, vice president for telecommunications; Webb McKinney, general manager, Home Products Division, Hewlett-Packard; Rouzbeh Yassini-Fard, president and CEO, LANcity; James M. Phillips, corporate vice president, Multimedia Worldwide Distribution and Marketing Division, Motorola; J.A. (Ian) Craig, president, Broadband Networks, Nortel; Robert Luff, chief technical officer, Broadband Communications, Scientific-Atlanta; Iwami Asakawa, senior vice president, Media and Entertainment, Toshiba America, Inc.; and William G. Luehrs, president, Zenith Network Systems Division.

At the event, CableLabs' chairman, Dr. John C. Malone, then president and CEO of Tele-Communications, Inc. said, "We applaud the research that went into today's modems, and obviously wish to use it to get into the marketplace first. But, in the next generation of modems we look for more commonality in cable so that vendors may enjoy mass market sales and cable customers can be assured that their devices work on TCI systems or Time Warner cable systems or any cable systems in the world."

There was also specific mention of cable companies and manufacturers working toward an open standard that included proprietary portions that would be available on a non-discriminatory basis, and CableLabs was tasked with coordinating the specifications process with its members for these connection points called interfaces. Further, it was said these specifications would then be presented to recognized standards setting bodies for approval as standards. High expectations were set as the industry indicated that several interfaces were within weeks of being specified, for example the connection between the cable modem and computer would be Ethernet 10BaseT (but not ATM 25 as some had hoped).

Immediately following the Western Cable Show the MCNS Partners issued an RFP for a Specifications Project Management Group and by the end of the year had engaged the consulting firm Arthur D. Little to develop a set of specifications. Slowly at first, the CableLabs High Speed Data Team (Jerry Bennington, Tom Moore, Brian Riley, Pam Anderson, and I were solicited for contributions to the specifications. This started a long-term relationship with the Arthur D. Little principles on the project including Stuart Lipoff, Roger Hey, Michael Litvak, and several others. Our contributions to the MCNS specifications were sparse at first, though quickly developed into a regular dialogue via frequent meetings.

At the time, we did not know then that our many suggestions to MCNS would end up back in our laps within one year's time.

---

[3] SPECS News from CableLabs, December 1995, Volume 7, Number 9.

Chapter Eleven

Lehr und Kunst, Tom & Brian's Magic, and Cyber Me meets Maxi Me

The end of summer, fall and winter of 1995 was the most magical time we had yet to experience. Tom Moore and Brian Reilly were working incredibly hard. Some might call what happened magic, others would call it solid engineering. The Germans call it Lehr und Kunst, meaning Theory and Practice.

Tony Werner's question of how many users could be supported on an upstream or downstream had an inescapable grip on our minds. Compared to all that we actually needed in order to suitably answer Tony's question, our work with the 60 second surf and Mini Me, while significant, felt like a single grain of sand on a beach or a drop in the ocean. We had a long way to go.

We needed much more varied traffic sources and we needed to send traffic bi-directionally, meaning to users and from users. Our tests with RTBench only simulated one-way transmissions of data sent from a modem at the Headend location to a modem at the customer location. We needed much more flexibility in defining traffic types (email, World Wide Web, FTP, gaming, etc.) and a whole heck of a lot more modems to be simultaneously put under test—tens, maybe hundreds of modems sending two-way traffic, if possible. Mini Me needed to grow up—very quickly!

Our friends at Wandel & Goltermann had another testing product named Domino that was suited for embedding in strategic locations across widely distributed internetworks. We had a close look at Domino and other test equipment, and ultimately decided that a better fit for our need of a multi-port traffic tester was the Netcom Systems SmartBits® SMB-2000 and associated software. Jim Jordan, our sales contact from Netcom Systems in Los Angeles, was eager to work with us and got us a loaner system right away. A single SmartBits chassis could be loaded with up to twenty individual cards, each one capable of simultaneous traffic transmission and reception. And multiple chassis could be ganged (cabled) together to make an even larger multi-port tester. Bang! In an instant our vision for a 40-port or someday even a 120-port traffic tester was born. With the SmartBits tester we could grow Mini Me into Maxi Me (a Mini Me on steroids).

We had already proven that the traffic created with Mini Me was sufficient to overload and in some cases overrun and take down (disable) a cable modem system. And soon with Maxi Me we would be able to subtly, or not, wreak new types of traffic havoc on modem systems in order to understand nuances of their operation and behavior. Once again, the vendors were supportive in giving us more modems—so we could expand our tests and look at performance of multi-modem networks in actual two-way operation. The theoretical possibilities of multi-port testing tantalized our minds. Quantifying and understanding the behavior of Maxi Me was exactly the next step in practical benchmark testing that we needed to perform. In a few weeks, we had multiple modems under test sending and receiving traffic. And once again, the results showed many thought provoking and sometimes glaring differences between the various modem systems.

An insurmountable obstacle remained. The traffic we were sending and receiving to each modem, despite our efforts to make it as realistic as possible, was not very reflective of the actual traffic customers would send and receive on a cable modem network. The traffic we could create and a measure was largely limited to constant bit rate, while customers' traffic was much more variable, stopping and starting, bursting fast and then slow! We established long-term relations with Netcom Systems to make their traffic sources as realistic as their SmartBits hardware and software could allow. And as new, more powerful SmartBits cards were released by Netcom systems, we were provided sample cards and immediately put them to use.

About this time, Tom Moore got the idea of building a computer-based virtual model that would describe the behaviors of Mini Me and Maxi Me. Tom quickly found and made friends with Mark Cohen, one of two brothers who founded the company Mil 3[1] and sold modeling packages known as OPNET Modeler and OPNET Planner. Very quickly Tom attended class at Mil 3's Washington, D.C. headquarters and soon had his own copies of all of Mil 3's software running on a Sun Microsystems SPARC WorkStation on his desk at CableLabs. Tom worked around the clock building his virtual cyber model of Mini Me.

In developing the cyber model of Mini Me, Tom took small steps at first, such as building the virtual network segments for a dialup telephone network, then building the virtual segments for an Ethernet network. Given that LANcity modems had fared so well in our testing, it was decided that the first of the to-be-added virtual cable modem segments should represent the behavior of the LANcity modem system.

Very quickly, Tom's cyber model was a fascinating beast. Each LANcity modem was represented by a virtual state machine, meaning that the software model kept track of what each modem was doing at any instant in time. One modem would be in the transmit state, while the other modem would be in the receive state. For each tick of time, measured in microseconds, the model would predict the current state of the system and prepare itself to transition to what would happen in the next tick of time. Through an exhaustive analysis and education on the protocols involved and the operating states of the machines inside the LANcity modems, every bit, Byte, Ethernet frame, IP datagram, and TCP packet was accounted for at every instant in time. A modem would transition from the transmit state to the idle state when it had nothing more to send. Likewise, a modem would transition from the receive state to the idle state when there was nothing more to receive. Tom's state-driven model, Cyber Me[2], was incredibly complicated and magical at the same time. The magic was that Tom had convinced us, and we all believed, that Cyber Me could describe reality and even predict the future!

Before long, Tom was using Cyber Me to tell Brian what to expect on a particular test that Brian was planning to run on Mini Me. Think about the power of that for a moment. Tom was predicting the outcome before Brian ran the tests! At first, Cyber Me's predictions were incorrect, and Tom and Brian

[1] Mark was such a fine gentleman. He immediately understood what we were trying to accomplish and backed our efforts in every way possible. Mil3 later changed their name to OPNET Technologies.
[2] The model evolved to be called the "Common Simulation Framework", though Cyber Me imparts a good image for now.

would wrack their brains until they figured out exactly why. Then, Tom would tweak the virtual model inside Cyber Me and come back with a revised set of expected test results. Before long, Cyber Me would accurately and reliably predict actual test results from Mini Me.

Not to be outdone, Brian grew Mini Me as quickly as possible by adding additional modems each driven by traffic from SmartBits ports. Under Brian's great care Mini Me grew into a much more complicated Maxi Me—and soon, Cyber Me, was having difficulty accurately predicting all of Maxi Me's behavior. This galled Tom and, reminiscent of our college days, he pulled many all-nighters trying to figure out what was wrong with Cyber Me[3]. Tom was keeping in contact with our new friends at LANcity, John Ulm and Wilson Sawyer, and they could see how hard he was trying to succeed. I remember the smile on Tom's face when he said the Cyber Me model was working again! As true a gentleman as Tom is, he credited his success to the help of the LANcity team.

And so it was that Brian would try to do some yet to be tried, off-the-wall traffic test with Maxi Me, and Tom would try to predict the outcome with Cyber Me before Brian was finished with the test! Through his practical tests with Maxi Me, Brian would arrive at results for throughput, delay, jitter, and loss that Tom would confirm with the theoretical model Cyber Me. And vice versa, Tom would theoretically predict test results, which he would ask Brian to prove practically (empirically) in the laboratory. The spirit of Lehr und Kunst (Theory and Practice) was alive and well, and we were learning a lot!

The most exciting thing by far, which Tom realized first and we all greatly appreciated, was that Cyber Me could work with (i.e., virtually transmit and receive) much more realistic World Wide Web, email and FTP traffic patterns than Maxi Me. Unlike Maxi Me, Cyber Me was not limited to Constant Bit Rate (CBR) traffic, and as such, was able to provide much more accurate predictions of the effects of real users on real networks. We could see that Cyber Me's capability to reflect reality and even perform ”what if” analysis on future traffic patterns gave us enormous potential to explore the most subtle nuances of modem operations—which would serve us well over time!

Cyber Me was useful in comparing the behavior and performance of the recently popular 28.8 kilobits per second dial-up analog telephone modem networks versus; emerging 64 and 128 kilobits per second ISDN digital telephone networks versus; multi-megabit per second cable modem networks.

Cyber Me was also useful in preparing to analyze impacts of nascent emerging user behaviors such as the Super Bowl receiving six million World Wide Web hits in one day and the IBM ChessCam receiving five million World Wide Web hits in one day. Of paramount importance to cable operators was the cost per customer for a given quality of service level. Cyber Me would soon help us answer that question, too.

By the end of 1995, as we watched the Internet grow and we thought of all the ways modems could be used in the future, we were troubled on several fronts. One troublesome front involved the ability to create and enforce traffic shaping policies. Though the LANcity modem peer-to-peer modem system

---

[3] You will find on this earth no more committed worker or better storyteller than Tom Moore and the adventures we were on were like no other!

moved more traffic faster than any we had seen to date, a huge detractor in a network of "peer" modems that operate independent of a Master, is the inability to enforce any sort of policy to manage traffic during times of congestion. Left to making decisions on their own, each modem will do as best as it can, but lacking an overarching higher intelligence, an overview of the current and historical utilization of the channel as well as each individual user, and the authority to manage traffic appropriately as needed; a "peer" modem network could not easily make decisions to allow priority traffic to be identified and transported during times of congestion emergencies.

Everywhere we looked, Internet usage was growing through the roof and we knew that congestion emergencies would be real one day. The new Netscape[4] browser was driving an unprecedented increase in hits to their home page which had gone from five million hits a day in July 1995, to 10 million hits per day in October 1995, to 25 million hits per day in early February 1996, to 40 million hits a day in late February, 1996. All around the world there were increasingly bigger web pages with increasing changes of content on them—and content seemed to quickly fill the bandwidth capacity created by the powerful cable modem networks already being deployed by MSOs.

Nonetheless, we were now much closer to giving Tony Werner and the other cable industry CTOs an accurate answer to questions such as: "How many users can coexist and be supported simultaneously on a cable modem network?" and "How does the performance of a cable modem network degrade as more users are added over time and/or each user sends and receives more and different types of traffic?"

Over the course of 1995, we had actually learned a thing or two about modems and protocols though were still hamstrung by a lack of theoretical and practical knowledge of the actual RF operating environment that modems would experience in real life cable plants.

[4] Netscape was the commercial outgrowth of the original Mosaic web browser developed at National Center for Supercomputing Applications (NCSA) at the University of Illinois, Champaign/Urbana.

Chapter Twelve

Assessing the Trajectory of Time's Arrow and the Need for Operations & Business Support and Service Assurance

Nearly since cable TV's birth in 1948, engineering and equipment procurement has been a rather odd thing. Virtually all "new" engineering was done by one of two vendors: General Instrument (GI) and Scientific-Atlanta (SA). When a cable operator was planning to upgrade their service delivery system, GI or SA were usually deeply involved. In many ways, GI and SA had a lock on innovation in cable TV.

Once a cable operator started deploying equipment from GI or SA, they had to continue deploying equipment from that same vendor. Geographically speaking, a GI system, for example located in Metroville, would always be a GI system. The executive management of the Metroville system had no other choice than to continue purchasing equipment from GI. Consequently, price, performance, functionality, and schedule of availability were largely dictated by the cable industry vendors GI and SA.

As 1995 came to a close and we looked along the trajectory of Father Time's arrow[1] to assess and chart the future of the cable TV industry, the GI and SA duopoly looked like a real hindrance to innovation and prosperity—perhaps even to cable's survivability. GI and SA each had a competing conditional access system, and though overtures were made for making these companies work together, they mostly still didn't. From my novice perspective, the cable industry and consumers would benefit greatly from a single conditional access system that would work across all services—and all types of boxes in people's homes, whether set-top boxes or cable modems.

I had been promoted to Director of Data Applications at CableLabs, Tom and Brian came on board full time, and Pamela D. Anderson joined my High Speed Data Team. Pam was looking after network management issues related to operations support, business support, and service assurance. We were so pleased to have Pam looking after this extremely important area. Pam was well qualified and had 11 years of experience in information systems and architecture development. She had stints with NASA, Shell Oil Company, and BDM International where she designed local and wide-area data communications architectures and Integrated Network Management System (INMS) architectures, migration strategies, and deployment plans. Pam held a bachelors degree in mechanical engineering from Louisiana State University and a masters in telecommunications from the University of Colorado.

Pam had rightly identified that network manageability could make or break the data business in cable TV. Studies of corporate Local Area Network (LAN) Administrators indicated that one full-time technician was required for every 30-100 corporate LAN users. Such a costly support model wasn't going to work in cable TV! Cable modems needed extraordinary remote management capabilities.

[1] Time's Arrow was a term coined later by Tom Moore. See the following at www.timesarrow.net "The true entrepreneur is able to see down time's arrow and predict the changes necessary to engender new coherent form from old patterns or ideas and thereby create incredible value to society"

Other issues on which Pam was focused included essential behavior of modems, for example; automatic installation, configuration and even remote reconfiguration; support for roaming modems, computers and users; ability to automatically and remotely troubleshoot and benchmark performance of modems and one or more—or even all of—the segments of the network. Our vision included many potential benefits of network management such as distributed status monitoring, increased intelligence for operations staff, and even network and service self-healing capabilities.

Since my first week in cable TV, I personally had (and still have) a great interest in service assurance. Remembering the warnings of Mario Vecchi, my team and I wondered how an operator would ascertain if users were getting the level of service intended, or if users were experiencing some sort of degraded service such as slowdowns because of traffic congestion or HFC plant connectivity problems. We knew little about how an operator might measure if service levels were being delivered as intended. What did intended service levels mean anyway? Though there was a real-time report of "health" information from each LANcity modem[2], my "gut" was telling me there was still a lot of research, design, and development work to do in Service Assurance—what an understatement[3].

[2] HFC Plant Health was available from the Management Information Base (MIB) within each LANcity modem and could be queried in real-time using Simple Network Management Protocol (SNMP).

[3] I recall that much later in the fall of 1998, Rusty Picard, then Vice President of Operations at Road Runner, showed me a Service Level Agreement (SLA) that we at Road Runner were supposed to be able to deliver to our customers, Time Warner Cable and MediaOne, and in turn, help them deliver to their high-speed data customers. I remember saying to Rusty, "How the heck are we supposed to measure SLA?" Neither of us had much of a solution, and it took me years after leaving Road Runner to begin to figure one out!

Chapter Thirteen

The Best Bosses and Mentors I Could Ever Hope To Find

As we entered into 1996, we had very little idea of the wonderfully kind and gifted people we would meet.

I told you about my CableLabs hiring manager, Scott Bachman, and I should tell you more about this dear man. Scott grew up in Kansas and had wild tales of his teenage years. Having been in cable quite a long time, he had wild cable stories, too. In fact, all my favorite bosses and mentors in cable seemed to have wild cable stories—very wild—from testing atomic bombs to scouting for new microwave sites with snowmobiles, beer and whiskey …and everything in between.

Before coming to CableLabs, Scott most recently had been at Cox Cable Communications in New Orleans. While there, Scott experienced firsthand how big city politics and cable systems worked together. Scott's smile was broad and his hair stood almost straight up. He was a real "doer", empowering those around him with as much "rope" as they wanted. I remember Scott saying to me one day "Cruickshank, you are flying awfully close to the treetops!" The "buzz" I had from work felt terrific and I really had no idea what he was talking about.

Scott's department within CableLabs was known as Operations Technologies Projects (OTP). In 1994 and 1995, we were the fastest growing group within CableLabs. At that time, the entire CableLabs organization was about 40 people and had an annual budget of about $13.7 million, based on two cents per customer per month across some 57 million cable customers. Scott made sure that the CableLabs OTP and engineering departments had healthy rivalry, which kept staff productive.

As our responsibilities grew with the increased visibility and importance of cable modem testing, we began making jokes about what might happen if we screwed up. Unbeknownst to Scott, a favorite joke pastime was making up new meanings for our department name acronym OTP, such as, "Off The Payroll."

One day, Scott told me he was leaving CableLabs and taking a job at UPC in Europe. I was deeply saddened as I loved this man a lot. He was the best boss I had in a long time and I did not want to see him go.

About this same time, Dr. Jerry Bennington had recently come to CableLabs as an executive on loan from TCI Technology Ventures. Jerry had most recently run TCI's X*PRESS data business, which used narrowcast modems to distribute Reuters-like news feeds. He had a fascinating background including Ingenious™, What On Earth™, and TV Guide On Screen®. He had taught at North Carolina State University and the University of British Columbia, and worked at the MITRE Corporation and Bell Telephone Laboratories. He held two engineering degrees from the University of Michigan and an MS and PhD from Johns Hopkins University. Jerry's new job was "coordinating equipment interoperability in the cable industry and deployment of new digital technology for computers and the Internet[1]." Right away, I liked Jerry a lot. He was quick and a sage, though he intimidated many, perhaps because he had very little patience for anybody who said something dumb. I was thankful Jerry was my new boss. He gave me so many great gifts over time and I will never forget him.

Another soul whom I grew to love and who never stopped giving me an education and often much needed encouragement was Tom Elliott, previously CTO of TCI. Tom's contributions were extensive in the early operation of CableLabs, and he had such an interesting background! He grew up on a cattle ranch in Montana, pursued an interest in electronics that led him to Colorado Technical Institute and then back to Montana to work for EG&G Technical Services in their weapons test engineering department doing timing, firing and instrumentation of high-speed nuclear events (tiny bombs). Note that Tom's work at EG&G was located within the perimeter of the heavily fortified U.S. Air Force Base/UFO conspiracy hotspot known as Area 51. At TCI, Tom engineered and installed the first commercial fiber line for the U.S. Air Force at NORAD. Here is a quote that shows the "can do" spirit Tom engendered[2]:

Along the way I built the world's first commercial transportable [video] uplink, Elliot remembers. Boy, that was quite a job. We built that thing back in the early days of satellites, so we were out there pulling around the country, setting up for remotes in little towns here and there, doing baseball games, basketball games … whatever kind of thing. One time we're doing this small town in North Carolina when the damn power supplies quit. To my guys' credit, they went out and started taking the batteries out of everybody's cars, bringing 'em in. Got enough batteries collected that we were able to finish the game before the batteries went dead. Even the kids that were playing, their parents were like, 'sure, use our battery.' It was a big thing for these kids to be on TV.

Tom had done and seen everything and absolutely came alive when talking about almost anything. An endless supply of information on the cable industry, being with Tom Elliott was as rare a treasure as any I could hope to find.

Between the cartoon-like Tom & Jerry Show, there was never a dull moment. And if I ever needed any anything, both were always there for my team with a story from the past, or a vision of the future, or a quick exhaustive education [sic], or a pat on the back. Tom Elliott would later liken our modem research to the adventures of Lewis and Clark. Tom & Jerry were authentic. My thanks to you both!

[1] See "TCI Executive Bennington Joins CableLabs As Executive On Loan" from CableLabs SPECS News, November 1995 -- Volume 7 Number 8.

[2] Cable World Profile on Tom Elliot, June 3, 2002.

Chapter Fourteen

The RF Calvary Arrives

For many years, CableLabs and its member companies had been collecting data that characterized the radio frequency (RF) environment found in Hybrid Fiber Coax (HFC) plants around the United States and Canada. It was now time to put that data to good use in testing cable data modems. During the summer and fall of 1995, CableLabs built a major new facility called the Cable Test System (CTS) to be used for testing any vendors' digital transmission equipment. The CTS enabled controllable and repeatable laboratory testing of two-way cable equipment under simulated, yet real-world conditions.

At the outset of 1996, Dr. Richard S. Prodan, then CTO and senior vice president of engineering at CableLabs, focused some of his team's resources on evaluating the existing cable modem systems from LANcity, Zenith, and Hybrid Networks. His team's work was officially dubbed "Data Modem Transmission Testing" and was a perfect complement to the unimpaired modem testing, which was ongoing since the spring of 1995 by the High Speed Data Team: Pam Anderson, Tom Moore, Brian Reilly, and Bob Cruickshank.

The first modem to be evaluated by the CableLabs engineering department was Zenith. We helped Rich's team set up the Zenith system and attached it to the CableLabs CTS for some practical testing. In addition to a whole host of realistic Radio Frequency (RF) environmental conditions that could be created, the test apparatus provided RF test points at the input and output of both the headend and customer modems. A spectrum analyzer was used to capture entire packet waveforms, including the synchronization preamble. Packet waveforms of all zeros were compared with packet waveforms containing all ones. Average and maximum hold spectral power plots were charted for 256 byte packets transmitted at one megabit per second and 2.5 megabits per second.

Channel Symbol Duration, Packet Preamble, Inter Packet Gap, and Packet Duration were calculated for varying frame lengths and data rates. A mathematical function was developed that described packet length (duration in milliseconds) versus frame size (Bytes). Several curves of transmission performance showed Bit Error Rate (BER) and Packet Error Rate (PER) as a function of Carrier to Noise Ratio (CNR) for the Zenith modem. The test results summarized the theoretical limits that were possible to achieve and measured Typical Implementation Loss by plotting Zenith's BER as a function of CNR. To simplify extensive RF parlance, suffice it to say, the Zenith modem did not perform as well as other modems in the RF Data Modem Transmission Testing.

It was great to have the CableLabs engineering department engaged in the testing of existing vendors' modem systems. With a complement of the engineering team, as an industry, we had developed significant expertise with which to evaluate existing and proposed data modem systems. This expertise served us well over the short and long term.

Over the years, the CableLabs Engineering Department continued their testing and in late 1997, Daniel J. Rice published a comprehensive paper, “Modem Technology for Two-Way Cable Transmission.” Dan is a dear friend and an amazingly gifted man and leader at CableLabs.

Chapter Fifteen

A. D. Little, MSOs and Vendor Authors Collaborate in Writing the MCNS Specs

We entered 1996 with a great sense of wonder about many things; we had an amazing team and unprecedented motivation throughout the industry. It was clear we would be publishing specifications and were blessed with the help of Arthur D. Little, the MCNS Partners, and the CableLabs member companies all coming together in this important undertaking. Most exciting of all was the opportunity to work with the incredible team we had built inside CableLabs as we developed our understanding of high-speed data Internet access.

Left to our own devices, the CableLabs High-Speed Data Team deeply focused on the research we needed to do to understand where we could take the industry in the future—by designing more than just cable modems, including a complete Internet Protocol (IP) multimedia delivery system. My team didn't pretend to have all the answers and it really was the vendors that made the designs work, every time! My team was in a position to have a unique perspective across all the vendors, to learn about and leverage their individual strengths, and to orchestrate amongst them an incredibly powerful comprehensive approach to transporting IP over the cable television network. Again, it was a magical time. We would take an educated guess, point the way, and the vendors would save the day by working out the details.

After evaluating all the responses to the High-Speed Cable Data Services (HSCDS) RFP and participating in all the follow-up question and answer sessions during "hell week", we began developing a vision of exactly what the cable industry needed for the next millennium of multimedia services delivered over IP. But, we certainly didn't have the expertise to design such a next generation cable modem system. We made a lot of good contacts in the vendor community and realized that these vendors collectively could create something greater than any one vendor could create on their own.

It was as if we were looking at a huge Chinese restaurant menu of all the available functionality that could go into a cable modem system. We knew that the best people to advise us on selecting the items from the menu were the vendors themselves. But in the capitalistic society in which we lived, the vendors naturally wanted to (figuratively) kill each other. Or did they? Certain vendors were willing to take huge risks to make a standard cable modem come into reality. Certain vendors were smarter on the subject; in specific areas they knew what they were talking about more so than other vendors. The task before us was to figure out how to pull these vendors together in a way which they could collaborate, yet still remain competitive and able to differentiate their "standard" cable modem products.

Perhaps it was an idea whose time had come. Perhaps it was the overtures of some of the more willing cable modem manufactures. Perhaps it was encouragement from CableLabs and its members. However, it was likely all of these things along with the relentless support of Michelle Kuska and Steve Dukes at TCI that brought the whole Vendor Author concept to reality. Michelle and Steve led the

MCNS team and were the real doers in bringing together Arthur D. Little, Vendor Authors, MCNS partners, Rogers Cablesystems, Continental Cablevision and CableLabs.

On the cable modem interoperability specifications project, there already existed very pressing deadlines and the interface specifications were divided into three phases:

- Phase One involved the connections between the cable modem and a personal or business computer (Cable Modem Customer Interface, CMCI) as well as the cable headend to the wide area network (Cable Modem Termination System Network Side Interface, CMTS-NSI), and was released to the equipment suppliers for comment on April 15, 1996.

- Phase Two involved the Data Over Cable System Operations Support Systems Interface (OSSI) as well as the interface for a cable modem to use a telephone line for the return path (Cable Modem Telephony Return Interface, CMTRI), and was scheduled to be available for equipment suppliers' comment later in the second quarter of 1996.

- Phase Three specifications included the security management interface (Baseline Privacy Interface, BPI), the Downstream Headend RF interface, Upstream Headend RF interface and the Cable Modem RF interface, collectively known as the DOCSIS RF interface, or RFI, which was published at the end of 1996 at the Western Cable Show.

The Phase One interfaces were fairly straightforward and released on schedule. The Phase Two interfaces took some additional time and the OSSI interface was released for comment twice. By far, the greatest amount of uncharted new ground-breaking work was in creating and refining the Phase Three specifications, especially given that a guiding principle in the development of all the cable modem interoperability specifications was the requirement to use available technology. This was to prevent pushing the technological envelope too far, which would create undue schedule risk or development risk.

The two interfaces that contained the most new and interesting work were the RFI and OSSI. The RFI was an enormous undertaking and there were very few people in the world that could fully understand all of its many details and design nuances. From my perspective, the RFI involved every piece of cable technology and data/ telecommunications engineering that many of us had ever thought about, and much we had yet to ever think about. It was to be one of the great pearls we created in DOCSIS, the OSSI was another, and the Security Specification, BPI, a third.

In evaluating all of the responses to the HSCDS RFP, it became clear that the physical layer of the RF interface should be based on QPSK and 16QAM transmission in the upstream direction, and 64 QAM and 256 QAM transmission in the downstream direction—and should support multiple options for packet error protection (using trellis coding, interleaving and Forward Error Correction), as well as multiple upstream and downstream data rates. To design this specification, we would need the expertise of Broadcom and General Instrument. It also became clear that the foremost experts in the Media Access and Control (MAC) sub layer of the Data Link Layer were to be found at LANcity Corporation. LANcity also had the most practical knowledge of the use of SNMP for Operations Support

Systems (OSS) that monitored and reported the health of the cable infrastructure. A final skill that was needed was the area of telephony return. For this expertise, we chose 3Com Corporation.

Consequently, our vendor authors were selected. The remaining business issue to be worked out was the use of each company's respective intellectual property including their people and patents, trademarks, etc. The oath that was agreed to with each of these vendors was essentially this: "At my own expense, I agree to provide my staff, my patents, and my trade secrets to create the DOCSIS specifications and I further explicitly agree that I will not charge (or sue) any vendor for the use of this intellectual property." Think about that for a moment. What an enormous commitment to make. I so deeply respect the vendors that took this commitment and had the vision to see what it would do for our industry. The way I rationalized their decision was that the vendors were trading their market share, a large slice of a small pie, in exchange for a smaller slice of a very much larger pie.

We met monthly for about a year assisted by the capable minds and hands of Stu Lipoff, Roger Hey, Michael Litvak, and others from Arthur D. Little. The meetings were in the United States, some in Boston, some in Chicago, and some at CableLabs in Colorado. The meetings had a large attendance of the Vendor Authors, MCNS partners, Continental Cablevision, Rogers Cablesystems, and the CableLabs High Speed Data Team (Jerry Bennington, Bob Cruickshank, Tom Moore, Brian Reilly, and Pamela Anderson).

Broadcom sent Tom Quigley, Henry Samueli, and later Chuck Anderson. General Instrument sent Thomas Kolze. LANcity sent Gerry White, John Ulm, and Wilson Sawyer. 3Com sent Levent Gun, Jack Fijolek, and Beser Burcak. TCI representatives included Michelle Kuska and Steven Dukes. The Cox Cable Communications representative was Alex Best and later Jay Rolls. The Comcast representative was Steve Craddock and later Franklyn Athias. The Time Warner representative was Mario Vecchi and at one point, Peter Bates. Continental Cablevision sent David Fellows, John Leddy, Kip Compton, Lou Steinberg, and Rich Woundy. Rogers Cablesystems sent George Hart, Albert Kim, Esteban Sandino, and Shehab El Nawawi.

Not everybody attended every meeting, and a few only attended once, but we always had a quorum. The meetings were unlike any I had ever attended. We were writing specifications that were intended to be standards, so the very essence of what we were doing was similar to what you might find in an open standards meeting that is run by Robert's Rules of Order[1]. But, this was not an open standards meeting; it was more like "standards process" meets "customer is king".

Imagine David Fellows and the other cable operator CTOs quietly and intently reviewing their prior month's stack of trade magazines, such as Communications Technology, Multichannel News, and CED magazine that each had lugged along to the meeting in anticipation of some quiet time to catch a few news bytes. Picture David and the other MSO representatives having been quiet for twenty minutes or more while the A. D. Little Team and vendors progressed in the rigorous standards-setting decision-

[1] Robert's Rules of Order is the recognized guide to running meetings and conferences effectively and fairly. Please see http://www.robertsrules.org.

making process. Now imagine a couple of vendors get stuck arguing over whether the least significant bit should come first or the most significant bit should come first when transmitting a Byte. In an open standards setting environment, such a bout could drag on for some time depending on how entrenched or how far along a certain vendor was with their design and there silicon chip(s) implementing their design; this is exactly the kind of arguments the CTO group would quickly nip in the bud. Fellows or Craddock would look up from their trade magazine or laptop and ask, “Does it really matter which bit comes first?” This would often defuse the argument and a sensible resolution could be reached very quickly, and if not, the reply from the MSOs would be “Just pick one method [least first] or the other [most first] and move on.” And so it was that the writing of the new standard interface specifications went, relatively speaking, very quickly.

Certainly, there were some questions that haunted us for many months. For example, whether or not we were writing just interface specifications or also writing functional specifications, whether the cable modem system should operate as a Layer 2 “bridging model” or a Layer 3 “routing model”, etc.

When word of what we were doing got out in the industry, some of the participants in the Institute of Electrical and Electronic Engineers (IEEE) got upset that we were usurping the cable modem standard they had been working on for over a year. The IEEE 802.14 Working Group was chartered in November of 1994 to create standards for data transport over traditional cable TV networks. In frustration, several participants said, “You shouldn’t be writing cable modem specifications. That is our job!” We were thankful for the efforts of the IEEE 802.14 Working Group, but we just didn’t see them making progress quickly enough to meet the needs of the cable television industry– so we forged ahead in parallel, overbuilding their efforts.

Chapter Sixteen

Atlantis 1: The DOCSIS UNI

When designing DOCSIS, we included the concept of an easy to use yet powerful user-to-network interface (UNI),which has laid mostly dormant for nearly two decades[1]. To understand the significance of the UNI, think back to 1840 and the communications power created by Samuel Morse enabling a user to signal the network with only a dot or a dash. Now, fast forward to 1963 and the introduction of Touch-Tone® signaling in the telephone network, which enabled credit card authorization, telebanking, voicemail, pagers, and many more features we take for granted today.

The creators of DOCSIS had an appreciation for the power that could be unleashed by a thoughtful UNI and signaling method. Alas, the UNI seems to have gone the way of the lost city of Atlantis. Decades after the inception of the UNI, in this current time when the UNI could be such a great help, it appears few are thinking about it.

The concept of the UNI is simple yet very powerful. As envisioned, the UNI has the ability to use multiple service identifiers (SIDs) coming from each modem and works as follows; upon launching any application on a PC, tablet or Wi-Fi connected smartphone, the application would tell the network what data rate and Quality of Service (QoS) it required. For example, depending on default settings; email might request a best-effort service: web surfing might request a better best-effort service; Peer-to-Peer (P2P and P4P) applications such as BitTorrent might request a lesser best-effort service; streaming audio, video and voice would request a real time service; etc.

By signaling in such a manner across the UNI, the DOCSIS network could decide how best to allocate its bandwidth resources to meet the needs of every user and every application. Users would have the option to select (and maybe pay for) different service levels, or change the default settings used by their own applications, etc. Perhaps, most importantly, the DOCSIS network would pass the user requests to northbound networks that could, in turn, efficiently use their resources to better meet the needs and expectations of users.

Despite a lack of broad adoption, we have a good time-tested working example of DOCSIS UNI signaling that scales effectively and performs as designed. Considering that every time a cable telephone customer (MTA user) goes “off hook”, the UNI automatically signals the network to provide priority to voice traffic. As an industry, what we need to do is coordinate with applications developers such as Apple, Microsoft, Google, BitTorrent, and others to leverage the unused signaling capabilities in the UNI. By doing so, cable could be “the good guys” by creating a positive public relations opportunity—ahead of the telephone companies who certainly are on their way to figuring out this opportunity on their own.

---

[1] The Author owes a debt of gratitude to John Leddy, Comcast Senior Vice President of Networks, for encouraging an awakening of the DOCSIS UNI during Fair Bandwidth Management discussions at Comcast in the spring of 2008.

The UNI transfers communications power to users by giving them control to signal their intentions to, and expectations from, the network on an application-by-application, or even session-by-session basis. In so doing, the UNI encourages users to “self-select the treatment of their traffic”, thus contributing to fair bandwidth management practices.

Chapter Seventeen

Atlantis 2: The DOCSIS PHY that Could, but Doesn't

Data transmitted to and from cable modems ride on the physical layer much like trains ride on tracks. Without tracks, trains don't go anywhere and without the physical layer, modems can't send data anywhere. When designing a railroad bed, great care is taken to make sure the tracks are laid down in a manner that supports the highest speeds possible. It is the same with cable modems; the physical layer should be designed to support the highest data transmission speeds possible.

Recent testing by the CableLabs engineering department indicated that the choice of physical layer modulation formats used by existing and many proposed cable modem systems was not taking full advantage of the potential speeds of the HFC plant. This was cause for great concern on our parts! With the explosion of Internet activity around us, we knew the physical layer needed to transport data to and from cable modems as fast as possible. Another thing we knew was that the physical layer should be a completely independent subsystem, which means that the physical layer could be changed in the future[1] without disrupting operations that are performed by higher layers of the cable modem system[2].

The extensive work to characterize the actual RF operating environment that would be experienced by cable modems was discussed during a CableLabs conference on the HFC return path in Dallas on February 9, 1996. "HFC Return Path: Paving the Way" brought together CableLabs members and the CableLabs engineering department and helped quantify the varying RF environments that cable modems would encounter.

At the conference, reports came in from all over North America confirming that cable TV HFC RF environments differed significantly in both one specific area of the country as well as across the country. Just like the network of roads and highways that crisscross a town and the world, some portions of the RF network could support very high speeds, other portions could support medium speeds, and yet others only very slow speeds. Just as in driving a car, a road that will support high speeds, like the German Autobahn, will also support low speeds. However, a road that will support low speeds, maybe one made of dirt that is full of potholes and ruts, simply will not, in any way, support high speeds.

---

[1] We were successful in designing an independent physical layer; DOCSIS 2.0, DOCSIS 3 and DOCSIS 3.1 are examples of physical layer changes that did not disrupt higher layers of the cable modem system.

[2] Where possible, DOCSIS was designed to follow the Open Systems Interconnection seven-layer reference model (aka OSI model), which is an abstract description for communications protocol design. From top to bottom, the OSI model consists of the Application, Presentation, Session, Transport, Network, Data Link, and Physical Layers. For example, TCP and IP are the transport and network layers, and the DOCSIS MAC protocol is a data link layer.

Considering the millions of HFC plant segments around the world, we came to realize that the DOCSIS physical layer must be able to support a very wide range of speeds depending on the actual RF environment in any particular HFC plant segment. What we came up with was an extremely powerful 100-speed transmission for DOCSIS 1.x that could reliably and expediently transport data anywhere in the continuum of the worst of RF environments to the best of RF environments. Just as a ten-speed bike has two sprockets in the front, and five sprockets in the rear, which allow it to go slowly or quickly depending on terrain, the DOCSIS PHY had different speeds for the different HFC plant segments it would encounter. To calculate the number of speeds on a bicycle, the two sprockets in the front are multiplied by the five sprockets in the rear resulting in 2 x 5 = 10 speeds.

The DOCSIS physical layer transmission "gearing" that the working group ultimately specified was more complicated, but may be described using similar logic. In this case, imagine a bicycle built-for-two that has an additional intermediate sprocket between the front and rear sprockets. In the DOCSIS 1.x physical layer there are two sprockets in the front (for selecting modulation type), five sprockets in the middle (for selecting channel width), and ten sprockets in the rear (for selecting appropriate level of Forward Error Correction); resulting in 2 x 5 x 10 = 100 speeds. This DOCSIS 1.x physical layer is quite amazing and faithfully moves data error-free in a wide range of HFC plant operating environments.

The Atlantis-like part of this story is that the intended operation of the DOCSIS 1.x physical layer was largely misunderstood and unused around the world. Instead of taking advantage of the 100-speed transmission, operators instead ran their plants in one, two or three speeds, leaving the remaining 97 speeds unused. For illustrative purposes, consider that cable TV operators for almost 20 years were using either 8th speed or 38th speed. And in DOCSIS 2.0, where there are thousands of speeds from which to choose; still only two or three speeds were typically used by cable TV operators.

What are the downsides of only using a few speeds? The economic impact alone is enormous and the customer impact is even greater. Economically speaking, some estimates indicate that the HFC plant could transport data 40 percent faster than it does today. Not five percent, or ten percent, but 40 percent! Said another way, the existing HFC DOCSIS plant today can move 40 percent more data than it currently does. Operators spend millions of dollars each year expanding the DOCSIS footprint and today there are approximately 250,000 HFC plant segments in North America alone. The aggregate capacity of those plant segments, fed by HFC nodes, could be increased by 40 percent if the DOCSIS physical layer was used as intended. Fortunately, DOCSIS 3.1 is coming to the rescue!

To understand the impact on customers from cable TV operators not fully using the physical layer, consider that today some HFC plant segments are congested with high volumes of user traffic, and others prone to ingress of noise, are recklessly corrupting data packets from customers. Data packets from customers can be thought of as railroad cars that make up a train on a track. If the train is moving too slowly and there are too many railroad cars, or in this case packets, they cannot be moved expeditiously because of the slow speeds along the tracks, and congestion results. Conversely, if trains go too fast, they come off their tracks and the data packets, or railroad cars, become strewn and undeliverable.

An undeliverable data packet has so many errors that it no longer represents what a user was trying to send. As such, the packet must be discarded. In the case of best effort traffic, such as web surfing and email, the higher layer protocols like TCP will automatically resend that packet, but with a resulting cost in network inefficiency and a negative user experience. The network inefficiency is due to old packets having to be sent more than once; taking up resources that would otherwise be used for delivering new packets. In the case of real-time services such a voice on a telephone call, the lost packets are indeed lost forever. The packets are not resent because the conversation has moved on.

As we move to DOCSIS 3.1, and the DOCSIS physical layer is used efficiently as intended around the world, there will be significantly lower cost of network operations, along with better service provided to customers because of the alleviation of network congestion and the overall reduction in errored packets.

Chapter Eighteen

The DOCSIS MAC and the Lobotomy That Got DOCSIS 1.0 "Out the Door"

Considering all the successes we were fortunate to create in DOCSIS, the MAC protocol is likely the greatest accomplishment. Like a grain of sand that irritates an oyster and in so doing creates a pearl, we relentlessly thought about what to do with the DOCSIS Media Access and Control (MAC) layer. It had been almost a year since we started our evaluation of existing vendors' modem systems and we had learned a lot about the way a MAC protocol should behave in an efficient cable modem system. The behavior of the MAC, more than anything else, would make or break performance—and the business case.

We poured over the inner workings of the existing MAC protocols that were in existence, not just in cable modems but also across many data networking technologies. In 1973–1975, Robert Metcalfe had co-invented a protocol commonly known as Ethernet that used an algorithm known as Carrier Sense, Multiple Access with Collision Detection (CSMA/CD), which allowed data transmitting/receiving stations, or in our case cable modems, to transmit on a common wire. The operation of Ethernet can be likened to a conversation among a gathering of people. Before speaking in a group, one listens to see if anybody else is speaking; this is the Carrier Sense portion of Ethernet. If nobody is speaking, anyone in the group can speak up; this is the Multiple Access portion of Ethernet. If two people start talking at exactly the same time they sense a collision in their conversations; this is the Collision Detection portion of Ethernet. Once a collision is detected the speakers wait a short period of time[1] and re-engage by starting the process over again. That is, they listen to check if anybody is speaking and if nobody is speaking, they start to speak while sensing for collisions.

Stations in an Ethernet network act much like you and I when conversing in a group. As long as most of the group is quiet, although the group can be enormous in size, one or two people would have no trouble finding an opportunity to speak. However, similar to the difficulty one has when attempting to speak amongst a large talkative group, Ethernet often has difficulty when there are a large number of stations, or in our case cable modems, trying to find an opportunity to transmit.

Another network that we studied very carefully was Token Ring. This network can be likened to the Native American tradition of the "Talking Stick", where the person holding the stick is responsible for speaking and then passing the stick the next person in the group. Token Ring networks operate under a similar principle– the station holding the token is authorized to transmit for a fixed amount of time and then is responsible for passing the token to the next station in the network. If the next station, in our case a cable modem, has nothing to transmit, it passes the token to the next modem and so forth. In

[1] Ethernet's behavior during this waiting period is governed by the use of a random number generator that results in each station waiting a different amount of time before attempting to transmit again. This "wait" algorithm is known as binary exponential back-off.

very large networks, it may take a relatively long time for the token to make a complete trip around the network back to the first modem that has been patiently waiting for an opportunity to transmit.

We knew the cable network was going to have hundreds of customer modems per HFC plant segment with as many as half of the modems actively transmitting and receiving data during peak usage and a much smaller number of modems active during off-peak. Translating what we learned about Ethernet and Token Ring networks, we were able to ascertain the following: 1) An Ethernet network with a high number of stations with only a few wanting to transmit, would perform extremely well and deliver high data throughput at very low latency for individual users; 2) But, as the number of active Ethernet stations grew, performance of the network would degrade since simultaneous transmissions from stations would collide and force many stations to back-off, resulting in poor data throughput and very high latency for individual users; 3) Likewise, a Token Ring network with a high number of stations with only a few wanting to transmit would spend an inordinate amount of time passing the token and asking each station "do you have anything to send?" resulting in low data throughput and high latency for individual users; and 4) A Token Ring network with a high number of active stations would fairly distribute bandwidth, and in aggregate would move traffic very efficiently, though would have moderate to high latency for individual users.

We needed some way to combine the high individual throughput and low latency strengths of Ethernet when a few modems wanted to transmit with the high aggregate throughput of Token Ring when a lot of modems needed to transmit, while at the same time, providing some modems with deterministic low latency for applications such as voice and streaming audio and video.

This was the nucleus of the problem that we pored over again and again like an oyster surrounding its irritating grain of sand. We developed the solution by creating a new protocol for cable networks that would meet all these objectives—and are forever indebted to our LANcity friends who coached us and ultimately made us successful.

From the outset, we designed the DOCSIS protocol to include the ability to have contention-based "Best Effort" customer traffic simultaneously coexist with reservation-based low latency isochronous[2] traffic. None of the protocols we had studied could perform this important task, and this key capability in the DOCSIS MAC has been copied many times for use in other networks, including wireless networks.

In order to enable rapid time to market, we then simplified the DOCSIS 1.1 MAC protocol by giving it a "lobotomy", so that it would only provide Best Effort contention-based services for the initial release of DOCSIS 1.0. Later, once DOCSIS 1.1 was released, we proved many skeptics wrong about DOCSIS' ability to provide high quality telephone service, and enabled the efficient voice/video/data Triple Play!

[2] Email and web surfing are examples of Best Effort traffic where slight delays in data packet arrival times are imperceptible to users. Telephone and Audio/Video streaming are examples of Isochronous traffic, where data packets must arrive at equal time intervals.

Chapter Nineteen

Atlantis 3: DOCSIS Security and the "Remove-able", Removable Security Module

The story behind the DOCSIS Security System Interface (SSI) and Baseline Privacy Interface (BPI) is fascinating from several perspectives. The powerful MCNS partners had decided that the DOCSIS SSI would use the new security system they had been developing over the past year or so. I was initially relieved that the security system was already defined, as it resulted in less work for my team, none of whom were security "spooks" or desired to be security experts[1]. What an unpleasant surprise when I realized the MCNS security system specified an expensive removable security module in every modem. I balked at the insanity! To me, the removable security module put the entire DOCSIS project at great risk of failure.

The MCNS design had grown out of the need to protect video assets such as movies that to date had been delivered via one-way transmission, meaning that the receiver did not acknowledge (in any way shape or form) receipt of the transmission from the transmitter. Unlike DOCSIS, a one-way transmission system did not have the benefit of the evolving challenge and response[2] security approaches where the master in the cable headend can authenticate the identity of the slave set-top box in a customer's home. The MCNS design was based on a long history of "died in the wool" one-way thinking and as such, specified a separable (i.e., removable and hence renewable) security system. It was separable in the sense that the credit-card-sized security module could be removed from the set-top box, and renewable in the sense that a new decryption module could be reinserted at any time in the future to thwart would be hackers attempting to illegally steal cable service.

While I understood the desire and critical need for a security system that would protect movie studio, cable operator and personal assets, I did not agree the right approach was a removable security module. My quick analysis indicated the removal security module would add $35.00 in cost, effectively doubling the price of the modem, which was unacceptable. What I found most offensive was the removable security module would insidiously allow operators to promote a long-lasting duopoly in security systems. Whereby, each operator specified a security module that supported whichever security system the operator already had in place from General Instrument or Scientific-Atlanta. To me, the removable security module was a treacherous "Trojan Horse" that would allow the duopoly of General Instrument

[1] "Spook" is the nickname given to experts in security systems, many of whom have either worked for or with the United States National Security Administration. David Harrison of the CableLabs OTP Department was a spook, but could tell us no more unless we jokingly agreed to his killing us afterwards.

[2] Challenge and response algorithms involve the "master" asking the "slave" questions such as, "Are you there?" to which the slave replies, "Yes I am." A number of questions also possible included: "When were you last updated? Have you had a high number of resets? Have you accepted any new security software? What is the current version of the security software you are running?" Challenge and response significantly empowers security system designers to make systems more secure.

and Scientific-Atlanta to promulgate (potentially forever), much to the detriment of customers and the cable industry.

I was passionate about the insanity of this proposal and fought as hard as I could to sway the opinion of the supporters of the separable security model. My team and I argued that a two-way cable system could be made secure enough without requiring a separable security module, that a two-way cable system could use a challenge and response mechanism to authenticate identity of the endpoints and secure the privacy of communications.

Privacy of communications and protection of content assets will always be paramount. Content providers, operators and users all need to be assured that the data being sent to and from a modem is not observable by other modems on the network. Those unskilled in the art might intuitively think that a shared access network of cable modems could easily be hacked, so that one modem (or set-top box or MTA) could see another modem's traffic. As such, we were taking great care developing industrial strength designs for the security system interface.

Despite our efforts to promote a cable modem that did not have a separable security module, the MCNS partners were adamant in having one, and it was clear we were wasting our energy and were not going to win the argument. After much frustration, someone (perhaps this author), came up with the idea of introducing an optional built-in security system that would be renewable and not require the separable security module. The rationale I argued was that cable operators could select either the separable or renewable specification when they were purchasing modems from vendors.

We needed to come up with names that would distinguish the two security systems. Internally, we were starting to refer to them as Security Heavy and Security Light. Security Light sounds like an oxymoron so we knew we needed a better name. I don't recall who came up with the name Baseline Privacy, but it caught on and worked just fine. Almost overnight, we had added another security specification that we dubbed the Baseline Privacy Interface (BPI).

As history would have it, no MSO ever ordered anything but baseline privacy and the removable security module never made it to market; it was permanently removed!

The sad part of this story is that in future years, other designs in the cable industry reinstated the separable security module that we had worked so hard to remove. The One-Way Cable Card was an example of a disaster that you could see coming a hundred miles away– it promulgated the duopoly, and among other flaws, would never support switched digital video. Today, there is a waning discussion about how set-top boxes should be made to work with the Motorola conditional access security system or the Cisco Scientific Atlanta conditional access security system. The duopoly that we worked so hard to topple managed to stay in power much longer than I ever thought they would. The net effect has been a loss of functionality to subscribers—as evidenced over the years by the many unsightly onscreen guides from these two companies.

Chapter Twenty

Atlantis 4: DOCSIS OSSI

Designing and deploying a Cable TV Operational Support System is not easy, and if anybody tells you it is, they likely do not understanding the challenge. In designing the DOCSIS Operational Support Systems Interface (OSSI) we thought long and hard about the work of those that had come before us, both in the cable modem community and the Internet community at large. From what we could tell, the network management possibilities before us were endless. If we did our job well, each cable modem could be like a proverbial canary in a coal mine, able to warn of impending disaster. Just as a canary droops its head in the presence of deadly methane gas, so a cable modem could warn of customer-affecting issues in the HFC network such as traffic congestion, packet errors, ingress noise, and RF power fluctuations. The OSS represented a cornucopia of HFC plant support capabilities for cable TV operators. Gone would be the days of customers experiencing some sort of outage and having to call their cable operator to notify of a service-affecting issue. With a properly designed and implemented OSS, there would be no stranded cable customers hidden in the labyrinth HFC network– no stranded coal miners buried alive in the mine.

On my team, Pamela D. Anderson, CableLabs Project Manager/Enterprise Management Technologies, had the daunting job of being the CableLabs expert on the DOCSIS Operational Support Systems Interface. Pam worked diligently to assemble and revise materials that described what was possible and what we in the industry should try to accomplish. From January through March of 1996, Pam researched operational support systems, consulted with CableLabs members, assembled OSSI proposals, and developed a two and a half hour tutorial on Simple Network Management Protocol (SNMP). Decades before the days of Internet Protocol Data Records (IPDR), SNMP was the workhorse that would retrieve information from cable modems and headend equipment. Through rigorous correlation and analysis of SNMP readings, the OSS would remotely point to customers having service troubles without requiring customers to call-in and report difficulties to their cable company.

Pam delivered her SNMP tutorial to CableLabs members at the conference “Providing Internet Services” held on April 9, 1996 at the Regal Harvest House in Boulder, Colorado. Her goal was to get members thinking about what we as an industry could do with SNMP to provide operational support to staff throughout cable TV organizations such as: install and service technicians, HFC plant technicians, telephone customer service representatives, workforce dispatchers, staff in the network operations center, and management throughout the organization. Her presentation was extensive and included a high-level overview of network management objectives, a detailed review of TCP/IP and SNMP, guidelines for collection of information from network devices, collection architectures, the essentials for managing data over cable services, industry trends, and a hand-on demonstration of browsing the Management Information Base (MIB) of a working cable modem.

The business case for network management had a lot to do with the evolution of cable television systems, from delivering single one-way broadcast services to delivering multiple two-way interactive services. It was clear to us that interactive users would require more and different support than passive TV-watchers. This was demonstrated in data from the corporate local area network (LAN) industry as well as in feedback from the Rogers Cablesystems trial of residential High Speed Data (HSD) service in New Market, Ontario.

Pam's vision was that the headend was evolving to become a data center where automation would play a key role in the determination and management of Faults, Configuration, Accounting, Performance, and Security (FCAPS) in the HFC plant and service delivery infrastructure. Fault management involved identifying, isolating, correcting, and preventing failures. Configuration management involved activating, deactivating, setting, resetting, and maintaining information on network devices. Accounting management involved measuring network (resource) usage by individuals, groups, geographic areas, etc., and establishing metrics, quotas, costs, and user billing records. Performance management involved predicting network degradation, preventing network congestion, and sustaining consistent service levels to users. Security management involved authorizing service and protecting management and user information held or transmitted by any device on the network.

Pam's tutorial included a section on the current Internet standards process that she hoped the cable TV industry would leverage heavily in the creation of the DOCSIS OSSI. She described the Internet Architecture Board (IAB), which has policy setting and decision review authority over the Internet. She reviewed the role of the Internet Engineering Task Force (IETF), which is the technical working body of the IAB and the primary body that develops new standards such as network management for the Internet. She described the Request For Comments (RFC) process that was used to document and develop standards, procedures, and specifications.

Pam provided an example of the MIB objects that should be in a cable modem. These were broken into four different areas: 1) hardware parameters such as MAC address, transmit and receive frequencies, signal to noise ratios and power levels, broadband plant size, and modem make, model, and version control; 2) software parameters such as IP address, IP subnet mask, community name, read-write and read-only permissions, remote reboot and remote configuration; 3) Cable RF broadband statistics such as bits, Bytes and packets received and transmitted, receive and transmit errors, and collisions sensed, and; 4) Ethernet statistics which mirrored the broadband statistics.

The good news is that most all of Pam's aforementioned statistics and MIB objects did indeed wind up being included in the DOCSIS MIB and are widely used throughout the industry today. The bad news is there was an entirely more important level of Pam's tutorial that seems to have gone the way of the lost city of Atlantis. Pam focused on having an OSS manager, which would, for example, trigger an alarm based on an error rate threshold that exceeded acceptable levels. Pam's vision was that the Network Management System (NMS) would poll the cable modems at regular intervals for parameters such as, a count of total packets received and a count of those packets received that contained user affecting errors. Pam even included the exact calculation to derive the percent of packets received with errors, which is the sum of errored packets divided by the sum of total packets (since last poll). For almost two decades, this elegant and short calculation and associated alarm features were poorly implemented and underused used by staff to proactively find customer-affecting problems throughout the vast majority of cable systems in the world!

Pam also studied the Telecommunications Management Network (TMN) architectures and shared much guidance during the tutorial, which also seems to have been lost. The TMN architecture is part of the International Telecommunications Union, ITU-T recommendation M.3000 that describes a comprehensive telecommunications management network that applies to all types of telecommunications facilities. The TMN model includes a Business Management Layer (BML), a Service Management Layer (SML), A Network Management Layer (NML), an Element Management Layer (EML), and a Network Element Layer (NEL)[1]. The BML involves the management of business units such as, a residential or business telephony unit, or one or more geographic units such as Metroville. The SML involves the provisioning of services, the management of workers, customer interactions, and plant and customer troubles. The NML involves network monitoring and configuration, alarm correlation, performance monitoring and analysis, and RF spectrum management. The EML involves element management systems, protocol translation, and logs of statistical information. The NEL involves HFC plant devices, customer devices (modems, MTA's, and set-top boxes), routers, bridges, network analyzers, and proxy agents. Pam included extensive examples of the architectures that included element management systems and a rules-based "expert system" alarm correlation engine that would intelligently pull together and correlate alarms from multiple sources and present actionable results to cable TV operations staff in a consistent manner.

To this day, delivering Pam's vision remains the "holy grail" in Cable TV network management systems.

[1] For further details, please see: http://www.hit.bme.hu/~jakab/edu/litr/TMN/tmn.pdf.

Chapter Twenty-One

Tag, You're It—and—Oops, My Team is Gone

Talk about a roller-coaster adventure! In a very short time, I was promoted to vice president, was given responsibility for the whole DOCSIS project, and everybody on my team left CableLabs!

Jerry Bennington came in my office one day and said, "I just told Dick [Green] that you can run the department without me. I am stepping out of the way." I was shocked. What a vote of confidence from this wonderful man whom I greatly admired. I thanked Jerry profusely. In an instant, I was a vice president in the cable industry.

At the same time, it seemed the MCNS partners were getting tired of paying Arthur D. Little to do specification development. One day, I came to work to find out that the specification effort was now my responsibility. It was relatively easy to provide input to the specifications process, but to run the specifications process, that was an entirely different matter. In addition, to run the development of specifications without the capable and dedicated minds and hands of Stu Lipoff, Roger Hey, and Michael Litvak was inconceivable. I hated writing and was not good at it, and I hated having too many details.

Somehow, I knew we'd manage. However, next thing I knew, a huge "earthquake" struck which really left me flying solo. For a host of good and varied reasons, Tom Moore, Brian Reilly, and Pam Anderson all decided in the course of a week or two they would be leaving CableLabs. My team now consisted of me alone and my responsibility included finishing the design and ensuring interoperability of the next-generation cable modem system. Oh boy.

From left to right: Bob Cruickshank, Tom Moore, Pam Anderson and Brian Reilly

I really was stuck, not quite sure where to turn, and for a moment… at a complete loss for what to do next. There I was, the only person in the cable industry working fulltime on what became known as DOCSIS, the Data Over Cable Service Interface Specifications.

A few things were immediately clear. I did not have the expertise to finish the design of a next generation cable modem system, CableLabs did not have the expertise to finish the design, and the cable MSOs did not have the expertise to finish the design. In fact, the only people who had the required expertise were distributed far and wide among the vendor community, and most of these individuals were as yet unidentified and did not even know each other. Yet, there had to be some way to pull this group together, for they were the only ones who could finish the design.

And so well, too!

CONGRATULATIONS

Bob, Thanks for always leading the gang in such as "camp"-like attitude... It's been enjoyable... John McGeer G.I

Bob - Good luck in your new venture! The Cisco Crew

Bob! I really really appreciate all your help! Good luck! KAZ

Best of luck, Lynne

Good Luck Bob

Congratulations Bob on your new position. I will miss you! Paige

Bob - Best of Luck in your new job. Who is going to wear the pilot uniform next Halloween? Kathi Jack

Bob, Good luck. John C.

Bob, you're a great guy to have around. Best luck to you!

...hange in everything ...will be. Best wishes - Pat

Bob Good luck Kellie

Go for the gold (the real stuff)!! ME TOO! Thanks Bob! Casey Lyd :) Jerry Bevington

Bob Great job in pushing MCNS & then DOCSIS up the hill; Enjoy the other side. Mike Schwartz

Try to keep Roadrunner in line & don't let media one goof this one up! Doug

Good luck!! SO

Bob Thanks for making the Best team under your leadership to make the History of DOCSIS cable modem. without your great ownership we couldn't do it. Raybel

Bob Don't let anyone with SGS ever get in your way! You are a great manager and friend, and I hope our paths cross yet again. Thanks for the opportunity! Bob Ward

Great knowing Bob. Best wishes at your new adventure! Bryan

Good Luck Bob from the gang from Zenith

Bob - I didn't get to work with you for very long, but I'm sure our paths will cross again. Best of luck to you on the East Coast! Diana

Bob God bless you!! Shin Kay, Samsung

Wishing you + your family the best of luck! Have fun! Tasha

Bob, It's been a great opportunity to be involved with CL DOCSIS project. You did great keeping things moving forward. Best of luck in your new adventure. Bill Snyder

Bob, luck to you in your new venture. Good luck. Bob Jack Chen, Cisco

Bob - Good luck with your new venture! Gaby

Good luck on your new adventure! I hope Virginia is good to you + your family. Best wishes, Julie

Bob Good luck at your new job. We'll miss you. Stay in touch. Dan

Bob, Best of luck to you + your family in your move to VA. May the force be with you. Elaine

Bob, Good luck and have fun. Fred

In pursuit of a network to improve global energy efficiency by delivering real-time pricing of electricity, from 1994 to 1998, Robert Cruickshank explored and ultimately led cable modem development at Cable Television Laboratories, Inc. Under the guidance of CableLabs CEO Dr. Richard Green, in 1995, his team benchmarked the customer experience and overall performance of existing cable, telephone and satellite modems. In 1996, unsatisfied with proposals for next generation modems, they developed extensive mathematical models for a new world standard cable modem which became known as DOCSIS®. In 1997-98, in collaboration with cable operators and vendors, they created and guided specification development, refinement, modem interoperability testing and certification to ensure the Industry's success.

Over the next 20 years, Cruickshank provided strategic direction to the design, development, and deployment of DOCSIS Operations and Business Support Systems. Cruickshank's focus has been helping cable operators maximize available bandwidth, service levels, and field operations resources. Prior to R&D in Contact Center and Field Operations at Cablevision, Cruickshank joined ARRIS in December 2007 as part of the acquisition of C-COR, Inc. He joined C-COR as part of the acquisition of Stargus, Inc., which Cruickshank co-founded in June 2000 to help MSOs monitor and manage—and ultimately unleash—the power of cable networks using CableEdge ServAssure. Cruickshank's career also includes exciting work at Road Runner High Speed Online, where as vice president of Engineering and interim CTO, he was responsible for the design, deployment, and Y2K upgrades of broadband networks and data centers across the United States.

In addition to his broadband career, Cruickshank has held a number of distinguished positions with AT&T Bell Telephone Laboratories, CyberLYNX Gateway Corporation, and the Wyoming Air National Guard. Cruickshank attended Worcester Polytechnic Institute, tutored students in energy systems and thermodynamics, performed basic research in optics, and earned his Bachelor and Master of Science Degrees (with honors) in Mechanical Engineering. Later, while at the University of Colorado at Boulder, Cruickshank completed five years of postgraduate research in Energy Management, Programming, Telecommunications, Computer Science, Civil Engineering and Intelligent Buildings. Cruickshank is a fourth generation inventor engineer and retired United States Air Force pilot who loves operating machinery. Helping preserve his family's tree farm in the Catskill Mountains of upstate New York, he and friends gather and camp there every year around July 20.