

EXHIBIT G

Case No. 1:14-cv-00857-TSC-DAR



This book is provided in digital form with the permission of the rightsholder as part of a Google project to make the world's books discoverable online.





This book is licensed under a Creative Commons license. By using a Creative Commons license, the rightsholder chose to give you more freedom to share or re-use the book than would otherwise be possible under copyright law.

The rightsholder has waived all copyrights to this book and dedicated the book into the public domain. Terms available here: <http://creativecommons.org/publicdomain/zero/1.0/>

About Google Books

Google's mission is to organize the world's information and to make it universally accessible and useful. Google Books helps readers discover the world's books while helping authors and publishers reach new audiences. You can search through the full text of this book on the web at <http://books.google.com/>








26.95

50 A

Exploring the Internet

A Technical Travelogue



Exploring the Internet

A Technical Travelogue

Carl Malamud



PTR Prentice Hall

Englewood Cliffs, New Jersey 07632

22

TK
5105875
.I57
M361
1993

Library of Congress Cataloging-in-Publication Data

Malamud, Carl, 1959-

Exploring the Internet: a technical travelogue/Carl Malamud.

p. cm.

Includes bibliographical references and index.

ISBN 0-13-296898-3 : \$26.95

1. Internet (Computer network)

I. Title.

TK5105.875.I57M36 1993

384.3—dc20

92-16450

CIP

Acquisitions editor: Mary Franz

Jacket design: Bruce Kenselaar

Cover design: Joe DiDomenico

Prepress buyer: Mary E. McCartney

Manufacturing buyer: Susan Brunke

Copyright © 1993 by Carl Malamud



Published by Prentice-Hall, Inc.

A Paramount Communications Company

Englewood Cliffs, New Jersey 07632

The publisher offers discounts on this book when ordered in bulk quantities. For more information, write: Special Sales/Professional Marketing, Professional & Technical Reference Division, Prentice Hall, Englewood Cliffs, NJ 07632.


All products are the trademarks of their respective manufacturers. Carl Malamud is a trademark of Carl Malamud.

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.


Printed in the United States of America

10 9 8 7

ISBN 0-13-296898-3



90000



9 780132 968980

ISBN 0-13-296898-3

Prentice-Hall International (UK) Limited, London

Prentice-Hall of Australia Pty. Limited, Sydney

Prentice-Hall Canada Inc., Toronto

Prentice-Hall Hispanoamericana, S.A., Mexico

Prentice-Hall of India Private Limited, New Delhi

Prentice Hall of Japan, Inc., Tokyo

Simon & Schuster Asia Pte. Ltd., Singapore

Editora Prentice-Hall do Brasil, Ltda., Rio de Janeiro

I151
846-845X
Libsc
4-27-94

Contents

Foreword and Preface, in which we secure funding for a voyage which will take us three times around the world in six months and give new meaning to the phrase "many fine lunches and dinners."

Prologue: The Document Liberation Front, in which we delve into the abyss of international politics and wade through hordes of Jehovah's Witnesses to examine a Bulgarian FEP.

Geneva · Prague · Zürich · Geneva · Paris · Boulder

Round 1: From INTEROP to IETF, in which we meet the Internet Samurai and the Uncle of the NSFNET and go to the zoo to witness the birth of EBONE.

San Jose · Honolulu · Tokyo · Fujisawa · Akihabara
Hong Kong · Macau · Singapore · Dublin · Amsterdam
London · Tampere · Paris · Geneva · Nice · Geneva
Ithaca · New York · Washington, D.C. · Santa Fe · Boulder

Round 2: From Christmas to Cleveland, in which we encounter a massively parallel bicycle in Mt. View, find magic boxes in Oz, and go swimming in the Sea of Acronyms in Europe.

Berkeley · Mt. View · San Francisco · Moffet Field
Wellington · Dunedin · Auckland · Melbourne · Sydney
Canberra · Adelaide · Singapore · Kuala Lumpur · Bangkok
Amsterdam · Utrecht · Bonn · Brussels · Paris
Washington, D.C. · Cleveland · Chicago · Boulder

Round 3: In Search of a Standards Haven, in which we meet the world's most intelligent building, try to convince people to do the bloody obvious, and learn the origins of the Internet.

Marina del Rey · San Francisco · Tokyo · Seoul
Taipei · Hong Kong · Bombay · Madison · Boulder

Index

Geneva

Sunday night, I took the TGV back to Geneva. Feeling refreshed after a day in the south of France and a weekend in Paris, I was ready to sink into the bureaucratic abyss of Geneva.

Sink I did. I spent the next three days battling the ITU bureaucracy, trying to stop a rear guard action that was threatening to kill the Bruno project.

In four weeks, the Bruno server had been a remarkable success. Twenty-one servers on four continents had cloned the file system and were distributing the Blue Book. Bruno was getting as many as 35 packets per second. Over 500 hosts in 27 countries had retrieved over 65,000 files. We had no statistics from the other servers, but it was not unreasonable to think that several hundred thousand files of the Blue Book had made their way out to people who were actually reading them.

How did this compare with paper copies? This was hard to say, as profits from documents had served as a sort of discretionary fund for the previous Secretary-General of the ITU. Knowledge of publications was highly dispersed; only finance seemed to have sales data, and they kept this information closely guarded.

Nonetheless, it appeared that the Bruno experiment had increased the distribution of the Blue Book by at least one order of magnitude, and probably two or more. Tony had documented all this in his "Friends of Bruno" newsletter and had papered the ITU—paper being the only medium that appeared to work there.

Yet, despite all this, the high-level Information Systems Steering Group had met the previous Friday to decide the future of Bruno. Rumor had it that the outcome of the policy group was that the experiment had not been successful and was over.

Stopping the experiment was, of course, not an option. The server was in Colorado in a locked room and I had no intention of stopping operation. Besides, twenty-one other servers had the data. Tony and I had carefully structured this project so there would be no turning back.

Geneva

There were, however, some important factors at stake. I had hoped to begin putting other ITU documents on the servers. A policy decision that the experiment was over would mean that we might have to bypass the ITU and start scanning paper copies.

More importantly was the role of the ITU in dissemination of standards over the network. The logical outcome of the Bruno experiment was to have the ITU put itself on the Internet and take over this function.

Tony had set me up for three days packed with meetings. Many of these people were division directors or other Very Important Bureaucrats (VIBs) who sat on the Information Systems Steering Group.

One meeting, in particular, stood out over all the rest. I was scheduled to meet with Walter Richter, director of something important. He was 45 minutes late, so I spent the time looking at the stacks of file folders on his secretary's wall-to-wall bookcase.

They had wonderful labels such as "The Preparatory Committee on Restructuring of Subsidiary Machinery" or "The Administrative Committee on Coordination." One was simply labeled "High Level Committee" and took up several folders high up on the top shelf. The committee that seemed to take the most wall space was the "Consultative Committee on Substantive Questions."

Finally, Richter strode in. Speaking with a heavy Austrian accent, he preceded to tell me how my experiment "was not a success and has been terminated."

He seemed very certain that the experiment had not been a success, so I asked why. It appeared that this Internet of mine (the ITU considered the Internet to be some private project run by Tony Rutkowski and myself) just didn't reach the right sort of people. By the right people, he seemed to mean those who were on the Administrative Council of the ITU or those that worked on the consultative committees like the CCITT or the CCIR.

The conclusion that the Internet had the wrong sort of people was odd, since I had not analyzed the data on who was accessing my server. In fact, anecdotal evidence was pointing to just the opposite conclusion. I had received personal messages from places like AT&T, Bell Labs, and Telecom Finland.

Exploring the Internet

I found out later how Richter came to this view. Richter had a buddy on the radio side of the ITU, the CCIR. His buddy had a pal in Canada to whom he had spoken.

"Ever hear of this Bruno thing?"

"Nope."

"Ever hear of this Internet business?"

"Yeah, but we checked it out a few years ago and it was too expensive."

Well, there you go. Can't argue with a few personal anecdotes when making a high-level policy decision. I tried trotting out a few of my own anecdotes, but Richter had already assembled the data he needed.

Richter had something in common with most of the other VIBs that I met in three days at the ITU. He was very, very sure of himself. For example, he was absolutely convinced that the entire ITU network architecture was fatally flawed.

I must confess, this was certainly my working assumption when starting to deal with the ITU computer group, but the reality turned out to be that they had a fairly decent network architecture in place. Not what I would have chosen, but adequate for the job.

I asked Richter to tell me what was wrong.

"The Ethernet," he replied. When his PC had first been installed, it was a diskless machine. A mistake, of course, but it had been fixed. He was convinced that all ITU network problems had at their root Ethernet saturation, because it had once taken several hours after pressing a key to see the character appear on his screen.

Based on this anecdote, he was ready to completely micromanage some fairly talented engineers that worked at the ITU computer department. Rather than set broad policy (an area that had been sorely lacking), he was convinced that the answer was to roll up his sleeves and dig into the bits and bytes.

Another curious aspect of the three days of meetings was this idea of the Internet as some academic toy that real people didn't use. I met with one staff member who expressed this view and waved a piece of paper at me that had the names of delegates he was working with as proof.

Geneva

I picked up the piece of paper and started going through the list. Many of the places, such as the Centre National des Etudes de Telecommunications in Paris, were clearly on the Internet. In fact, the vast majority of institutions on the list appeared to be connected in one form or another.

There was another more fundamental issue that started to focus and helped explain this reluctance among the VIBs. Printing documents was a big empire at the ITU, and building empires was the name of game. My project was not a good way to build big empires (efficiency never is).

The printing department at the ITU was truly an impressive place. I walked past the "keep out" signs and gave myself a private tour. There were seven offset presses, four state-of-the-art, top-of-the-line Xerox 5090 copiers, and a dozen or so other large copiers. The ITU's own facility generated only a fraction of the total output. Swiss printers had a long and cozy relationship with the ITU bureaucracy.

Things would be printed with no relationship to demand under the assumption that larger print runs meant a lower per-unit cost. True, of course, but if you throw away most of the units, your average costs can be considerable.

One of my underground sources gave me an example. For several years, the ITU had produced a beautiful four-color "charts in profile" document. Each time, 10,000 copies would be printed at a cost of several hundred thousand dollars.

Of this print run, 2,000 copies would be given away and roughly 100 sold. Yet, every few years, a new edition would be put together and 10,000 new copies would be printed. Strolling around basements and subbasements, I saw enough paper to start a firestorm. Pallet after pallet was loaded with boxes and boxes of documents that nobody would read.

This being the decade of the environment (or was it the children?), I naively asked about the ITU recycling program. Needless to say, one didn't exist.

The Bruno experiment directly threatened this paper empire. The bureaucracy had framed its argument very cleverly. Every year, the ITU had received several million Swiss Francs in revenue from

Exploring the Internet

selling documents. The official "profit" from the Blue Book had gone to fund programs in the developing world.

In other words, Bruno was depriving the ITU of revenues that would fund vital infrastructure. My selfish little project meant that people who needed to call a doctor wouldn't be able to. Project Bruno, baby killer.

Nobody actually accused me explicitly of killing babies, but I certainly felt that undercurrent. After donating several months and several thousand dollars to putting ITU standards online, I had somehow not expected this type of reaction.

The donation was the single most difficult concept for the VIBs to understand. Why was I doing this? What was my motive? What was in it for me?

Of course, donations to the common infrastructure are how the Internet was built. Even formal standards bodies like the ITU run on donations. Corporations work in the standards process as a volunteer effort.

In the Internet community, volunteer efforts are the norm. The IETF has many people who attend as private citizens, paying for the privilege three times per year out of their own pockets. Paying for the privilege of getting the Blue Book online was not remarkable, but VIBs didn't know what to make of it.

The "Bruno, baby killer" aspect was a difficult one. Profits from document sales were virtual at best, and the simplest solution would be to redo the accounting system to look at the total costs of the inefficient document production process, but that proved to be dangerous.

I thus attacked the widespread unease with giving away copies on the Internet. Tony was advancing the novel theory that by giving copies away, you increased the market and thus increased sales. Such an argument, although bearing a few logical flaws, seemed to stop the VIBs, at least for a few minutes.

While the ITU was criticizing the Bruno experiment, they were attempting to move forward on their own electronic document handling system. Evidently, the Bruno situation had impressed the Secretary-General enough that he had presented himself at the meeting

Geneva

of the Information Systems Steering Group and suggested that they should do something.

The committee had thus spawned a task force. The task force had formed a small working group. Their initial inclination was to start using X.400 as a way to send out working documents, but only to members of the committees.

Tony and I tried valiantly to switch the focus to getting the ITU on the Internet and sending the documents out to as wide an audience as possible. This was not meeting with much success. The computer department was worried about the resource implications of such a move and wanted two additional staff members (in addition to their current network staff) to support the effort.

Basically, the bureaucracy desperately wanted to get back to a world they could control. In order to control documents, however, you need to own them. Nobody at the ITU wanted to admit that there was a possibility that the ITU didn't own its own documents.

Tony Rutkowski had made an analysis of the issue of copyright and had come to the conclusion that the ITU didn't have a sustainable basis for asserting copyright protection. Many of the other VIBs, however, felt that the issue was cut and dried.

There are no apparent legal cases in which somebody has challenged copyright on a standards document. There are many factors that must be weighed before a court will uphold a copyright claim, and it was naive to think that the issues are so simple that the ITU could confidently claim they would win in a court of law.

In order for a document to have a copyright applied to it, it must, among other criteria, be original and not previously published. Since almost all standards start out as public domain working documents, even this fundamental requirement is not often met.

Many jurisdictions do not allow protection to be granted on official or governmental works. Even a private standards body might be considered by the courts to be quasi-governmental. Many places, such as the U.S., make standards a procurement requirement, making copyright enforcement questionable at best.

Even if standards are copyrightable, only the representation of the standard, not the contents, can be protected. Tony's conclusion was that in almost any jurisdiction, running the paper through a

Exploring the Internet

scanner and OCR software and posting ASCII text would be defensible.

Many standards do have graphics, of course. The graphics have a stronger basis for copyright, since the representation is everything. As we had seen with the Blue Book, though, in many cases the graphics were not absolutely essential, at least for getting a rudimentary understanding of the standard.

With these factors in mind, Tony and I walked across the street to meet Larry Eicher, Secretary-General of ISO. My feeling was that even if there was no copyright on standards, it was certainly easier to work with ISO than against them.

The fact that Tony accompanied me was meant to send the message that my efforts enjoyed at least some support from the ITU. I brought along a copy of STACKS; Tony brought the slides from his presentation to INTEROP in which he concluded that it was unlikely that any standards organization could assert copyright on documents.

"Do you think that's diplomatic?" I asked.

"Nothing wrong with pushing forward the state of the art," he said with a smile.

We met Eicher and Mike Smith, one of the leaders of the task force which supports the OSI effort. Both turned out to be very reasonable people.

I gave a little speech about the moral necessity of disseminating standards. I advanced the view that the reason that OSI had taken so long to come to market was simply because it cost so much to find out about it.

We then started talking about applying Bruno to the ISO world. Eicher was quite frank: 25 percent of ISO revenues came from the sales of standards documents. How did I propose to replace that revenue? Even more importantly, ISO was controlled by its member organizations, which also made much money from standards sales. How did I propose to convince groups like ANSI that posting standards for free would help them?

Simply put, it was a question of financial survival. Interestingly enough, Eicher was clearly unwilling to argue his case on copyright grounds. When I ventured the theory that copyright protection for

Geneva

ISO documents was legally weak and that some radical might just go ahead and post the standards, Eicher said "it's not a question of copyright protection, it's a question of fair business practices."

We began searching for a potential solution. I proposed my high resolution/low resolution compromise. The plan would post low resolution versions of documents for free on the network and allow ISO and ANSI to continue to sell the high resolution versions, either on paper or electronically.

Low resolution might mean ASCII text and 200 DPI bitmaps of graphics, formulas, and other elements not well suited to representation as ASCII text. Some document format such as ODA could be used to tie the pieces together. Using ODA would help ISO by spurring the development of the standard by giving people a substantial base of documents worth reading.

The crucial assumption was that people with the free version would then pay for documents. I argued that free distribution of standards would increase the base of people who read documents by at least a factor of 10, maybe even more. Many of these would want the paper documents. Giving away standards would lead to increased revenues.

I then offered to test this theory on Bruno at no cost to ISO. Eicher agreed to at least consider a formal proposal, so we went back to the ITU and dashed off a formal letter. Kind of a long shot, I figured, but certainly a first step. (I never received a response to my letter, but that was no surprise. I did, however, publish the offer in *Communications Week* just in case ISO had misplaced my letter and needed a reminder.)

Tony and I had one more item of business to attend to. A seminar had been scheduled for Monday morning to give people at the ITU a briefing on Bruno. Tony had sent electronic mail on Friday, but by Monday, none of the mail had arrived and therefore nobody showed up at my lecture.

Turned out that the entire ITU mail system was running off a VAXmate with a very limited amount of memory. If you sent mail to everybody at the ITU (only a few hundred people), the system crashed.

Exploring the Internet

Rather than remove the offending mailing list or even move the message handling system up to an appropriate host, the issue of the "mail server situation" had entered the bureaucracy and a heated debate had begun, focusing on whether or not to expand memory on the VAXmate. I was a bit incredulous. Running a mail system for 900 professional bureaucrats off a VAXmate is kind of like using a Volkswagen Beetle to haul timber out of the Amazon jungle.

We rescheduled the seminar, booking Tarjanne's personal conference room for the occasion. The secretaries were worried that the room would be too small (it could only hold 50 people or so), but Tony and I insisted that the venue had the appropriate symbolism and that having to turn away people wouldn't be all bad.

Since electronic mail wasn't going to do the trick, we had only one alternative: ElevatorNET. Posting notices in the elevators was about the only effective means of communication at the ITU, the organization that invented X.400.

The ITU has some of the strangest elevator manners in the world. When you enter a lift, custom requires you to greet everyone. Everybody then choruses back a hearty "bon jour."

When you leave, you say goodbye and everyone responds with their own "au revoir." Nice custom, but what it means is that during a busy period, it can cost you a dozen hellos and goodbyes to go up just a few floors.

Tony and I pushed the button for the elevator and caught the first one, respecting the elevator protocol. Every time we caught an elevator, we rode a few floors, long enough to post the notice. We then got off, pushed the button again, and hoped that a different elevator would start. Finally, dozens of bon jours later, we had caught the last elevator and posted the last notice.

Wednesday morning, Tony had prepared all sorts of handouts for the eager crowds. I nervously sat in the corner and prepared my talk.

Two people came. One was our ally in the computer department, another a gentleman I had already briefed. Nobody else bothered to show up.

We all chatted for a few minutes, had a cup of coffee, then went back to Tony's office. I bid Tony goodbye.

Geneva

Thursday, I took a series of four flights from Paris to London to New York to Ithaca, home of Cornell University. Sitting on the planes, I had plenty of time to reflect in wonder at the ITU. Many people were grateful that Tony Rutkowski had put the time into the bureaucracy, but you had to wonder how somebody that talented could survive in such a labyrinth. (He didn't for very long. Tony is now an employee of Sprint International and Vice President of the Internet Society.)