Telecommunications Electronic Reviews (TER)

ter
telecommunications
electronic reviews

Volume 3, Issue 9, December, 1996

Contents

ter issues (/lita/publications/archive/ter)

REVIEW OF: J. P. Claude, ed. Advanced Information Processing Techniques for LAN and MAN Management. Amsterdam, NY: North Holland, 1994.

by Steve Cavrak

In the olden days, say a decade or so ago, local area networks (LANs) had narrow service goals. Perhaps a user could login to a host computer or catalog system, send files to a dot matrix printer, and maybe, if there were another host computer on the network, transfer files between them. If the network were advanced,

users might be able to send e-mail to each other, even if they were on different hosts! The now-retired BITNET (Because It's Time Network), founded in 1981, operated along those lines. It supported everything except remote login, and it did so on leased telephone lines running data at 9600 bits per second! Network management was almost totally concerned with the issue of connectivity--do we have a network? If you were connected, you had a network; if the network didn't work, you must have become disconnected.

Since then, of course, the picture has become more complicated, and is becoming more complicated even as this review is written. The global Internet is forever facing meltdown, and still more users are piling on and on. Under the umbrella of "digital convergence," cable-based television networks are gearing up to provide data, data networks are being used to carry voice and video sessions, and the plain old telephone companies are trying to offer video and data services. At the same time an emerging wireless industry is scheming to provide anything, anytime, anywhere. Network management begins to remind us of the old joke about government: "It takes angels to govern men, but if men were angels, they wouldn't need government."

In April 1993, before the World Wide Web became a household word, the IFIP TC6 (International Federation for Information Processing Technical Committee 6--on communications systems) organized a conference aimed at understanding how computational techniques could be applied to network management. The resulting conference, co-organized with Laboratoire MASI (Methodolgie et Architecture des Systemes Informatiques), Versailles, collected 21 papers.

The speakers identify a variety of advanced information processing techniques--automation of management functions, the use of reliable software engineering techniques (object oriented programming), distributed processing (e.g. local intelligence in devices, delegated management functions), human computer interface design (for configuration and diagnosis), knowledge-based systems (versus algorithmic designs), and advanced architectures (such as an "intelligent" network). Some of the papers are tutorial--discussing, for example, European experience in international network management, the OSI (Open Systems Interface) network management model versus the Internet management philosophy, a design of a telecommunication network management platform, the use of delegation methods. Others are presented in the context of specific network management projects such as the proposed tax collection network for Poland, an automated local area resource management project at the hosting Laboratoire MASI, an object oriented approach to managing a network gateway at TELECOM Paris.

Despite the technical nature of the conference, the proceedings are easy, though definitely not light, reading. Acronyms and jargon are generally defined anew in each presentation. The style is generally straightforward, and you don't have to be a computer programmer to understand what's going on.

These proceedings probably won't help the user whose new modem won't work on his new computer with the old university's dial-in lines and whose life seems to have been reduced to an interminable game of phone tag with different vendors in different time zones. In the daily scheme of things, this concern with advanced techniques might seem beside the point--any information processing at all would be an improvement. For those somewhat more behind the scenes, maybe even those sitting behind a help line or reference desk, the conference proceedings provide evidence that some progress is being made to make network life easier. It wasn't really all that long ago, for example, that making a long distance call was a major production; yet today, it's all mostly invisible-- except on Mother's Day. Someday, in the digital future, everything will be just a mouse click away.

Steve Cavrak (STEVE.CAVRAK@UVM.EDU (mailto:STEVE.CAVRAK@UVM.EDU)) is with Academic Computing Services in Computing and Information Technologies at the University of Vermont. He is a contributor to the upcoming ALA Editions publication, The Cybrarian's Manual.

REVIEW OF: Mike Hendry. Practical Computer Network Security. Norwood, MA: Artech, 1995.

by Shawn M. Collins

Early in Mike Hendry's book on this very popular topic, its orientation becomes clear. Network security, for the purposes of this work, means more than simply creating firewalls against intrusion; it also means the creation of procedures and protocols that allow users of a network to be secure about the state of their data--its integrity and confidentiality.

Directing the book to end users, managers, and mid-level technicians, Hendry strives to balance technical discussions of procedures and technologies with a clear introduction to the issues involved in creating a secure network environment which can serve as a guide to those who must make sensible decisions about investment in and deployment of security measures. The scope of its coverage, which falls between beginner and intermediate, makes it not only an interesting introduction to the topic of security in networked computing systems, but also an introduction to the wide variety of environments in which computer security is a concern. However, that same breadth of scope may make it unsuitable for those needing in-depth coverage of specific topics. The book is constructed around three parts, "Requirements and Risks," "Technology," and "Applications." In the first of these, Hendry outlines the broad range of issues that must be considered when developing a security plan, and emphasizes attention to the costs and benefits of managing security risks. In a chapter on the preconditions of a secure environment, he suggests that the business process be examined closely to see what impact it has on security.

In this chapter, he asks a question which, while it might make some Information System managers a little nervous, bears asking: do the security needs of the business require that management of the computer systems be in house, or should they be contracted out? While this may seem to offer fuel to an outsourcing-happy manager, Hendry makes clear that in outsourcing computer services, the culture of both the contracting company and the contractor must be such that security will be maintained. Relinquishing maintenance of security to a firm which does not hold it as dearly as you would can spell disaster, and Hendry suggests that security protocols be clearly spelled out in any contract. In a similar vein, Hendry addresses what he refers to as organizational integrity. While it may seem odd to discuss the ethics of the Chief Executive Officer, or the relationships of workers in the context of computer network security, Hendry's practical approach to security excludes no element of the environment in which the network exists. As he suggests, if the boss, or the employees are crooks, you've got a real security problem.

Similarly, problems can be generated by an environment in which computer expertise cannot speak to the problems of the business process, and those managing the business process cannot comprehend the nature of computing. After a discussion of the need for a clear understanding of all the objects at risk, the directions from which those risks can come, and the costs of protection, Hendry moves into a discussion of hardware, software, and networks, and offers a clear statement of the raison d'etre of the book.

[Security] problems are aggravated by the barrier that still exists between the computer people and the rest of the world. Many decisions that properly relate to the running of the business are left to computer staff whose knowledge of the business is at best imperfect. Or managers must make decisions concerning the selection of hardware, software, or encryption tools with only the sketchiest understanding of the relative merits of the available products. It remains to be seen whether a fresh computer-literate generation will still suffer from these problems. (p. 33)

This statement serves as an opening to a discussion of issues which might bridge some of the gaps it describes. Risks to physical hardware, as well as the complexities introduced by software are laid out in terms which relate business processes to computing processes. Viruses, software licensing, supplier selection, and problems specific to applications like databases and communications are discussed, with a focus on points of risk which are introduced by different applications and technologies. Application issues in environments ranging from real-time control systems in manufacturing to banking and the military are touched on here, and this broad scope characterizes the rest of the book.

The second part of the book is focused on technologies, and includes a nice introduction to the principles behind, and the technologies that implement, encryption. A discussion of encryption keys and key management is followed by a discussion of hardware oriented systems such as smart cards, encrypting modems, and key generators or "guns." These tools are then compared with available software tools and the risks involved in their use. In keeping with earlier portions of the book which focus more on principles than on technical specifics, these discussions are oriented toward making clear the issues involved in use of each tool, and exploring the "risk calculus" which has to be applied before making a selection decision.

The third section of the book focuses on applications, and explores the manner in which various kinds of transactions and data have different security needs and implications. Banking and financial systems, online subscription services, software distribution systems, and satellite telemetry, monitoring, and control are all covered in this section, with a distinct international perspective. While the scope of this section is so broad it can be difficult to read, with a great deal of information unlikely to serve most readers' needs, it does offer enough information to make it potentially useful in a classroom environment.

I read this work at the same time that I was reading Internet Firewalls and Network Security, from New Riders Publishing. While this other book offered none of the technical implementation details of Hendry's work, it did a much better job of outlining the concept of computer network security. I found it refreshing to see a book with this title that dealt so effectively with the psychological needs of network users to know not only that their data is safe from intruders, but also that it is safe from general blunders and glitches.

Mike Hendry's Practical Computer Network Security, however, does accomplish the task it set for itself, in a readable and informative manner. With an index, glossary, bibliography, plenty of helpful graphics, and a detailed table of contents, it could serve as a decent reference work for those seeking concise descriptions of specific concepts like public-key encryption. It also serves well as an introductory text for those needing to learn about the issues involved in creating a secure computer network environment. At several points in the work, Hendry makes it clear that he sees the book as a tool for those needing to make informed decisions which balance the costs and benefits of security implementation protocols, and it can serve well in this role. In the hands of a good instructor, it could serve as a fine text for a course. In the hands of a manager needing a better understanding of the proliferation of security threats and solutions being discussed in the computing press, it can serve as a readable introduction to a complex topic.

Shawn M. Collins (scollins@utk.edu (mailto:scollins@utk.edu)) is Computer Services Coordinator of the School of Information Sciences at the University of Tennessee.

table of contents | ter issues (/lita/publications/archive/ter)

REVIEW OF: Rolf Oppliger. Authentication Systems for Secure Networks. Norwood, MA: Artech House, 1996.

by Thomas Dowling

Despite the title, "Authentication Systems for Secure Networks" does not contain much general information on network authentication, and still less on broader security issues. Instead, it is a highly technical exploration of six authentication systems, or authentication built into larger security packages: Kerberos, in both versions 4 and 5; NetSP, developed by IBM; Digital Equipment Corporation's SPX (pronounced "Sphinx" and different from the Sequenced Packet Exchange used in IPX, or Internetwork Packet Exchange, environments); TESS, The Exponential Security System; the European Computer Manufacturer Association's Project Sesame; and the Open Software Foundation's Distributed Computing Environment, or OSF DCE. An opening section introduces some of the technical aspects of cryptography and security, and a concluding section compares some of the features of the authentication systems described.

Oppliger's book is clearly not intended as a general guide to security, and it will not help casual readers who want to set up PGP (Pretty Good Privacy) or pick a good password. It is intended for technical staff at organizations running, or considering, one of the packages above. A systems administrator at an institution planning to implement Kerberos or DCE will undoubtedly find it helpful, but the benefits to more general readers are hard to place.

The introductory chapter is certainly a rich source of information on the technical cryptography environment, including a thorough rundown on the relevant standards and specifications from the ISO (International Standards Organization), IEC (International Electrotechnical Commission), ITU (International Telecommunications Union), and National Institute of Standards and Technology. It also describes some features of a full-fledged security environment that will not be immediately apparent to newcomers, such as secure methods for distributing private keys and the security mechanisms enumerated in the OSI (Open Systems Interface) security architecture. [1] By way of suggesting the technical level of this book, the introduction has a 15 item bibliography, from sources like Bellcore technical reports, Proceedings of the IEEE (Institute for Electrical and Electronics Engineers), and the ACM (Association for Computing Machinery) Operating Systems Review.

A casual survey suggests that, of the systems described in the book, Kerberos is the subject of greatest interest and use. Kerberos, which developed as a means to provide security on the Massachusetts Institute of Technology's Athena network, is a pervasive security mechanism which typically requires customized (or "Kerberized") network applications. Using an IP (Internet Protocol) network with Kerberos would necessitate a Kerberized Telnet application, and possibly Kerberized Web servers, Web browsers, e-mail clients, and

other applications. Kerberos also requires a substantial investment in servers to distribute encryption keys and session tickets. Oppliger does a very thorough job of describing exactly what protection a network receives in exchange for this investment--even more helpful, he does a thorough job in describing the criticisms of, and known weaknesses in, both Kerberos V4 and V5.

Having devoted more than a third of the book to the introduction and the discussion of Kerberos, Oppliger returns to that material as a building block for the other chapters. To some extent, that means that a reader's understanding of OSF DCE, for example, will have a lot to do with its similarities and contrasts with Kerberos. While this may lead to a certain amount of imbalance in the book's coverage, it is a fair reflection of the current authentication discussion at many institutions (perhaps more so in the academic world than the corporate world): name brand recognition suggests that once a site determines a need for secure authentication, the next question is not so much "What system do we use?" but "Is there a reason not to use Kerberos?" Within that context, it makes sense to bring all other systems around to a comparison with Kerberos.

Throughout the book, Oppliger maintains a high level of technical detail and thoroughness. The limited audience this book aims for will find it valuable and very informative.

Notes:

[1] By "OSI security architecture," Oppliger refers to the material in the ITU X.800 and ISO 7498-2 standards. He also makes note of authentication issues in ITU X.509 and key distribution issues in IEEE 802.10 and ANSI (American National Standards Institute) X9.17.

Thomas Dowling (TDOWLING@OHIOLINK.EDU (mailto:TDOWLING@OHIOLINK.EDU)) is Assistant Director of Systems, Client-Server Applications, at OhioLINK.

table of contents | ter issues (/lita/publications/archive/ter)

Growing Up in the Computer Age

by Thomas C. Wilson

I would like to start by saying that I believe that Netscape is in many ways an incredible product and an amazing company. But it is not the only incredible company with a series of products. The recent history of microcomputing suggests that they rank with many other companies that have bright ideas and the energy and financial backing to pursue them.

That having been said, Netscape also provides an illustration of the youthful arrogance, naivete, and bullying nature that has pervaded the computing industry since the mid-1970's; these traits also characterize the analysts that report on happenings in the industry. The Wall Street Journal recently reported that one analyst believed that Netscape was growing up because they are now pursuing competition head-to-head with Microsoft Windows on the desktop, instead of just in the browser arena. [1]

I previously wrote in this publication (http://gold.ohiolink.edu/ter/ter-3-2.html#wilson (../../../../../../../ter-3-2.xml#wilson)) about my concern that Netscape would get side-tracked if they continued to pursue the idea of creating yet another operating system. Lesson one to indicate growing up in this case would be careful consideration of the words used to convey meaning, instead of falling prey to the purposefully ambiguous language of marketing. What Netscape really appears to be doing--or at least what can be read between the lines--is creating another "shell" for the Windows environment. If that is the case, such a momentous announcement is a bit overstated, as there have been multiple shells (i.e., replacements for Program Manager) for the Windows environment for years. In fact, with Windows NT 4.0, Microsoft allows the user within the product itself to select either the Program Manager or the Win95 interface. Perhaps Constellation, Netscape's up-and-coming product, will eventually be another interface a user could select--but it would not be another operating system.

Netscape's attempt to break Microsoft's hegemony on the desktop is laudable in a culture that glorifies competition. But then again, what is real competition? Does the user gain anything by trading one master for another--except perhaps a younger generation at the helm? Certainly Microsoft has for years thrown its weight around to influence the industry and users in ways that positively affect its bottomline. Do you honestly think that Netscape doesn't or wouldn't?

Perhaps one could see a benefit in submitting to a master that cuts the most efficient pathway to the ultimate goal--a sleek cross- country runner rather than a bulky linebacker--if one believes that we are on a path that inevitably leads to a Web interface and metaphor for all computing. In this context, however, there are two major difficulties with the ideas expressed in the clauses of the previous sentence: (1) the notion of efficiency for the long- term is antithetical to how things occur in the American computing industry, and (2) the definition of progress implies a misguided, unquestioned belief.

First, I would challenge any reader to identify and document any significant American company in the computing industry that has been able to remain lean and mean for at least ten years after hitting it big. The problem: the American business culture assumes that a company must grow or die; it cannot remain constant. With growth in market share comes growth in size and the concomitant complexity of managing large organizations. Thus, it is more likely that a company will grow fat and sassy rather than remain lean and mean. In other words, efficiency is fleeting. There would not be a barrage of literature and seminars on down-sizing, right-sizing, and re-engineering businesses if this problem were not common.

The second difficulty mentioned above actually constitutes a much more far-reaching conceptual challenge, not just to computer companies, but to every user. The belief that we are on some progressive path that necessarily leads to one ultimate goal not only can be dangerous in terms of current purchasing or implementation decisions, but also can produce intellectual laziness in considering strategic planning and decisions. Not questioning our own assumptions, or those of the players in the computing industry, including company officers, pundits, and advertising firms, weakens our ability to strike a path that represents and facilitates the interests of our constituents.

One of the most remarkable examples of an unquestioned path is the pursuit of a single interface for all computing--Constellation or any other. Many people believe that is really what we need. Why? Are we prepared to suggest that everyone can absorb information in exactly the same fashion? Or that the same individual always wants information and the machine interface to look and feel the same for all purposes? Or that all types of data should be displayed similarly? I think not. Many people also believe that the Web is giving us the ability to create such an environment. Talk about hegemony!

What the Web is giving us--besides an opportunity to consume donuts and coffee while we wait for a response--is the possibility of building platform-independent, distributed applications and documents, not a ubiquitous interface. Please note that I used the word "possibility"; the Web does not require that any of us do this. We are not given a single interface with the Web--have you visited any sites recently? Each site is different in design and content; and that is the flexibility that we seem to want in the aggregate. One could easily argue that the interest in Java and ActiveX illustrates this point: Designers want to be able to create and control unique environments.

There may be good reasons for selecting a Web-related interface for some applications, but let's not use the single interface argument any longer. It's not becoming to us--nor is it intellectually honest.

When analysts speak of growing up and define that process in terms of Netscape competing more directly with Microsoft, I am disappointed. [2] I've come to see growing up as making peace with ambiguity and realizing that there are few truly new things under the sun. When we succumb to youthful thoughts, our expectations may become unrealistic.

Here's to a new year filled with grown-up computing!

Notes:

[1] Rigdon, Joan I. & Bank, David. (1996, November 20). Netscape Aims to Nab Desktop from Microsoft. The Wall Street Journal, p. B6.

[2] If you find yourself irritated that I have used Netscape as an illustration in this manner, feel free to substitute the name of your favorite (or not-so-favorite) computer company for each occurrence of "Netscape."

Tom Wilson, Editor-in-Chief TER, TWilson@uh.edu (mailto:TWilson@uh.edu)

table of contents | ter issues (/lita/publications/archive/ter)

About TER