<<TCP/IP详解 卷1：协议（英文版· >

图书基本信息

书名：<<TCP/IP详解 卷1：协议（英文版·第2版）>>

13位ISBN编号：9787111382287

10位ISBN编号：7111382285

出版时间：2012-5

出版时间：机械工业出版社

作者：（美）Kevin R. Fall,（美）W. Richard Stevens

页数：1017

内容概要

《TCP/IP详解》是已故网络专家、著名技术作家W. Richard Stevens的传世之作，内容详尽且极具权威，被誉为TCP/IP领域的不朽名著。

本书是《TCP/IP详解》的第1卷，主要讲述TCP/IP协议，结合大量实例讲述TCP/IP协议族的定义原因，以及在各种不同的操作系统中的应用及工作方式。
第2版在保留Stevens卓越的知识体系和写作风格的基础上，新加入的作者Kevin R. Fall结合其作为TCP/IP协议研究领域领导者的尖端经验来更新本书，反映了最新的协议和最佳的实践方法。
首先，他介绍了TCP/IP的核心目标和体系结构概念，展示了它们如何能连接不同的网络和支持多个服务同时运行。
接着，他详细解释了IPv4和IPv6网络中的互联网地址。
然后，他采用自底向上的方式来介绍TCP/IP的结构和功能：从链路层协议（如Ethernet和Wi-Fi），经网络层、传输层到应用层。

书中依次全面介绍了ARP、DHCP、NAT、防火墙、ICMPv4/ICMPv6、广播、多播、UDP、DNS等，并详细介绍了可靠传输和TCP，包括连接管理、超时、重传、交互式数据流和拥塞控制。
此外，还介绍了安全和加密的基础知识，阐述了当前用于保护安全和隐私的重要协议，包括EAP、IPsec、TLS、DNSSEC和DKIM。

本书适合任何希望理解TCP/IP协议如何实现的人阅读，更是TCP/IP领域研究人员和开发人员的权威参考书。
无论你是初学者还是功底深厚的网络领域高手，本书都是案头必备，将帮助你更深入和直观地理解整个协议族，构建更好的应用和运行更可靠、更高效的网络。

作者简介

Kevin R.
Fall博士有超过25年的TCP/IP工作经验，并且是互联网架构委员会成员。
他是互联网研究任务组中延迟容忍网络研究组（DTNRG）的联席主席，该组致力于在极端和挑战性能的环境中探索网络。
他是一位IEEE院士。

W. Richard
Stevens博士（1951—1999）是国际知名的Unix和网络专家，受人尊敬的技术作家和咨询顾问。
他教会了一代网络专业人员使用TCP/IP的技能，使互联网成为人们日常生活的中心。
Stevens于1999年9月1日去世，年仅48岁。
在短暂但精彩的人生中，他著有多部经典的传世之作，包括《TCP/IP详解》（三卷本）、《UNIX网络编程》（两卷本）以及《UNIX环境高级编程》。
2000年他被国际权威机构Usenix追授“终身成就奖”。

书籍目录

章节摘录

版权页： 插图： 1.1.2 The End-to-End Argument and Fate Sharing When large systems such as an operating system or protocol suite are being designed, a question often arises as to where a particular feature or function should be placed. One of the most important principles that influenced the design of the TCP/IP suite is called the end-to-end argument （SRC84）: The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the end points of the communication system. Therefore, providing that questioned function as a feature of the communication itself is not possible. （Sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.） This argument may seem fairly straightforward upon first reading but can have profound implications for communication system design. It argues that correctness and completeness can be achieved only by involving the application or ultimate user of the communication system. Efforts to correctly implement what the application is "likely" to need are doomed to incompleteness. In short, this principle argues that important functions （e.g., error control, encryption, delivery acknowledgment） should usually not be implemented at low levels （or layers; see Section 1.2.1） of large systems. However, low levels may provide capabilities that make the job of the endpoints somewhat easier and consequently may improve performance. A nuanced reading reveals that this argument suggests that lowlevel functions should not aim for perfection because a perfect guess at what the application may require is unlikely to be possible. The end-to-end argument tends to support a design with a "dumb" network and "smart" systems connected to the network. This is what we see in the TCP/IP design, where many functions （e.g., methods to ensure that data is not lost, controlling the rate at which a sender sends） are implemented in the end hosts where the applications reside. The selection of which functions are implemented together in the same computer or network or software stack is the subject of another related principle known as fate sharing （C88）. Fate sharing suggests placing all the necessary state to maintain an active communication association （e.g., virtual connection） at the same location with the communicating endpoints. With this reasoning, the only type of failure that destroys communication is one that also destroys one or more of the endpoints, which obviously destroys the overall communication anyhow. Fate sharing is one of the design philosophies that allows virtual connections （e.g., those implemented by TCP） to remain active even if connectivity within the network has failed for a （modest） period of time. Fate sharing also supports a "dumb network with smart end hosts" model and one of the ongoing tensions in today's Internet is what functions reside in the network and what functions do not. 1.1.3 Error Control and Flow Control There are some circumstances where data within a network gets damaged or lost. This can be for a variety of reasons such as hardware problems, radiation that modifies bits while being transmitted, being out of range in a wireless network, and other factors. Dealing with such errors is called error control, and it can be implemented in the systems constituting the network infrastructure, or in the systems that attach to the network, or some combination. Naturally, the end-to-end argument and fate sharing would suggest that error control be implemented close to or within applications. Usually, if a small number of bit errors are of concern, a number of mathemati cal codes can be used to detect and repair the bit errors when data is received or while it is in transit （LC04）. This task is routinely performed within the network. When more severe damage occurs in a packet network, entire packets are usu ally resent or retransmitted. In circuit-switched or VC-switched networks such as X.25, retransmission tends to be done inside the network. This may work well for applications that require strict in-order, error-free delivery of their data, but some applications do not require this capability and do not wish to pay the costs （such as connection establishment and potential retransmission delays） to have their data reliably delivered. Even a reliable file transfer application does not really care in what order the chunks of file data are delivered, provided it is eventually satis fied that all chunks are delivered without errors and can be reassembled back into the original order. As an alternative to the overhead of reliable, in-order delivery implemented within the network, a different type of service called best-effort delivery was adopted by Frame Relay and the Internet Protocol. With best-effort delivery, the network does not expend much effort to ensure that data is delivered without errors or gaps. Certain types of errors are

usually detected using error-detecting codes or checksums, such as those that might affect where a datagram is directed, but when such errors are detected, the errant datagram is merely discarded without further action.

媒体关注与评论

“我认为本书之所以领先群伦、独一无二，是源于其对细节的注重和对历史的关注。
书中介绍了计算机网络的背景知识，并提供了解决不断演变的网络问题的各种方法。
本书一直在不懈努力以获得精确的答案和探索剩余的问题域。
对于致力于完善和保护互联网运营或探究解决长期存在问题的可选方案的工程师，本书提供的见解将是无价的。
作者对当今互联网技术的全面阐述和透彻分析是值得称赞的。
”—Vint Cerf, 互联网先驱对本书第2版的评论： 本书第1版自1994年出版以来，深受读者欢迎。
但是时至今日，第1版的内容有些已经比较陈旧，而且没有涉及IPv6。
现在，这部世界领先的TCP/IP畅销书已经被彻底更新，反映了新一代基于TCP/IP的网络技术。
这本书仍保留了Stevens卓越的写作风格，简明、清晰，并且可以快速找到要点。
这本书虽然超过一千页，但是并不啰嗦，每章解释一个协议或概念，复杂的TCP被分散到多章。
我很欣赏本书的一个地方是每章都描述了已有的针对协议的攻击方法。
如果你必须自己实现这些协议，并且不希望自己和前人一样遭受同样的攻击，这些信息将是无价的。
这本书是日常工作中经常和TCP/IP打交道或进行网络软件开发的人必需的，即使你的工作并不基于IP协议，这本书仍然包含很多你可以用到的好想法。
”——摘自Amazon读者评论对本书第1版的赞誉： 这本书必定是TCP/IP开发人员和用户的圣经。
在我拿到本书并开始阅读的数分钟内，我就遇到了多个曾经困扰我的同事及我本人许久的难题，Stevens清晰和明确的阐述让我豁然开朗。
他揭秘了此前一些网络专家讳莫如深的许多奥妙。
我本人参与过几年TCP/IP的实现工作，以我的观点，这本书堪称目前最详尽的参考书了。
——Robert A. Ciampa，3COM公司网络工程师 《TCP/IP详解 卷1》对于开发人员、网络管理员以及任何需要理解TCP/IP技术的人来说，都是极好的参考书。
内容非常全面，既能提供足够的技术细节满足专家的需要，同时也为新手准备了足够的背景知识和相关注解。
——Bob Williams，NetManage公司营销副总裁

编辑推荐