Mißbach, Gibbels, Karnstädt, Stelzel, Wagenblast

Adaptive Hardware Infrastructures for SAP®

SAP PRESS

Contents

Foreword

In safeguarding their competitiveness, many companies today face the challenge of having to constantly adapt their business processes to the ever-changing requirements of the market. The main issues that companies must address are the increasing internationalization of markets and production locations, the creation of new partner networks (due to decreasing vertical integration within companies, for example), and the timely introduction of new products and services.

With the continuing development of its software solutions—from R/2 and R/3 to the multi-tier, Internet-enabled architecture of the mySAP Business Suite—SAP makes a significant contribution to helping companies meet this challenge. In particular, SAP's latest measure, in which it brought together its existing solutions in a Web-based services architecture, was intended to address the aforementioned needs for adaptation. This new architecture enabled the combining of individual function modules (for example, a module for creating a customer order) in order to create new solution modules and therefore support changing business processes quickly and cost-effectively while maintaining high quality.

Today, IT managers in most companies are faced daily with the question of how to reconcile the need for cost-cutting in the implementation and operation of SAP solutions with the demand to make these solutions more flexible. The most frequently asked questions are:

- What significance do the changing SAP solution architectures have for underlying IT infrastructures?
- What does an SAP IT infrastructure have to do in order to be able to adapt quickly to changing requirements?
- How do we ensure that issues such as high availability, performance, and scalability are not neglected?
- What can we expect from operating these 'adaptive SAP solutions' in terms of their running costs?

This book provides real-world answers to these questions. The team of authors, under the leadership of Dr. Michael Mißbach, Senior Consultant at Hewlett-Packard (HP), has several years' combined experience dealing on a daily basis with international clients and various SAP development and consulting departments in solving similar questions. His wealth of experience is reflected in this book, which focuses on current developments in platform and network technologies, and on technologies and concepts for operations optimization such as virtualization, IT Service Management (ITSM), and Total Cost of Ownership (TCO).

Moreover, this book highlights the close cooperation that has existed between SAP and HP for 15 years now, both in various technological developments and the daily joint support of thousands of clients worldwide.

Dr. Wolfgang Oskierski
SAP Business Manager EMEA

Foreword

The argument for using IT to support enterprise processes is stronger than ever today.

While managers in enterprises usually have little or no understanding of the concepts and specifics of information technology, they will still articulate what they expect from IT: to provide a measurable business value. However, they do so in their own specific "business-language" vernacular. The ability to translate these expectations expressed in business-speak into IT requirements is an art in itself; however, it is one that we must master if we want to position technology in its rightful place in our enterprises.

A situation where employees have to spend an indefinite amount of time in front of a static input screen to search for the information they require (and possibly never find it) is not the way to demonstrate the efficiency that is always demanded of IT in day-to-day business operations. A far better way to show the usefulness and value of IT is to have the interface open up a communications portal that can adapt itself to very specific individual requirements. SAP refers to it as a "role-based user interface" and it is considered to be a result of an Enterprise Services Architecture (ESA). This kind of architecture clearly demonstrates the value proposition of IT, because it enables users to access the information they require quickly and directly.

That's the theory.

For this kind of application layer to be feasible in the real world, a stable and solid basis is a must. SAP has its own name for this too—SAP NetWeaver—which is the foundation of the ESA structure. SAP NetWeaver forms a basis that consolidates and provides information about the people, information, and processes in an enterprise. Only if the basis is strong enough does the added value of the overlying application and communication landscape become apparent. An added difficulty is that besides being solid, the basis must be also highly flexible—two factors that may initially appear to be mutually exclusive.

In business-speak, this means that in the coming years we'll see an increased demand for solutions that promise quick interchangeability and ease of integration into existing landscapes. Shorter roll-out times and lower process costs will be expected as well, and the demand for contributions from IT solutions to value creation and the fulfillment of enterprise strategies will be, in a word, uncompromising.

This book is intended to help you lay the foundation for successful process design and visible added value in your enterprise. It will assist you in mapping out your own strategic path, along which enterprise visions can become a reality. Economizing in the wrong places in this context simply provides fertile ground for risks and errors rather than value creation, and this book thus makes a significant contribution to the success of ESA strategies.

May you have many hours of enjoyment reading this book and may you have success in creating solid solutions that bridge the gap between IT and business.

Andreas Kerbusk

Chairman of the German SAP User Group (DSAG e.V.)

Introduction

Adaptive IT infrastructures for agile enterprises

While most IT projects in recent years have been dominated by the need to cut costs, today enterprises are enhancing their competitiveness by using IT to adapt their business processes to markets that are changing evermore rapidly.

This is because thinking purely in terms of costs makes sense only up to a point. Without any doubt the cheapest IT system is the one which is completely written of and is running without any change, however continually changing market conditions mean that enterprises are forced to modify their processes, including their IT, on an ongoing basis. Therefore, enterprises have a clear competitive advantage if their IT is flexible enough to be able to implement new processes quickly.

However, it is not just the markets that are changing. Company mergers and acquisitions indicate that enterprises themselves are becoming increasingly agile. For example, steel companies can turn into mobile telephone operators or tourism enterprises, and it is not always the big fish that swallow the small fish. In these cases, too, it is the speed at which the IT infrastructure can be adapted to the new circumstances that determines the success or failure of an enterprise.

SAP NetWeaver and ESA

With SAP NetWeaver, and especially the Enterprise Services Architecture, SAP has developed products and concepts that will have a dramatic effect on how IT is used to benefit enterprises. These concepts form the basis for quickly introducing and adapting business processes. However, for this to be possible, enterprises need an infrastructure that is as adaptive as the software, while still providing a stable technology basis.

Technologies for adaptive infrastructures are not a completely new concept. For some time, many hardware manufacturers have been working on consolidation concepts that include the "virtualization" of server and storage resources.

SAP Adaptive Computing Controller

One new concept, however, is the idea of incorporating the application, which creates an holistic, all-around solution. With the introduction of SAP Adaptive Computing, this kind of solution is beginning to make its mark. For the first time, the SAP Adaptive Computing Controller provides an interface between the infrastructure and the application, and therefore implements the preliminary steps toward integrating two previously separate worlds.

Availability The increasing predominance of SAP solutions in enterprise processes proves that enterprises are becoming increasingly more dependent on these systems. High availability is therefore gaining more importance than ever before. Because business processes in SAP NetWeaver environments are distributed across several SAP systems, the processes can work only if every system is functioning perfectly. Similarly, Enterprise Services Architectures (ESA), which can integrate functionalities from different systems in an overlying business process in a very short space of time, can work only if all the connected systems are equally available.

However, high availability does not depend on technology alone. Besides unified management and monitoring tools, running adaptive infrastructures also requires the consistent and carefully planned use of IT Service Management (ITSM) methods.

Adaptivity After many discussions with customers, the authors have learned that there is a demand for a guide to adaptive infrastructures. However, everyone has his or her own idea about the meaning of "adaptive"—from simple load distribution mechanisms to the virtualization of whole data centers or automatic recognition and monitoring of resources. Intrinsic to all definitions is the goal of enabling infrastructures to adapt flexibly to business requirements. This book provides an outline of all the aforementioned topics.

The great demand for our two existing books, which deal with the infrastructure and the operation of SAP systems, encouraged the authors to produce this third book, which deals with the latest developments and challenges and their corresponding solutions. We also outline trends for the future, where appropriate. Therefore, this book is an addition to the books we previously authored, and not simply an update.

To ensure that readers who do not regard themselves as "gurus" in each area can still derive benefit from reading this book, the most important technologies are explained in detail. Practice-oriented guidelines are provided throughout the book in order to make the reader aware of essential but often less obvious facts. The book concentrates exclusively on the technical aspects of IT infrastructure; the details of how to install and adapt SAP software to business processes are beyond the scope of this book.

The solutions presented here refer to the latest releases at the time of printing. However, although the laboratories are constantly producing new hardware and software, the underlying technologies and architectures change much more slowly, and so the concepts presented here can be used on a long-term basis. Also, many of the technical solutions presented here are also suitable for other enterprise-critical software systems.

The book has intentionally taken a neutral stance in terms of manufacturers. Nonetheless, the authors are employees of the SAP HP International Competence Center and their expertise is largely based on the numerous solutions that have been developed there since 1989. For this reason, SAP and HP products, and solutions from partner companies, of which the authors have positive experiences, are used as examples for a class of solutions. However, references to any product do not represent an evaluation of that product.

Structure of the Book

Each chapter of this book contains a short introduction that outlines the goals of that particular chapter. Wherever possible, the detailed descriptions of the solutions are addressed in terms of performance, availability, and flexibility, and are illustrated by real-world examples. The closing section of each chapter consists of a short summary of the main recommendations.

The first two chapters provide a brief description of the functionalities of the SAP software components and of the underlying technical components involved in each case from an IT point of view to establish the foundation for the subsequent chapters. **Chapter 1** gives an overview of the functionality and use of the SAP NetWeaver components. It also describes the software lifecycle of the SAP solution and user management solutions.

Chapter 2 presents the solutions of the mySAP Business Suite and places them in the context of the Enterprise Services Architecture.

Chapter 3 introduces the architecture of the SAP Web Application Server, which is a platform for the process execution of almost all SAP solutions. In doing so, it deals with the ABAP and the Java stack and considers the various aspects of Unicode implementation. Lastly, it looks at aspects of grid computing.

Chapter 4 deals with dimensioning computer systems. The focus here is on explaining the most important parameters used for designing a hardware landscape. This chapter ends with an examination of the level of exactness that can be achieved with a standard approach to sizing.

Chapter 5 presents computer systems for SAP applications. It describes the available technologies with particular emphasis on the design of the processor and main memory, and describes the advantages and disadvantages of blade concepts. It also presents the various operating systems, with a special focus on Linux.

Chapter 6 deals with the disk storage sub-systems of the SAP database server. It describes the various files in an SAP system and explains how to dimension and structure storage sub-systems. It also deals with how Network Attached Storage (NAS) can be used with SAP.

Chapter 7 tackles the subjects of availability and downtime in integrated systems of mission critical systems. All aspects of high availability are presented here, from protecting a computer center from disasters, to cluster technologies and shadow databases, to system operation.

Chapter 8 describes the various user interfaces of the SAP solutions, as well as print and output management solutions. It also presents the new Interactive Forms and SAP Web Dynpro.

Chapters 9 to 12 deal with network infrastructures for SAP system landscapes. The specific requirements that an SAP system has in terms of network bandwidth and latency, as well as the specific aspects of local-area and wide-area networks are each dealt with in their own separate chapters. Another chapter deals with protecting enterprise data and systems from unauthorized access via the Internet.

Chapter 13 shows how the virtualization technologies presented in the previous chapters can be used to build flexible SAP infrastructures. It also describes the SAP Adaptive Computing Controller and 10 different application scenarios for adaptive infrastructures.

Chapter 14 addresses the management of Enterprise Services Architectures. It also presents the IT Service Management (ITSM) reference model and the various management system concepts for monitoring, analyzing, and optimizing the infrastructure. An example of a vendor-managed inventory scenario illustrates the various points.

Chapter 15 explores the cost aspects of operating an SAP system. It briefly describes the most common scientific models and a model for structuring the overall operating costs. Lastly, using two practical examples, it discusses the costs and benefits of a system integration scenario, and compares the operating costs of a scale-up concept with those of a scale-out concept.

Acknowledgements

This book is the product of voluntary work done in our free time during many nights and weekends. We therefore dedicate it to our wives and children, who have had to spend a lot of time without our undivided attention.

We would also like to thank all the customers and colleagues who selflessly provided much help in the form of tips, contributions, and constructive criticism. Without their support, we would not have been able to write this book. In particular, we would like to mention Helmut Fieres, Sebastian Buhlinger, Marina Marscheider, Roland Wartenberg, and Markus Meisl at SAP; Monika Reitmeier and Michael Weber at Munich Re-Insurance; Friedel Manus at Capgemini, Paul Hammersley at EPI-USE Limited and Rob de Maat and Peter van Eijk at Deloitte Consulting; Uwe Hoffmann at Microsoft; Gerd Kammerath at Citrix; Nils Bauer from Network Appliance, Lothar Zocher and Andreas Epple at EMC; Dan Ellenbogen at Spaceline; Christian Schult at Norasia; and Andreas Schweizer, Jens-Uwe Walther, and Horst Jacobi at Carl Zeiss Jena (now HP Managed Services) and Andreas Kerbusk at STEAG.

Thanks are also due to those colleagues who patiently answered so many of our questions: Eric Martorell, Gene A. Burke, Laurie Ford, Chuck Desostoa, Fanny Osorio, Filip Van Grembergen, Michael Wiseberg at HP USA and Canada, Nigel Edwards at HP Labs Bristol, Heiderose Doms, Bernd Klopsch, Carsten Helmers, Jörg Schade, Mike Wenner, Peter Weiler, and Friedrich Kilian at the SAP HP Competence Center, Werner-Wolfgang Gaertner, Georg Storz, and Rudi Grom at the HP SAP Center of Excellence, David Adelmann, Arne Hartmann, and Holger Zecha at HP Managed Services, Peter Holzmann at HP Service, Michael Igel at LinuxLab, our Microsoft experts Horst Kanert and Erik Rieger, Dr. Christoph Balbach at HP Storage, Thilo Domsdorf, Engelbert Epple, Andreas Koch, Volker Empl, Bernhard Zimmermann, Markus Berg, and Dirk Schneider at HP Consulting & Integration, Matthias Precht, Dirk Benecke, and Peter Schenk at HP OpenView, and the countless others who contributed to this book with commitment and dedication.

We would also like to thank in particular Mr. Robert Riemann at Porsche, whose detailed comments made a major contribution to the content of many chapters, and Ms. Susanne Jansen at SAP, who provided much valuable information. The practical experience of all our advisors greatly added to the value of this book, and their support was a great source of encouragement.

Lastly, our thanks go to Florian Zimniak, Nancy Etscovitz, John Parker and the rest of the staff at Galileo Press and Wellesley Information Services, who made it possible for us to bring this book to life.

11 Local Area Network Solutions

SAP NetWeaver solutions on the internal data highway

Network infrastructures installed on a company's premises are called local-area networks (LAN). Modern LAN technologies are capable of providing extremely high bandwidths on an area of several square miles.

When our first book on SAP infrastructures was published,[1] many company networks were still based on proprietary terminal-based applications with their own cable and plug types, bus systems based on coaxial cables, and architectures that were built from hubs and bridges. With the advent of SAP, cabling and network technology often had to be completely redone.

Since then, a unified infrastructure based on full switched Ethernet fiber optics, or twisted pair cables, RJ45 plugs, and hierarchical architectures from department and backbone switches has become standard, which is in keeping with the requirements of SAP solutions.

This chapter will therefore deal with only those aspects of local networks that pertain to the end-to-end availability of SAP systems. For example, these include technical characteristics that are often ignored such as lightning protection and electric potential equalization. When planning new buildings or approving the facilities in rented buildings, we advise you to read the corresponding chapters in the book mentioned above.

At the same time, mobile network technologies have undergone further rapid development. Therefore, we have devoted a specific section of this chapter to the characteristics, and the advantages and disadvantages of the wireless network.

Bandwidth demand for end users

Today's LANs are integrated data highways that transport a variety of data flows. These data flows can be divided into three categories: business-critical, time-critical, and other data traffic. The average bandwidth requirement of a typical client application is:

- SAP GUI: 1.5kbps (depending on the content)
- VoIP: 12kbps (depending on the compression)

1 Mißbach, Hoffmann: *SAP Hardware Solutions—Servers, Storage, and Networks for mySAP.com*. Prentice Hall 2000.

- Web browser: 30kbps (depending on the content)
- MPEG video: 1.45Mbps (depending on the compression)
- File transfer: many Mbps but not time-critical

If you compare the bandwidth consumption of these typical online activities, you'll notice that an SAP system, when contrasted with all applications, creates the least amount of data load on the network. Since only a few users use SAP and video-on-demand simultaneously, a bandwidth of 10Mbps is more than adequate for the individual end-user connection.

The individual data flows are merged at the work group level. For 12 to 24 users, a bandwidth of 100Mbps will generally suffice. The data flows are merged a second time at the building level where a multiple of 100Mbps or 1,000Mbps is required as a bandwidth. In the computer center, this bandwidth is then distributed with 100Mbps or 1,000Mbps connections to the server systems on which the applications run.

Fast, Gigabit, and more Ethernet

Optical and electrical signals are subject to various negative effects on their path. As these effects increase in proportion to the product of frequency and distance, the achievable bandwidths and range of communication lines are subject to physical limits. Therefore, increasing the signal frequency alone will not help you to attain a higher bandwidth. Special encoding algorithms must ensure that the signal that was originally put into the medium on the sender's side can be correctly reconstructed from the signal that contains added noise on the recipient's side.

Fast Ethernet and Gigabit Ethernet could only be developed so quickly, because existing tried-and-tested coding algorithms could be reverted to. Fast Ethernet is based on technologies that were originally developed for Fiber Distributed Data Interface (FDDI), according to ISO 9314. For Gigabit Ethernet, the encoding algorithms of fiber channel (see Chapter 7) where adapted. Meanwhile, even switches with 10 gigabit uplinks are no longer considered exotic.

11.1 High Availability for Local Networks

As already discussed, the local network architectures generally used contain several points of failure regarding both the active components and the cable connections. The worst-case scenario is represented by a failure of the central backbone. This is similar to a blackout for the entire company. A disruption at the work group level affects the work of an entire department.

Beware of loops!

In order to achieve high availability of local networks, it is apparent that all cable connections and active components should be redundantly designed. Unfortunately, for Ethernet networks, this is not always possible because redundant connections would generate loops in the data path, which must be avoided under all circumstances within a broadcast domain.

To understand this problem, we must look at the functionality of switches. A switch learns the addresses of the hosts which are reachable through its different ports from the transferred packets that arrive at these ports. In an internal address table, the switch saves information regarding whose hosts can be reached via which connection.

When a data packet arrives at a port, a check is performed to verify whether the destination address already exists in the address table. If this is the case, the packet is forwarded through the corresponding connection. Otherwise, the switch replicates the packet through all its connections, with the exception of the connection through which the packet was received. This means that target hosts, which had been unknown until this point, can be reached, and due to their response, the address table can be updated. This is exactly where the problem of redundant configuration lies, as shown in Figure 11.1.

In this example, there are two potential paths from computer a to computer b, which, in each case, lead through a connection to switches 1 and 2. What happens if computer *a* sends a packet to computer *b*, but computer b is not yet listed in the address tables of either switch?

1. Computer *a* in network Segment A transfers a packet. Both switches learn that computer *a* can be reached through connection 1.1. or 2.1, and broadcast the packet through all the other connections as they do not yet know the address of computer *b*.
2. Thus, there are two identical packets in Segment B. Because both switches are connected with each other via Segment B, these packets also reach both switches *cross-over*, then through connection 1.2 or 2.2 respectively. Since the packet still contains computer *a* as the sender, both switches learn that computer *a* suddenly seems to be located in Segment B.
3. As computer *b* is still not recognized, the packet is once again replicated through connections 1.1 or 2.1 into Segment A, from which it originates and is then duplicated. As the switches do not know each other and each switch continuously broadcasts the packet into the other segment, an endless loop is generated.

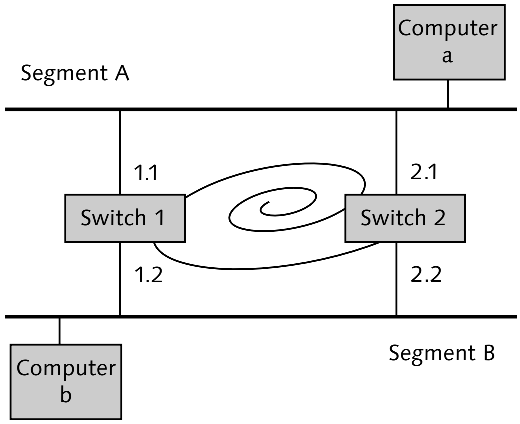

Figure 11.1 Network Loops caused by Redundant Paths

Broadcast storms and network meltdown

Due to the endless replication of the broadcast packets, the loop creates an avalanche effect. In a short space of time, one of the feared broadcast storms floods the network. In such a situation, no more communication is possible as the broadcasts use up the entire available bandwidth. This causes a network meltdown. The PCs connected are so strongly overloaded with interrupts that the systems freeze, which disturbs the entire data processing. For this reason, you should definitely avoid starting loops within a broadcast domain.

Spanning tree

To solve this problem, Radia Perlman developed the spanning tree algorithm (IEEE 802.1D). Switches exchange information using the spanning tree protocol, in order to recognize parallel paths. These paths are then shut down in sequence until only one is left. The remaining loop-free paths result in a tree structure that spans from the data center to the end devices, which is why it is called a spanning tree.

The disadvantage of the spanning tree process is that in redundant processes only one link can be used for data transport, while all other links are switched to standby mode. Investments in these cables and connections are therefore not exploited until there is a breakdown. In the event of a breakdown, the necessary recalculation of the spanning tree is also a relatively time-consuming process. During this time, the connections don't forward any more packets.

Modern switches limit broadcast storms

Therefore, the spanning tree concept doesn't play an important role anymore. Modern switches have mechanisms that limit broadcast storms. These mechanisms are based on the assumption that typically a certain ratio between user data and broadcasts is not exceeded. In the case of a broadcast storm, those broadcasts that go over the limit are distorted. You should, however, implement such mechanisms with caution. On the one hand, there is the danger that good broadcast packets could also be eliminated; on the other hand, the danger exists that a real broadcast storm could be disguised. In both cases, problems that are difficult to identify can emerge.

11.1.1 Link Aggregation

Different manufacturers provide different technologies that can bundle several 100Mbps or 1 Gbit/s connections (link aggregation). For the operating system, the bundled connections represent a single logical interface with a single MAC and IP address. Due to this aggregation, the load is distributed to the parallel connections. This provides higher performance and redundant paths. If a connection breaks down, the data traffic is automatically transferred to the remaining connections within the bundle.

However, link aggregation supports only parallel point-to-point connections between two devices. This means the switches that are linked through link aggregation still represent single points of failure. In addition, the different cables of a bundle are generally placed in the same position so that they're exposed to the same potential risks and can be simultaneously destroyed.

11.1.2 Highly Available Network Clusters for Business-critical Applications

Installing redundant cables and switches for each important work center in a company would lead to exorbitantly high costs and immensely complex configurations. An alternative approach, (which was developed by one of the authors) is based on the view of the business functions in a enterprise to design highly available networks for business-critical applications.

Human redundancy

From this hands-on approach, you can assume that a high availability is essential at the department level, but for individual work centers, a certain downtime can be tolerated. This is because at the department level there is always a functional redundancy since each user is assigned a substitute for leave, sickness, and so on. This substitute is often the colleague at the next desk. Even if the substitutes themselves are not there, the employee whose connection to the SAP system is disrupted can use the colleague's PC (or the colleague's connection wall socket). This fact can be used to grant also that work will be done even in case a network device is going down.

Rules for network clusters

In order to avoid SPoF on the business functions level, network clusters must be configured according to the following simple rules:

- An employee's PC that fulfills an important business function should never be connected to the same network switch as that of his or her substitute.
- For this reason, every network cabinet must contain at least two switches with separate connections (uplinks) to the data center.
- If possible, the connections should be executed on different paths.
- There must be at least two backbone switches installed in the data center (in different fire protection zones).
- Each clustered SAP server must be connected through separate network cards to these backbone switches.

Network clusters can most easily be implemented by intelligent patching, which consists of linking the end device connections on the patch panels of the network cabinet with the connections of the switches. You can generally assume that adjacent PCs are provided for employees with the same business functions that can mutually cover for each other. In order to meet the requirements of a network cluster, you only need to attach the end device connections with straight numbers on the patch panel to one switch and those with odd numbers to another switch.

When this concept is implemented methodically, single points of failure are avoided at the business function level as well as network loops and inactive standby connections. Network clusters can be implemented with plug & play components of any manufacturer.

Two simple examples illustrate the network cluster concept. In Figure 11.2, the sales departments PCs are located on the left-hand side and the logistics department PCs are located on the right-hand side.

- **Scenario 1:** One of the switches (or its uplink) in the sales department breaks down. Every other work center is dead, but the rest remain operational. In a typical network environment in which all PCs of the same department are connected to the same switch, the entire department cannot process any more orders.
- **Scenario 2:** One of the backbone switches in the data center breaks down. In half of the hosts the connection breaks down, but the rest, how-

ever, remain operational. If it takes too long to replace the backbone, switch cross connections between the switches can be used as a bypass.

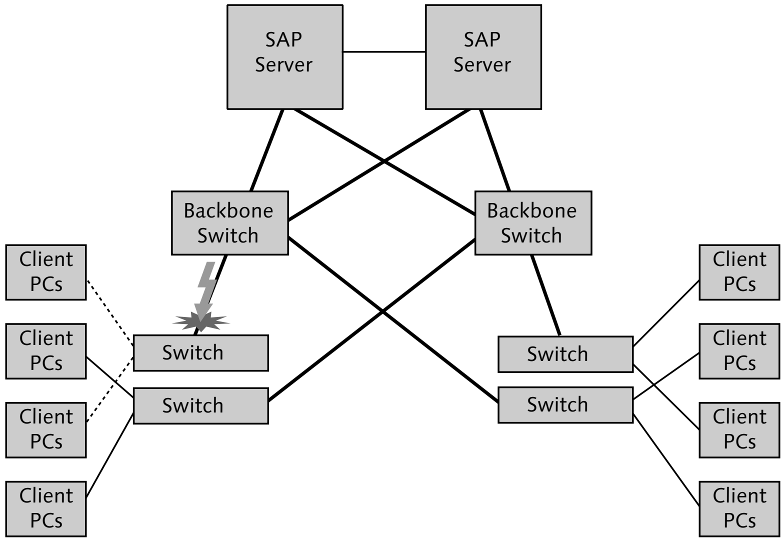

Figure 11.2 Highly Available Loop-Free Network Cluster

The network cluster concept ensures that at least one in every two PCs of a department has a connection to the SAP system at any time from a network perspective, and the business tasks of a department can be executed in any situation. The investments are the same as for a redundantly designed network based on the spanning tree concept. Alternatively, there is a network cluster but no convergence time, and the available bandwidth is substantially higher due to the utilization of all available links and connections.

11.1.3 Error-tolerant Meshed Networks

For a network cluster, even when a backbone switch or link breaks down (see Scenario 2 in the previous example), the operation of the enterprise can be maintained. However, in this case, up to 50 % of the work centers can lose their network connection. Due to switch meshing, the network cluster concept can be extended so that even the breakdown of a backbone component can be absorbed. This means that the local network is extensively error tolerant.

Switch meshing

Switch meshing is a technology originally developed by Hewlett Packard, which enables the creation of a completely meshed local network infrastructure without generating the risk of loops and subsequent network meltdown due to broadcast storms. Currently, other manufacturers implement this technology as well. All links and connections are always active. On the basis of load statistics, algorithms distribute the data traffic equally to all links and prevent broadcast landslides.

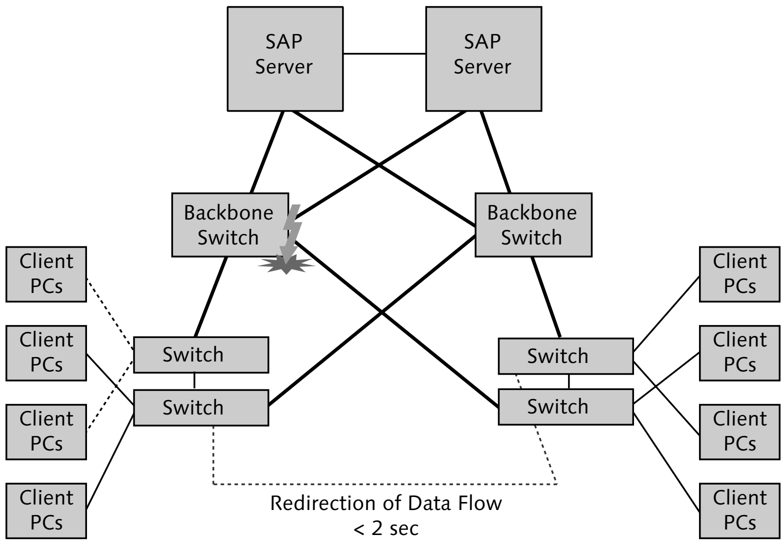

Figure 11.3 Error-Tolerant, Completely Meshed Network

If the "intelligent patching" described for network clusters is implemented with completely meshed switches, even in the case of a failing backbone component, the full operation of the network can be ensured. The switchover time in an SAP cluster test environment, after turning off a backbone switch, was under two seconds in a running operation. The switch took place transparently for the application and user without losing any transactions.

UPS Units Are Also for Network Cabinets

To ensure the availability of the SAP infrastructure, in areas where there are more frequent voltage fluctuations, Uninterrupted Power Supply (UPS) units should be used. The ability of a UPS unit to filter and stabilize the power supply voltage is more important than bridging long power outages. A switching operation in the high-voltage power grid of the electricity supplier only causes the office lighting to flicker, but the switches and routers may restart and cause network downtime. However, USP units are also active components that must be monitored and maintained.

11.2 Wires and Fibers

Today, basically two cable types are used for local networks: lines with twisted pairs of copper wires and fiber optics cables. Both have types have advantages and disadvantages, because of their physical characteristics.

11.2.1 Copper Cables

The area of end device connections is generally based on twisted pair copper wires. This type of cable has existed since the first telephone signals were transferred. Throughout the years, on the one hand, the transfer frequencies have become much higher; on the other hand, there are essentially more sources of disturbance. Fortunately, twisted pair cables were also developed to the same extent. Therefore, we can say with assurance that twisted copper cables actually do meet the requirements of high-speed data transfers. Compared with optical fibers, copper cables are much easier to install and, consequently, are more cost-effective.

Copper cables—"lets twist again"

Twisted pair cable consists of two copper wires. Each wire is encased in its own color-coded insulation, twisted around one another. Multiple pairs are packaged in an outer sheath, or jacket, to form a twisted-pair cable. The twist of the cable is essential for electrical noise immunity and must go as near as possible to the connectors of the wall receptacles and patch-panels. By varying the length of the twists in nearby pairs, the crosstalk between pairs in the same cable sheath can be minimized. The typical nominal impedance is 100 ohms.

For data networks in companies, structured cabling in accordance with EN 50173-1[2], ISO/IEC, or EIA/TIA-568 category 5 onwards has become standard.

The decisive quality attribute for top quality data cables is the symmetry of the cable. Different twist lengths of pairs that are placed next to each other avoid crosstalk. In this context, it is important that the cables are not only symmetrically stranded but also precisely finished.

2 General requirements for application-neutral communication systems in offices (2003).

High Bandwidth is not Equal to High Frequency

There only appears to be a connection between the transfer frequency and the achievable bandwidth. By using highly developed signal encoding processes, all high-speed technologies such as Fast Ethernet and Giga Ethernet, as well as ATM, don't exeed 310 MHz as the transfer frequency. Technologies with higher bandwidth are based on fibre optic cables.

Shielded or unshielded?

Contemporary data networks operating with frequencies that are located in the middle area of the VHF radio band (Very High Frequency). The metallic conductors act like antennas for these frequencies, for receiving as well as for transmitting. STP cables contain a metal shield to reduce the potential for electromagnetic interference (EMI). EMI is caused by alternating electromagnetic fields from other sources such as electric motors, power lines, high power radio and radar signals but also by flickering fluorescent tubes in the vicinity that may cause disruptions or interference, called noise.

There are, however, a variety of shielding solutions for data cables. From the most simple aluminum polyester compound film, through combinations of tin-plated twisted meshwork and compound film, to expensive metal shields, you can find all possible constructions.

The names of the various shielded (Shielded Twisted Pairs, STP) and unshielded (Unshielded Twisted Pair, UTP) cable types are quite confusing. STP also encompasses Screened Shielded Twisted Pair (ScTP) and Foil Twisted Pair (FTP) cables. Within UTP, there are paradoxically also Shielded Unshielded cables (S-UTP) with a complete external shielding, but without individual shielding of the pairs.

STP cables

At first glance, STP cables appear to be immune to any interferences, because of their shielding. But, unfortunately, this is not the case. As the grounded shield also acts as an antenna and transforms the incoming interferences into a current, which induces a current in the signal wires in the opposite direction. As long as both currents are symmetrical, they eliminate each other. Any discontinuity in shielding or asymmetry in the currents between the shield and signal wires acts as a source for electronic noise. Therefore, STP cables are effective only if the entire link from one end to the other is continuously shielded and properly grounded. However, this can in turn cause severe problems by amperage flow over the shield in cases, where the electrical supply grid has no separated grounding.

For UTP cables, the physical shield is replaced by improved variations of the twisting as well as sophisticated filtering techniques in the network devices. Disturbances are equally induced in both conductors and therefore eliminate each other. Throughout the years, the UTP cables have constantly been improved so that they now fully meet the requirements of category 5. **UTP cable**

Despite a heated debate over the years about the advantages and disadvantages of shielded versus unshielded twisted pair cables, a final conclusion has still not been reached. In Europe, STP is the main preference, not only to protect data signals against outside emissions, but because corresponding regulations require the protection of the environment against the emissions of data signals. UTP cables are used generally in the rest of the world. In any case, reliability is always determined by the quality of the cable manufacturing and proper installation.

Electromagnetic Compatibility (EMC)

Another factor to consider when choosing a cabling system relates to electromagnetic compatibility (EMC). In the U.S. and Germany, EMC regulations have existed for years. However, the implementation of the European EMC Directive 89/336/EEC in 1989 has refocused attention on EMC. With the increased amount of electronic equipment in the average workspace, EMC becomes increasingly more important. Excess radiation from one piece of equipment can adversely affect performance of another piece of equipment. EMC refers to the ability of an electronic system to function properly in an environment where several pieces of equipment radiate electromagnetic emissions. This means that every electronic system, which includes all copper based cabling systems, must meet this directive.

11.2.2 Fiber Optics Cables

A fiber optics cable consists of a bundle of optical threads (fibers), in which messages modulated onto light waves propagate along the direction of the fiber because of internal reflection. The reflection occurs due to the different refraction indices between the core and the coating. Fiber-optic network cabling is made up of at least two strands of optical fiber running parallel to each other in a plastic "zip-cord" jacket, or multiple fibers in a single jacket.

Multi-Mode Fibers

There are two different types of optical fibers: Multi-Mode Fibers (MMF) and Single-Mode Fibers (SMF). In this context, "mode" refers to the effect of spreading a light signal in an optical fiber, resulting in light-rays following different paths (or modes) down the fiber (modal dispersion). Multimode fiber (MMF), with a core diameter of 62.5μm allows a light signal to take various zigzag paths. This modal dispersion causes some light rays to arrive later at the end of the fiber. Due to the undesired runtime differences caused by these different modes, the range of MMF cables is limited to 1.5 miles (2 km).

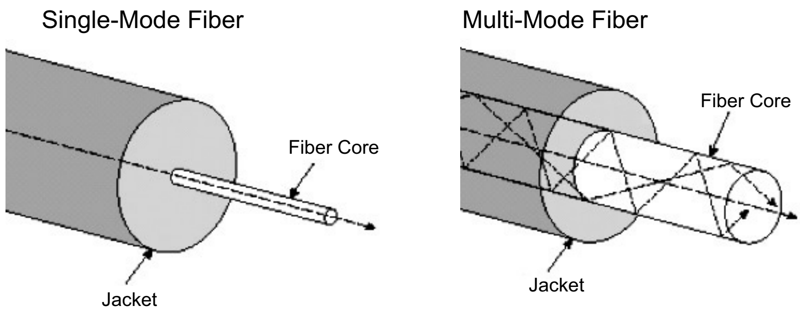

Figure 11.4 The Route of Light Rays in Single-Mode and Multi-Mode Fiber Optics Cables

Single-mode fibers

Single-mode fiber (SMF) with a core diameter of only 9 microns allows only one path for the light to take due to the fiber's very small diameter. Single-mode fibers and components are more expensive than Multi-mode fiber, but allow connectivity up to 12 miles (20 km). Cheap Plastic SMF fibers can be used only for very short connections, they absorb the light rays earlier because plastic is not as clear as glass.

Multi-mode fiber is designed for coupling light from low cost LED[3]-based transmitters. Single-mode fiber is only suitable for laser-based transmission.

There are varieties of connectors (FDDI-MIC, ST, SC-Duplex etc.) for MMF as well as for SMF. Be sure to use the same diameter and connectors throughout your infrastructure. Project deadlines are easily missed when plugs at fibers do not fit the active components, and adapter cables are not on hand.

3 Light Emitting Diodes.

Fiber optics cables can be used for all current network technologies. In comparison to network connections made of metal, fiber optic cables have numerous advantages:

Advantages of optical fibers

- Optical fibers enable a larger bandwidth than copper cables.
- Optical fibers are not sensitive to electromagnetic radiation and don't emit any by themselves. This means that all regulations are meet by default.
- Optical fibers are immune to lightning strikes and power line transients.
- Metal free Optical fibers cannot generate any ground loops.
- Optical fibers are much thinner and lighter in weight than metal wires.
- It is very difficult to tap eavesdrop on optical fibers without being noticed, making this very secure from electronic eavesdropping.

On the other hand, fiber optics cables also have some disadvantages:

Drawbacks

- Fiber optics cable connectors are high precision parts. An exact alignment of the fiber inside the connector housing and proper polish of the fiber end is crucial for connectivity quality.
- Installing connectors on site is time consuming and requires high precision work. The alternative of fusion welding strands with prefabricated connectors (called pigtails) on site needs expensive equipment.
- Together with the higher costs for the cable itself, the deployment of fiber-optic cable costs more than twice that of a category 5 copper connection. The "per port" price of active fiber optic components (hubs, switches, router modules) is typically twice the price of their copper counterparts. The prices for long range single-mode cables and components are even higher than for multi-mode
- Glass fibers are more fragile than wire and sensitive and age under the impact of hydrogen ions. They must therefore be protected against moisture through special coatings. However, this protective layer is also subject to aging.

Case Example: Mice in Cable Conduit

In a company, a complete administration building was suddenly without network based IT services. The reason behind this was a fiber optics cable that had been gnawed through in a cable conduit. Rodents like to build their houses in cable conduits, and their offspring like to test their teeth on the cables. Since fiber optics cables can only be spliced by using special tools, this led to a downtime lasting several days. Even rodent-safe cables and mousetraps are therefore investments that increase availability of enterprise service architectures and IT services in general.

11.2.3 Installation Guidelines for Cable Networks

Good craftsmanship, together with using high quality components, has a direct relationship to how long your cabling infrastructure will last. As mentioned before, high-speed data links have higher demands than plain-old telephone lines. To make matters worse, the effects of poor installation work may not be immediately evident!

Deformation or mechanical stress during installation causes most cable failures. Deformation changes the physical properties responsible for high frequency transmission. Even when the cable looks flawless from the outside, irreversible degradation of transmission properties is suspected when too much stress is applied to the cable. Mechanical stresses as well as temperature levels are part of the ISO/IEC 11801 standard "Generic Cabling for Customer Premises." The installation of network infrastructures should be dedicated to certified contractors, familiar to the special demands of data cabling.

However, even with equipment, which has been checked and conforms to the standards, problems still exist, as with increased demands on the networks, the reserves in the transfer parameters decrease.

Visual inspections

An example is the standard of fitting the data cables in the connection components. The standard for this does provide for a visual inspection, but this is rarely carried out. Generally, people content themselves with the test logs created by cable scanners. For instance, in distributor panels where the cable sections that have been stripped of the isolation are narrowly guided along the blank wire ends of the neighboring cable, this can lead to a short circuit between the wire and the foil shield if the latter changes, for example, because the temperature of the cabinet interior increases. During an acceptance test on cabling, it is therefore critical that you perform visual inspections.

With fiber-optic cables, you can run into problems later on as well, even if they have been installed according to the standards. Here the problems can mainly be found in the preconfigured plugs. These plugs are very sensitive to scratches, dust, and inept treatment when being mounted. We also recommend that you use a microscope when conducting checks. You can find additional information and further considerations when executing acceptance tests in Mißbach/Hoffmann,[4] which is also worthwhile reading for small changes or enhancements carried out by your internal electrician.

11.3 Potential Equalization, Grounding, and Lightning Protection

Frequently, sporadic disturbances and breakdowns occur in data networks without a clear reason. Connections become extremely slow, monitor screens flicker, assemblies burn through or after a thunderstorm, and entire facilities break down at once. Furthermore, individual employees may be marginalized, because strangely it is only always their PCs that go "mad."

In many cases, however, the real reason for these breakdowns can be found in the potential equalization and grounding. Apart from the data network cabling in every building, there is also cabling for the power supply. People often overlook the fact that these two cable networks are linked via the grounding, and massive disturbances of the data networks can occur if the power supply network is not designed as IT-compatible.

Potential displacements

In order to ensure an electricity flow, a wire (L) is required from the electricity source to the consumer, as well as a retracting wire (N). A third wire is stipulated as a ground wire or protective earth conductor (PE). As the 230V alternating current is tapped from a 380V three-phase network, this results in a 5-wire network or TN-S system with three conducting phases and a common neutral and ground wire for each phase.

The security function of the protective earth conductor is ensured, even if the protective earth conductor is connected to the neutral cable that is also grounded, and thereby forms a "combined" PEN conductor (that is, a retracting circuit (N) plus a protective earth conductor (PE)). Such cou-

4 Mißbach, Hoffmann: *SAP Hardware Solutions—Servers, Storage, and Networks for mySAP.com*. Prentice Hall 2000.

plings are permitted and are commonly used in building installations, because this means that a cable can be saved. This form of network is called a 4-cable network or TN-C system. This variant has no negative effects on the lamps connected, whereas, when connecting other electronic appliances, this can lead to considerable problems.

Interfering transmitters in the system

Since an increasing number of electronic pre-connection units are being used for fluorescent tubes and switching power supplies for computers, the current flow is not sinusoidal. Instead, it contains considerable high-frequency components. These can cause parasitic currents of several amperes, which results in a magnetic field that acts as if an irregularly functioning high-frequency transmitter is an integral part of the computer system.

To ensure a stable operation of the data network, an integrated 5-wire network with a clear grounding concept should be guaranteed. If necessary, a separate 230V power supply grid must be installed from the transformer of the sub-distributor. In this separate supply circuit, the neutral cable must not come into contact with the ground wire at any point (separate conduits for PE and N, so that no parasitic currents are possible through the data cables). The sockets for this network should be marked "only for IT devices." There should also be corresponding signs in the rooms of its installation to ensure that an inexperienced electrician does not create a connection between N and PE again.

Danger—high voltages due to lightning

In conjunction with a nearby lightning stroke, cable screens grounded on only one side, act as antennae into which high voltages are induced. The voltage will be discharged at the network board or the network socket, depending on where the disruption to the shield occurs. Even with a double-sided grounding due to the induction caused by a lightning stroke, a compensatory current flows in the data shield. This can lead to extensive destructions of the connected interfaces as a result of strong electric currents. Therefore, power sockets with integrated over voltage protectors are definitely advisable. The same is true for your house, where lightning bolts can induce voltage surges in the main telephone line, which destroy telephone systems and DSL routers, as one of the authors can attest from his own first-hand experience.

Parasitic currents through ground loops

Ground loops represent another cause of parasitic currents, and one often overlooked. In the network cabling system, problems with ground loops occur mainly when shielded twisted-pair cables (STP) are used between parts of a building complex. If these parts have different earth potential, parasitic currents of several amperes can flow over the shield if grounded on both ends. These ground loops can lead to a degradation of network performance, and even to a damage of network components.

When installing STP cables, you must ensure that grounding occurs only at one end of the grounded link. For UTP cables, ground loops are not an issue, because of the design of these cables. Fiber-optic cables are the most secure way of eliminating damage caused by parasitic currents of any type.

Fire Protection

The insulation of data cables generally consists of flammable synthetic materials, mainly polyethylene and PVC. To reduce the threat to people and material, Building Codes generally require cables that does not generate toxic fumes when burning, such as FEP (fluoro-ethylene polymer) in air ducts, air plenums, and other environmental air spaces.

In the event of a fire, these cables also don't generate corrosive gasesand the smoke gas density is considerably lower. Therefore, PVC cables are banned in building installations in an increasing number of countries.

It is for good reason that building insurers in recent years placed major emphasis on the issue of fire protection. According to the insurers, all cable conduits in ceilings that run at a right angle to an emergency route must have a full fire protection. This also applies to all later changes. You should therefore always ensure that the cable installations are in accordance with the regulations of the property insurers.

11.4 Wireless Networks

Wireless technologies in recent years have undergone rapid development in terms of both their bandwidths and their increased use. In 2004, 42 % of all notebooks were already equipped with Wireless LAN (WLAN) functionality, and, according to estimates from the IDC market research institute, this percentage will rise to 98 % by 2007.

11.4.1 WLAN Standards

Wireless local networks (Wireless LAN, WLAN) are defined in the norms IEEE[5] 802.11 and ISO CD8802-11. After an initial 2Mbps in the first

5 Institute of Electrical and Electronics Engineers.

802.11 standard without additional letters at the time, 802.11b followed with 11Mbps and 802.11g followed with up to 54Mbps (~30Mbps net) with three channels on the 2.4 GHz frequency. Some manufacturers already provide systems with transfer rates of 100 Mbit/s, with the assurance that modifications resulting from a standard that will subsequently be introduced can be imported in the form of a software update. Contrary to this, 802.11a and 802.11h each have eight channels in the 5 GHz band at their disposal. However, this band is also used by radar systems and earth observation satellites. A certain detection threshold in the interaction with Dynamic Frequency Selection (DFS) is supposed to ensure that these radar systems and satellites are not disturbed by WLAN.

The 802.11 alphabet

The 802.11 alphabet also knows a range of other letters, which reflect modifications made according to country-specific regulations or modifications made for functional enhancements.

For example, variant 802.11h was introduced to comply with the requirements of some European countries for a automatic adaptation of the transmission power (Transmission Power Control, TPC), which is supposed to further reduce the probability of interferences. Without TPC and DFS, 5 GHz radio networks in Europe can be subject to very rigid obligations that limit the operation to ranges of less than 20 meters in buildings.

The letters e, l, and f, are used for function enhancements. IEEE 802.11f defines the Inter Access Point Protocol (IAPP) for roaming between Access Points of different manufacturers. The 802.11i standard is supposed to protect wireless networks against unauthorized access by implementing encryption processes and user authorization. The 802.11e attempts to enable a prioritizing of certain applications such as voice-over IP (VoIP) in the WLAN. However, no bandwidth can be guaranteed here; rather, the access point tries to implement the different priority levels in as far as possible.

Standard 802.11a uses Orthogonal Frequency Division Multiplexing (OFDM) for transmission with up to 12 channels depending on the country, while standard 802.11b is based on Direct Sequence Spread Spectrum (DSSS). The development of 802.11g was made possible after the Federal Communications Commission (FCC) had released the OFDM technology in the 2.4 GHz band.

802.16a

In addition to the 802.11 family, the IEEE has ratified norms for even higher speeds. 802.16a or Wimax[6] uses the frequency range of 2 to 11

GHz for transfer speeds of 70Mbps with coverage of almost 30 miles, while 802.16b is supposed to enable even 134.4Mbps in the frequency range of 10 to 66 GHz. The Wimax Forum[7] aims to have an unlicensed band for 802.16a with 5.8 GHz and two licensed bands with 2.5 and 3.5 GHz. Other bands are to follow later. The radius of a cell should realistically be set to 20 miles, regardless of whether there are obstacles between the sender and the recipient. Market researchers expect a wide availability of these technologies by 2008.

A Practical Example of WLAN

In the feeder railway of a lignite mine, information about shunting tasks, the sequence of wagons, wagon data, and the currently covered length of shunting tracks is transferred to the handheld PCs of shunters and locomotive drivers via 802.11b WLAN.

To minimize the number of access points, the antennae were aligned along geographically defined lines of sight, and the locomotives were equipped with an access point repeater that ensures radio coverage for the shunter's handheld in the area of his shunting unit. In addition, the locomotive drivers enter operating data such as engine and compressor runtimes, or the fuel consumption directly through pocket PC and WLAN on the driver's console.

While using WLAN for mobile applications is indispensable, the rapid development of radio technologies to ever-increasing bandwidths and the growing use of laptops with integrated WLAN interfaces beg the question of whether avoiding cabling altogether, for stationary work centers as well, is a real possibility. There are, however, some constraints to consider that you won't typically find in the manufacturers' brochures.

Bandwidth is dependent on the distance

Due to physical laws, at a given transmitting power, the bandwidth is dependent on the distance. As you can see in Table 11.1, the rules are as follows: The larger the distance, the lower the actual bandwidth. For increasing frequencies, obstacles cause mounting problems to electromagnetic waves.

---

6 Worldwide Interoperability for Microwave Access, the corresponding name of the European ETSI is HiperMAN as this technology is generally regarded as suited for Metropolitan Area Networks (MAN).

7 Go to *www.wimaxforum.org.*

	11Mbps	5.5Mbps	2Mbps	1Mbps
Open country	70 yds. / 66 m	100 yds. / 91 m	135 yds. / 125 m	187 yds. / 171 m
In buildings	30 yds. / 28 m	38 yds. / 35 m	47 yds. / 43 m	58 yds. / 53 m

Table 11.1 Correlation between WLAN Coverage and Bandwidth

In addition, all WLAN protocols cause considerable overheads, so out of a bandwidth of up to 11Mbps, approximately only 7Mbps can be used for user data and this can only be reached when you are close to the transmitter (i.e., access point).

ISM band

Most WLAN products use the ISM (industry, science, and medicine) band in a range of 2.4 GHz. However, the transmitting power in the ISM band is limited to 500 mW so the highly sensitive medicinal diagnosis systems don't get disturbed. For this reason, mobile telephones that have an essentially higher output power must not be used in hospitals, and you should resist the urge to download your latest email to the computer via your cell phone while sitting in the waiting room of the intensive care unit. The low transmitting power reduces the range correspondingly. Similar restrictions also apply to the 5 GHz band.

In addition, a basic disadvantage of radio networks is that all participants must share the available bandwidth. It is precisely because the ISM band is license-free that it is also used by many other systems. Examples of other systems include the wireless control of erection cranes and the hobby area. RC cars and planes are therefore potential sources of inferences for WLANs. The Bluetooth short distance radio technology, which allows mobiles, headsets, handhelds, and printers to communicate with each other, uses also the 2.4 GHz ISM band. Even a microwave can pose a possible source of disturbance since it also often works in the 2.4 GHz range.

5 Gigahertz band

Like the 2.4 GHz range, the 5 GHz band is similarly utilized by radio applications. However, because of the higher number of channels, more users can share a radio cell.

All end devices share the bandwidth

All end devices set to a specific channel necessarily share the bandwidth. Therefore, the 11Mbps that exist on paper can easily become only 600 kbit/s in actuality, and even this cannot be guaranteed. In certain circumstances, this leads to drastically increased response times for an SAP user who is connected to an SAP system through WLAN if a coworker is currently loading a large email attachment on the same channel. If access points are used as radio bridges, the range increases but not the bandwidth as the traffic from the neighboring cell also has to be transferred.

Power users in fixed cabled network connections

In order to cover larger areas and user numbers, more cells must be installed. The access points required for this, in turn, need a conventional cabling, which means that you cannot avoid providing a fixed cabled network. Practice has shown the benefit of equipping power users with both stationary work centers and a fixed cable network connection, and a hotspot for the shared desk area in the office.

11.4.2 Installation Guidelines for Wireless Networks

In order to determine the locations for WLAN access nodes (so called WLAN basis station), the construction drawings for the building should be inspected for hidden metal constructions such as steal reinforcement and water pipes, which shield the radio waves like a Faraday cage and therefore disrupt the WLAN connection. But, even a larger number of people have a negative influence on the performance of a radio network (the high water content of human beings damp the radio waves). Due to starkly reduced prices for access points, you can simply install some of them on a trial-and-error basis.

Directional Antennae for Improved Radio Coverage in Warehouses

In warehouses and manufacturing, wireless mobile terminals with bar code readers are frequently used to compile data along with the mobile data entry interface of SAP Materials Management (MM-MOB).

The steel racks and the steel reinforced concrete walls in high-bay racking, however, absorb the transmitting energy of omnidirectional antennae. Directional antennae that radiate into the warehouse alleys can ensure a stable connection.

WLAN channel layout

Using the Direct Sequence Spread Spectrum (DSSS) technology, the IEEE 802.11b standard provides 13 channels for transmission; however, because these channels overlap each other, they can't be used in direct proximity (side by side). At the end, there are three triples (channels 1, 6, and 11; channels 2, 7, and 12; and channels 3, 8, and 13), which don't overlap each other.

This means that in an ideal scenario a maximum of three access points with a total bandwidth of 33Mbps can cover a room without any disturbance. Anyone can send and receive on a different frequency without any interference provided there's a sufficient distance between the sender and the receiver. To ensure complete redundancy, the radio field of an access point must also be covered by the radio field of a second access point.

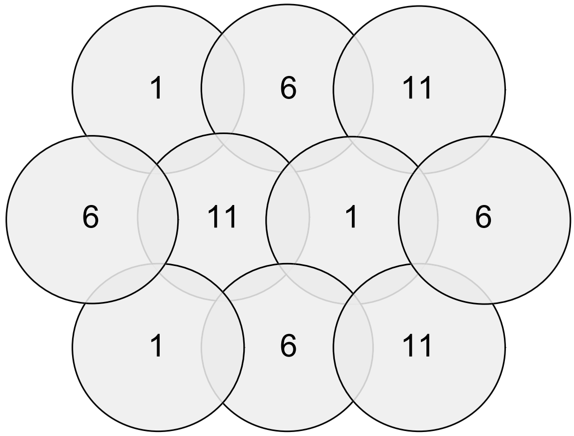

Figure 11.5 Channel Layout for WLAN

For larger WLAN installations, you must ensure that access points, which are situated next to each other, are configured with different channel numbers. Otherwise, they would mutually disrupt each other. Here you must consider that access points necessarily radiate through several floors. As the individual radio channels also partly radiate into neighboring frequencies, for instance, when node A transmits on channel 1, the directly adjacent access node B should be set to channel 6, and node C should be set to channel 11. Therefore a carefully designed channel layout is necessary for larger WLAN installations.

In larger office buildings that are used by several companies, there are also problems if the WLANS of individual companies overlap and thereby cause interferences. If services in the form of hot spots are provided for third parties, in some countries a concession is required, which is currently still free of charge.

WLAN—a paradise for hackers?

Experience has shown that many WLANs are insufficiently secured. According to a study by Ernst & Young from 2003, over 50 % of users do not change their default passwords to access points, 25 % configure the Service Set Identifier (SSID) in such a way that it reveals the network name, and in many cases, even the company name or the IP address. But, at least 48 % of WLAN users use a Virtual Private Networking (VPN) to protect their data. Only a third of users implement a firewall between wired LAN and WLAN. In total, WLANs are included in only 33 % of the companies in the technical and regulatory regulations for the security concept.

> **WLANs also Threaten Wired Networks**
>
> For companies, the growing use of notebooks with built-in WLAN connectivity increased the likelihood that the security of their networks was being compromised. The reason behind this was and is that badly configured notebooks function as WLAN access points when they are connected to a company LAN by cable without disabling WLAN functionality. In most cases users are not aware of this security hole and breach security unintentionally.

Hiding the SSID

In addition to encryption, the access procedure can also enhance security. Each WLAN has Service Set Identifier (SSID) as a name. So clients can communicate with the radio network, they must know this SSID and enter it when logging onto the radio network. In hot spots, the SSID is often sent out as a broadcast. If this is prevented, the clients must already know the SSID in order to be able to create any connection. All other participants are excluded from communication with this WLAN.

However, during the authentication process, each client sends the SSID in plain text to the access point, which can easily and most assuredly be eavesdropped on by an attacker. Unfortunately, this is unavoidable, because several different radio networks can exist within one footprint.

MAC-address access control lists

Some manufacturers have integrated access control lists (ACL) in their access points so they can only permit those clients with known MAC addresses to communicate in the WLAN. Although this excludes participants with unknown MAC addresses from using the network, this mechanism can also easily be overcome by attackers with simple methods. During communication in a radio network, the MAC addresses must be transferred unencrypted. This enables the attacker to tap valid MAC addresses, which they can then configure in their own WLAN cards by using the corresponding software.

Fixed keys do not provide sufficient security

Technologies like Wired Equivalent Privacy (WEP) which the key is stored in the access point and the notebook, generally do not provide sufficient security, because they can be relatively easily cracked by scanning the data traffic.

Therefore, we advise you not to implement any WLAN-based encryption; instead, you should establish a secured connection between the client and the firewall with a powerful IPSec encryption in a VPN. In addition, an overall concept from authentication, authorization, accounting, and encryption is necessary.

Security from end to end

In "typical" access point concepts, only the WLAN-side "air interface" of the access point is encrypted while the data in the cabled part is transferred unencrypted. So called WLAN switches can be positioned in such a way that their network port is logically immediately connected to the firewall or the VPN server. As data traffic on the cable route between the WLAN switch and the antenna systems is encrypted in the same way as in the air interface, security is guaranteed from one end to the other without the end user having to install a VPN client.

For big installations, a large number of access points means that configuration and administration becomes time-consuming and costly. These difficulties were overcome on classic, cabled networks by automatic, rule-based switching on network levels 2 and 3. For WLANs, there are corresponding concepts of Wireless LAN switching.

Wireless LAN-switching

To do this, a WLAN switch is installed (for example, from HP, Nortel, Extreme Networks, or Proxim), from which access points and access to the network can be administered centrally. Thus, the decentrally installed access points become pure antenna systems that convert only radio signals to Ethernet packets. The "intelligence" of the WLAN is concentrated in the wireless switch. In general, access points don't even need an IP address. Their power supply can be ensured through "power over Ethernet" according to the 802.3af standard, so that, apart from the Ethernet cable, no further installation is necessary.

For real mobile users, which roam within a WLAN network between the footprints of different access points, a wireless switch provides a single sign on (SSL) and roaming times that are typically under 30 seconds. However, strictly speaking, this is not the kind of roaming we know from mobile phones that roam between the networks of different providers. Instead, it is an interruption-free handover from one radio cell to another.

Furthermore, many WLAN switches offer functions such as automatic channel selection—where the layout of the radio cells is automatically optimized—and preemptive roaming (wireless load balancing).

11.4.3 Ad-Hoc Networks

Bluetooth

Ad-hoc network technologies such as Bluetooth[8] were developed to enable a dynamic connection establishment between mobile devices such as wireless DECT phones, laptops and PDAs. Recently, this list has been complemented by hands free speaking systems and headsets. Originally, the Bluetooth concept was only intended to replace the cables between the phone handsets and their peripheral devices with a radio connection. However, the user spectrum was very quickly expanded to include the world of the personal computer. The developed is controlled by the Bluetooth Special Interest Group.[9]

Network configuration on the fly

While WLANS require a fixed configuration, ad-hoc networks are based on a master-slave system where a master device controls the changing connections in a Piconet cell. As the type and number of devices in the cell can change unexpectedly, the routing protocol used by Bluetooth must be capable of dynamically reconfiguring the network "on the fly."

The designers of Bluetooth, too, have decided to use the license-free 2.4000 GHz–2.4835 GHz ISM frequency band. Since this band is already used by so many other wireless services, Bluetooth uses an Advanced Frequency Hopping Technology (AFH) to avoid interference problems, which have made life difficult for other ISM band users. The AFH concept uses 79 different radio channels from among which it switches 1.6 times per second. Thus, a channel is used only for 625 microseconds before the switch is made to the next randomly selected channel.

Bluetooth currently allows a transfer rate of up to 1Mbps which corresponds to a real throughput rate of approximately 720 kbit/s. Power management in Bluetooth is divided into three different performance classes: Class-1 devices work with 100 milliwatt (mW) and have a range of up to 110 yards (100m). Class-2 devices work with 2.5 mW and have a range of up to 10 yards (10m). Class 3 manages with 1 mW and reaches between 5 inches (10cm) and 1 yard (1m). This relatively short range has the advantage that the transfer channels are not blocked by Bluetooth devices operating from a greater distance.

8 Named after the Viking prince Harold Bluetooth, who unified Denmark, Sweden, and Norway in the 10th century.

9 Go to *www.bluetooth.com.*

11.4.4 Mobile Communications

For the mobile business applications, in particular, which are provided by SAP with its NetWeaver product, the data services of mobile communications providers are an interesting alternative for replication—that is, replication between the mobile client (mostly a Personal Digital Assistant, PDA), and the SAP Mobile Engine Infrastructure Server. The data quantities to be transferred are typically so small that the mobile technologies currently available have no problems with them. However, even these relatively small volumes of data, which are usually ignored when designing a network, can become an issue, especially when it comes to connection costs.

11.5 Voice—Data Convergence

One area in which the infrastructure consolidation has rapidly developed in recent years is the merging of voice and data services. After all, this is not very surprising because the transmission of information through electronic signals is really integral to both concepts.

All through one cable

One of the reasons why Ethernet has become more popular than technologies such as TokenRing was the development of 10BaseT by Hewlett-Packard, where, instead of coaxial cables (10Base2 and 10Base5), twisted-pair cables of category 3 could be used, which at the time corresponded to the existing telephone cables used in the US.

American-type phone cable consists of two pairs of separately twisted wires. Alternatively, the telephone cables predominantly used in Europe consist of four wires that are twisted together (see Figure 11.6). This structure results in a stronger crosstalk that obstructs a usage for the Ethernet.

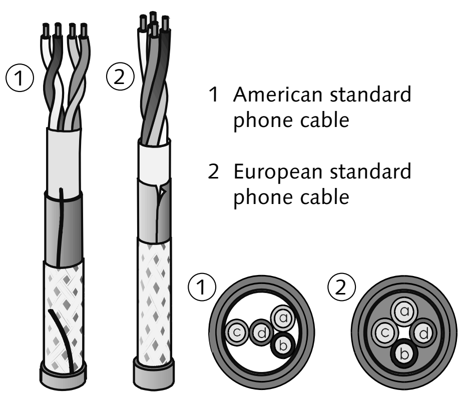

Figure 11.6 American and European Telephone Cables

Meanwhile, at least in company networks, the quantity of bytes for data transfers has exceeded that for voice communication by far. It is no longer about transferring data via a modem through proprietary telephone networks, but rather about transferring voice through open IP infrastructures (Voice over IP, VoIP). Here, one advantage is that, due to suitable compression algorithms (codecs), the necessary bandwidth for a telephone conversation is so low that it can be easily "saddled" on the normal Ethernet connection of an SAP user. However, the particular requirements of language services must be considered here, especially with regard to latency. Therefore, the use of VoIP technologies depends on the constant availability of Quality of Service (QoS) in the IP infrastructure of a company.

Power supply for IP telephones

Another important difference between language and data networks is that in conventional telephone systems the end devices are generally provided with the necessary operational voltage through the connection cable. Even if a PC with headset was perfectly sufficient (and offers substantially more functionalities), experience shows that users don't like to be without their familiar phones on their desks; admittedly, these phones don't have to be booted. These problems can be solved through patch fields, which superimpose a direct voltage on the high frequency data signals (Power over Ethernet, PoE) to supply power to the IP telephones. If, on top of that, the IP phones are also daisy chained into the connection of the PC, only one Ethernet connection per user is necessary.

Voice-data convergence needs high availability

In this way, the consolidation of voice and data transfer can drastically reduce the costs for the local network infrastructure. However, we know from experience that, in order for VoIP to be accepted by the users, availability must be guaranteed, which is akin to that of the familiar telephone, and it can only be achieved with the concepts described above.

11.6 Summary

Modern LAN technologies provide sufficient bandwidth to connect a large number of users to an SAP system. However, there are certain requirements to be considered regarding reliability:

- Design the network backbone as redundant, but be aware of the threats caused by network loops.
- Implement highly available network clusters and error-tolerant meshed networks together and patch intelligently.

- Do not forget to equip all network cabinets with an uninterrupted power supply.
- For connections between buildings and in the rising mains area, all fiber optics cables are required due to their lack of sensitivity to lightning strokes and ground loops.
- The quality of the installed cables and the proper installation has a significant impact on the performance of your network which is usually underestimated. The wiring of a floor or entire building is a major investment. Using low quality cables or unqualified installers can void this investment to a great extend, leading to significant cost in future. A wiring investment should be planned as carefully as a hardware investment project.
- The cabling for power supply also plays a pivotal role for disruption-free operation. Neutral wires and protection wires should never be used together (PEN); between building parts with different grounding potential, only fiber optics cables should be used.
- Radio networks are suitable for connecting individual mobile SAP applications. For large numbers of users, a distribution across several access points is necessary, which requires a well devised channel layout plan.
- WLANs must be integrated into the security concept of the company.

Index

D

M

T

U

V

W