SAP®
Document

SAP Technical Infrastructure

Network Integration of SAP Servers

SAP AG
Neurottstr. 16
D-69190 Walldorf

Copyright

Icons

The following icons are used in this document as visual aids.

Icon	Meaning
	Caution
	Example
	Note
	Recommendation

Version

June 2001
by SAP Technology Development – Landscape Implementation Development

Contact

For more information, see `http://service.sap.com/network` or mail to `network@sap.com`

Contents

1 Introduction

Client/server applications, such as the SAP R/3 System, communicate with each other mainly using network connections. Therefore, the quality of the network connection has a very strong influence on the stability and performance of the entire SAP System.

For this reason, integrating an SAP System into an existing corporate IT infrastructure is a demanding task.

It is particularly important when introducing or expanding an existing SAP System to begin the necessary planning early.

You need Basis specialists, and operating system and network specialists for this task. These three areas are often independent of each other.

This document tries to connect these areas and focuses on the following target groups:

- Basis specialists
- Network specialists with SAP contact
- IT project managers and consultants in the SAP environment

1.1 Structure of this Document

The Network Integration Guide is divided into the following main subsections:

I. Introduction

The section Basics of SAP Communication [Page 6] describes the architecture of an R/3 System and the network requirements that result from it, as well as the different communication connections for an SAP System.

II. Starting with network planning

The section Planning and Conception [Page 16] describes a general procedure for SAP network planning, and an extensive list of all the important points during the planning. There is also an explanation of network topology [Page 19] in this section.

III. Comprehensive documentation

This section describes in detail the following topics:

Network integration of SAP servers [Page 22]

IV. **Examples of network topologies**

In this chapter, you can find examples of network topologies for small, mid-sized and large SAP Systems. [Page **Error! Bookmark not defined.**]

2 Basics of SAP Communication

2.1 R/3 Architecture

The SAP R/3 System has a three-tier client/server architecture. All data is stored in a database, and the data is processed in the application layer on the application servers. The SAPgui frontend (presentation layer) is the interface to the user. All three layers are connected to each other with networks. The following graphic depicts the client/server architecture of the SAP System, and the communication requirements between the presentation and application layers and between the application and database layers:

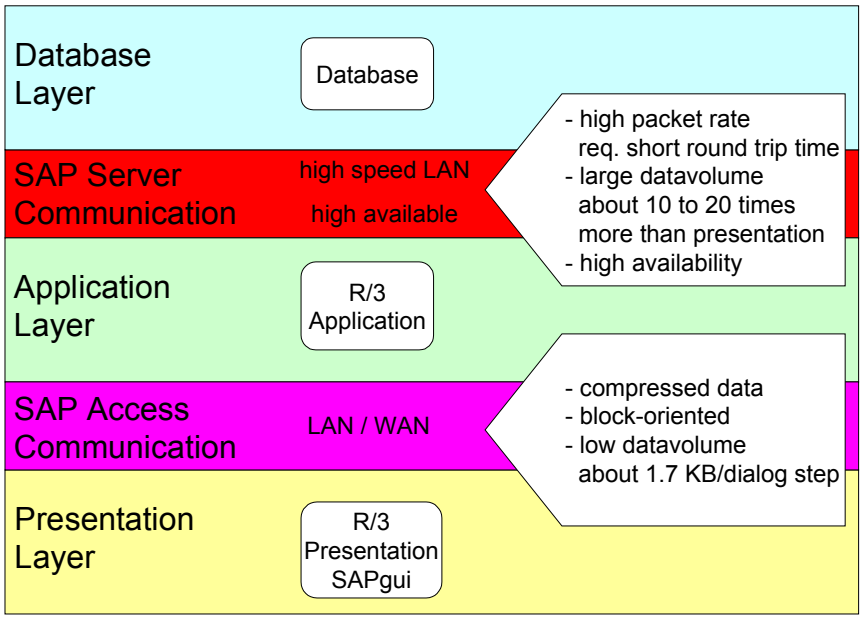

Depending on your requirements, you can distribute the services to different hosts.

Smaller applications keep the database and the SAP application on the same host. The large volume of data that passes between the SAP application and the database (SAP server communication) is processed locally and not through a network.

The presentation layer is usually made up of PCs on which the SAPgui frontend is installed. The SAPgui is not a terminal emulation but an application program that displays SAP application data graphically. This means that there are no great demands placed on the connection between the SAPgui frontend PCs and the SAP application (access communication).

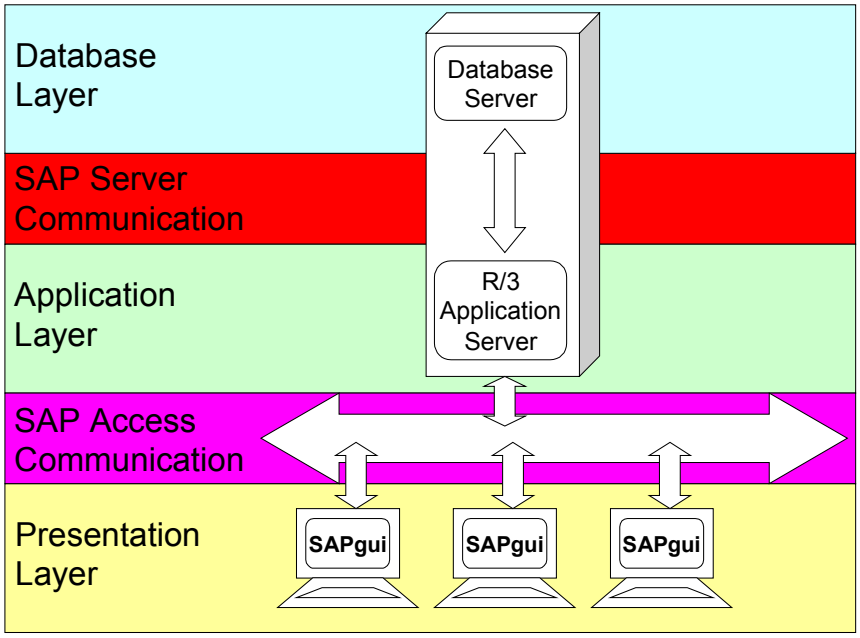

Higher processing demands on the SAP application can be realized by additional application servers (application servers are hosts on which the SAP application runs).

Very high demands are made on bandwidth and delay time between the application servers and the database server. You need to set up a suitable network connection to meet these demands.

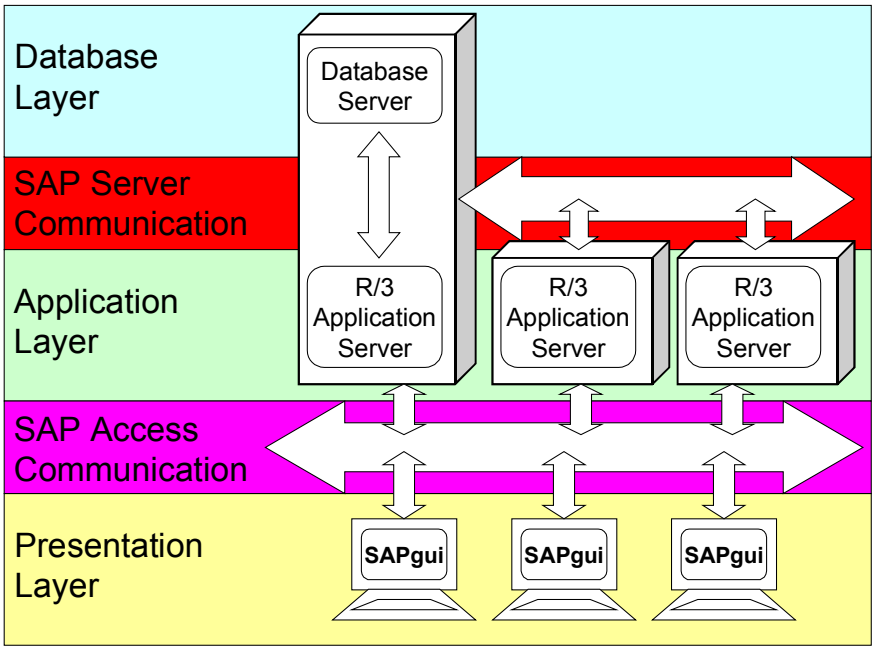

You can speed up and secure data throughput to the database by placing the database on a separate host. The database server host then communicates only with the SAP application servers. By isolating the database completely from the rest of the corporate network, you prevent unauthorized access to sensitive data and ensure high performance.

For data backup purposes you may need to connect the database server to a dedicated network (SAN = Storage Area Network).

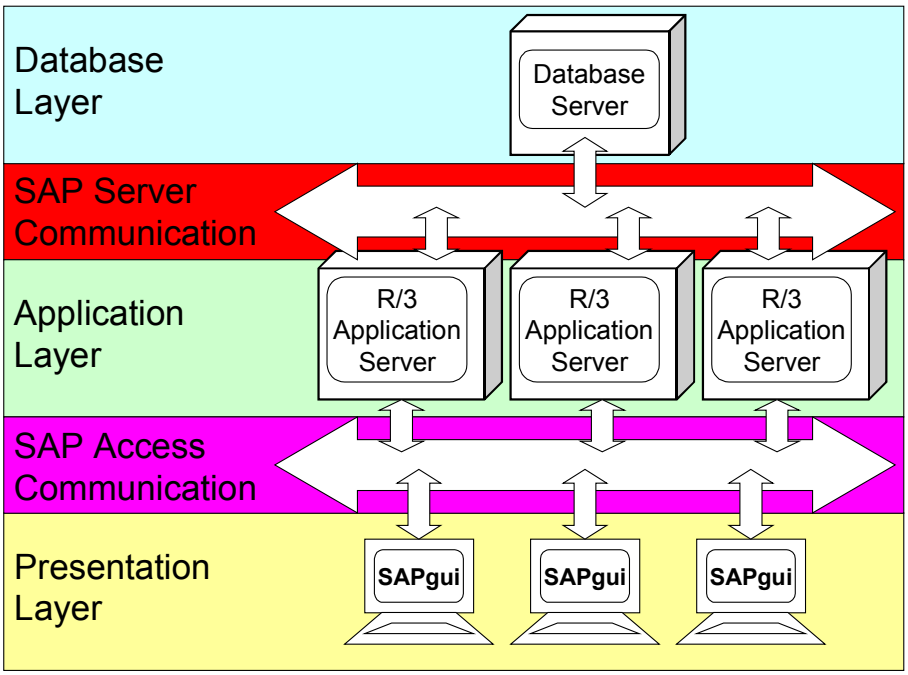

2.1.1 Access Communication

Access communication covers all access to the SAP System. This includes user access through the SAPgui, and also links to other SAP Systems and external applications.

An access network is not a dedicated network segment, but includes all network segments through which the SAP System is accessed.

Important Design Criteria for the Access Network:

- Block-oriented data traffic → delay times in the network are relatively non-critical for the SAPgui
- Bandwidth must be determined separately for each location, depending on the number of users and their activities
- The availability of the network must also be specified according to the user group or location

2.1.2 Server Communication

Server communication covers all communication between the application servers and the database server, and is of great importance for the SAP System. In each individual case you must decide whether you want to process the server communication through its own physical network (server network), or whether you want access and server communication to share a physical network. You must remember that any worsening in server communication also has a negative effect on the performance of the SAP System.

A server network is the network connection between the servers (application servers and database servers) of an SAP System.

Important design criteria for the server network:

- High throughput of data or high bandwidth
- Minimum delay time (round trip time)
- Includes the servers of the SAP System only
- High availability
- No non-SAP data traffic (for example, data backup over the network)
- Direct server connection → no expensive cabling
- Greatest possible security against unauthorized access to the database server if the server network is set up as a separate segment from the rest of the corporate network

2.2 Communication Connections of the SAP System

The SAP System supports the following communication connections:

- Presentation frontend to the SAP System
- Connections from the SAP System to printers
- Connections to other SAP Systems
- Connections to external applications

You can also use these connections with SAProuter.

The following communication connections exist within an SAP System:

- From the application server to the database
- From the application server to the message server
- Communication between application servers

2.2.1 Representation of External Communication Connections

The following graphic shows you all the different ways an SAP System can communicate. The graphic includes TCP connections only. The SAP System can consist of one or more different hosts. The communication connections within the SAP System are described later.

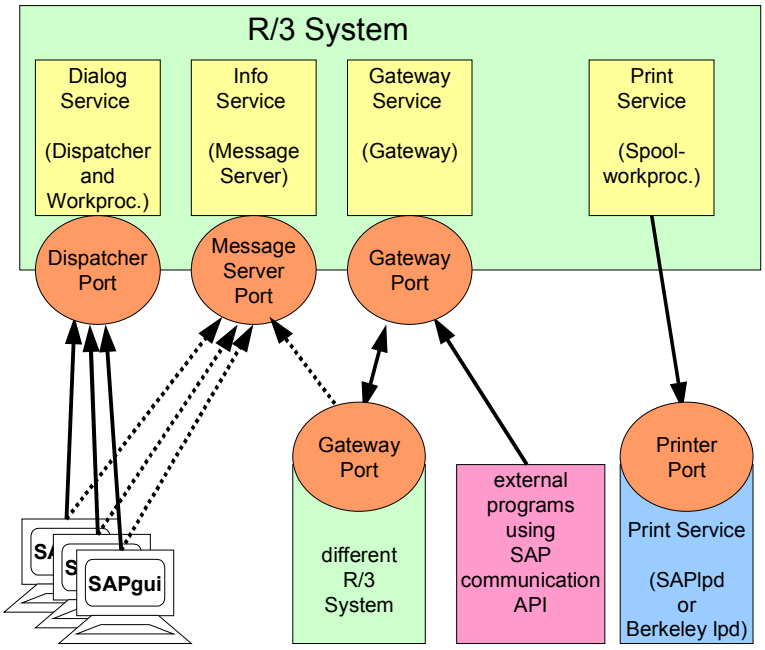

The arrows show in which direction the TCP connections are set up. The following sections describes these connections in detail.

Connection to Frontends

A dispatcher runs on each application server that can connect the SAPgui clients with dialog work processes, if needed. You can access this dispatcher under the TCP port `sapdp<nr>`, where `<nr>` is the instance number (00 to 99) of the application instance. The default value is `sapdp00` and the corresponding TCP port number is 3200.

In each SAP System there is one info service that can be used for the variable assignment of SAPgui frontends to an application server (dispatcher). This service is provided by the message server. You can use this info service to organize the frontends into groups (logon groups) and distribute them according to the load on the various application servers. The message server normally runs as part of the central instance of the SAP System and you can access it with the TCP port `sapms<SID>`. `<SID>` is the system ID of the SAP System (for example, C11). You can choose different TCP port numbers for different SAP Systems. They usually start at 3600 and increase by one each time.

Default values for instance number 00:

	Service	**TCP port**
Dialog (Dispatcher)	`sapdp00`	`3200/TCP`
Info (Message server)	`sapmsC11`	`3600/TCP`

Connecting Printers

The SAP System uses spool work processes to process print requests. There can be one or more spool work processes, which themselves can run on one or more application servers. The way in which the spool work processes print depends on your configuration. The system can print in one of the following ways:

- Output to local operating system spooler, on the application server host (no network connection)
- Output to the standard "line printer service" (lp) of the target host (TCP service printer with the TCP port number 515)
- Output to the SAP printer daemon (SAPlpd) that is running on the target host (TCP port number 515)
- Print output using the dialog connection of the SAPgui frontend. The existing SAPgui connection is used, which means that no specific port is needed.

2.2.2 Connections to External Partners And Other SAP Systems

Each SAP application server has an SAP gateway that it uses to communicate with other SAP Systems, or with other applications that use the SAP communication interfaces for CPI-C or RFC.

- **CPI-C** (*Common Programming Interface for Communications*) is a standardized protocol for the simple transmission of data between two programs. CPI-C uses half duplex transmission. This means that only one party at a time can transmit data. The receiving party must keep receiving data until it is its turn to transmit.
- **RFC (*Remote Function Call*)** is SAP's own communications interface. RFC communication always involves a caller (RFC client) and a receiver (RFC server). The RFC server provides one or more function modules that can be called. An RFC client can call one of these function modules, transmit data, and then read the results of the function module. Both the server and the client can be either external programs or the SAP System.

You can also use the gateway for communication between two applications within an SAP System. You can access the SAP Gateway on every application server under the TCP port **`sapgw<nr>`**, where **`<nr>`** is the instance number of the application instance.

Default values for instance number 00:

	Service	**TCP port**
SAP Gateway	`sapgw00`	`3300/TCP`

2.2.3 Connections with SAProuter

The SAProuter program helps you to set up and control remote connections easily. SAProuter functions as an intermediary between the source system and the target system. First, the source system sets up a TCP connection to SAProuter and tells it to which target system it wants to connect. SAProuter then sets up a connection to the target system. A SAProuter connection from a source system to a target system consists of multiple TCP connections.

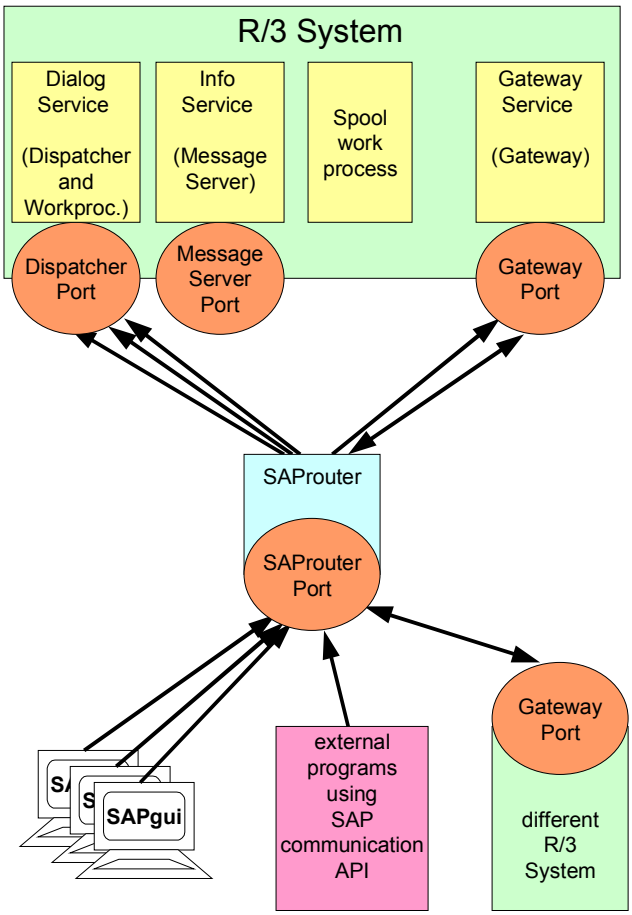

If necessary, a connection can be made using multiple SAProuters in sequence. If this is the case, all SAP connections between the SAProuters are relayed through defined TCP ports between defined partners (SAProuters). This makes it much easier to implement mechanisms that protect access to the connections.

Separate organizational units that administrate their networks and security separately can also take advantage of double SAProuters. Each unit defines its own SAProuter as a communication port for all SAP communication that goes outside its limits. SAProuter then becomes an **application level gateway** that can run, for example, on a firewall host. You can use authorization tables to control who may set up a connection with whom.

Each organizational unit can allow or disallow connections as it wants.

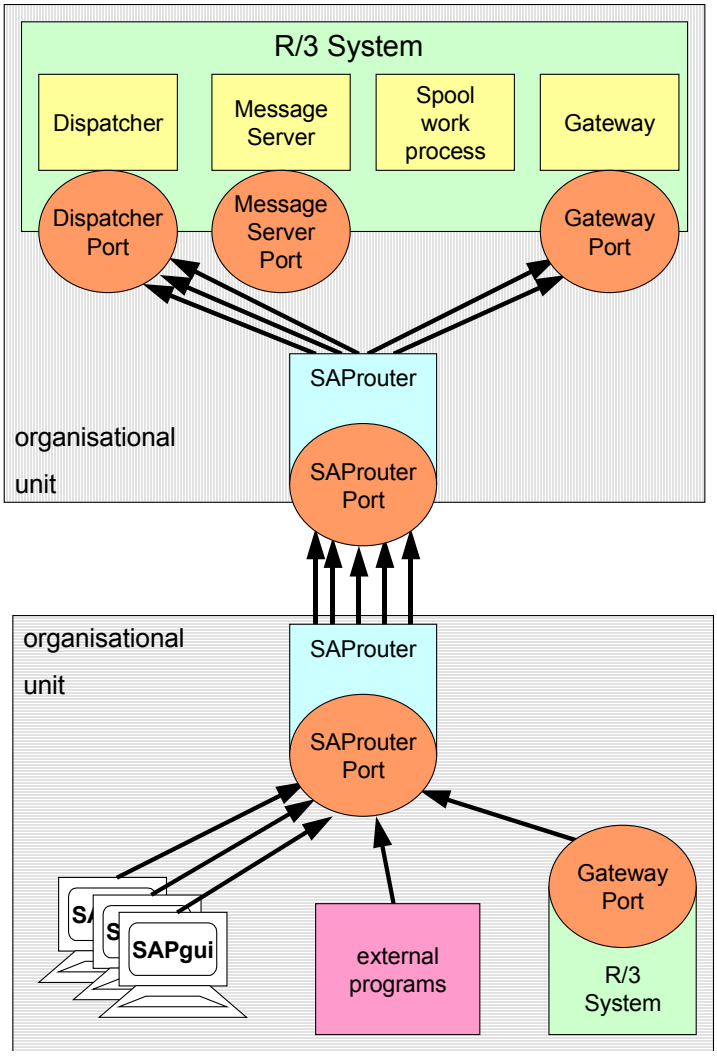

When you start the program, you can determine under which port you want to access the SAProuter.

Default values:

	Service	**TCP port**
SAProuter	`sapdp99`	`3299/TCP`

2.2.4 Internal Communications

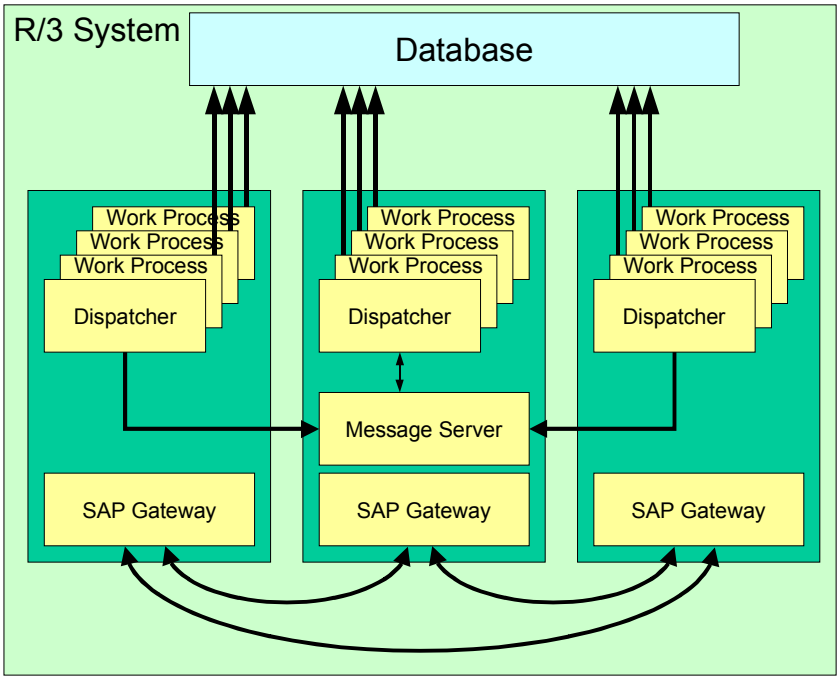

The following section describes these connections in detail:

Connecting to the Database

The connection between the SAP application servers and the database server is based on a Remote SQL database interface.

The TCP port you use depends on the individual database system:

Database system	Service	TCP Port
Oracle	`tlisrv`	`1527/TCP`
Informix	`sapinf<sid>`	`3800/TCP`
MS SQL Server	`ms-sql-s`	`1433/TCP`
SAP DB	`sql30` `sql6`	`7200/TCP` `7210/TCP`
DB2 Universal Database for UNIX and Windows	`sapdb2<SID>` `sapdb2i<SID>`	

Connecting to the Message Server

The various instances (application servers) use the message server to access central services, such as the enqueue or update services. The message server also receives regular information from all instances about current system load and the services that are on offer. This means that the message server can provide information on load balancing for different applications, such as the SAPgui or background applications.

When an SAP application server is started it sets up a connection to the message server of the central system (more precisely, each dispatcher sets up connection to the message server port sapms<SID>).

Connecting Application Servers

If application servers need to communicate with each other, a connection is set up between the SAP gateways of the application servers (port sapgw<nr>). Once it is set up, this connection remains until the application server stops running.

2.2.5 Table of Communications Connections

Service	Default TCP service name	Default TCP port number	Possible range for TCP service name	Possible range for TCP port number
Dispatcher	sapdp00	3200	sapdp00 – sapdp99	3200 – 3299
Message server	sapms<SID>	3600	sapms<SID>	Free choice
SAP Gateway	sapgw00	3300	sapgw00 – sapgw99	3300 – 3399
SAPlpd	-	515	-	Free choice
SAProuter	-	3299	Free choice	Free choice
Test program niping	-	3298	Free choice	Free choice
Oracle database	tlisrv	1527		
Informix database	sapinf<SID>	3800		
MS SQL Server database	ms-sql-s	1433		
DB2 UDB for UNIX and Windows	sapdb2<SID> sapdb2i<SID>			Free choice
SAP DB database	sql30 sql6	7200 7210		

3 Planning and Conception

3.1 Planning a Network

When you plan a network you need to consider all possible demands that might be placed on the network infrastructure by the SAP System. Experience has shown that a completely new network is not usually set up for the SAP System, instead it is integrated into an existing network infrastructure. Here as well, you need to compare the demands of the SAP System with the characteristics of the existing infrastructure, and make any necessary adjustments or extensions.

The following graphic explains the areas that are affected, and where they overlap. Network planning needs to be a joint effort between several different specialist departments.

You should plan your network in two stages:

1. SAP application and Basis specialists decide the general and customer-specific network requirements of the SAP System (see *Important Aspects when Planning Your Network*).
2. Network and operating system specialists work to meet these requirements.

   a.) Network topology and network components (routers, hubs, switches, servers, if necessary with multiple network connections)

   b.) Network configuration of the serves (IP addresses, routes, host names)

   c.) Installation and configuration of the SAP System

   d.) Test.
   Not just basic network functions, but also the connections and conditions under full network load.

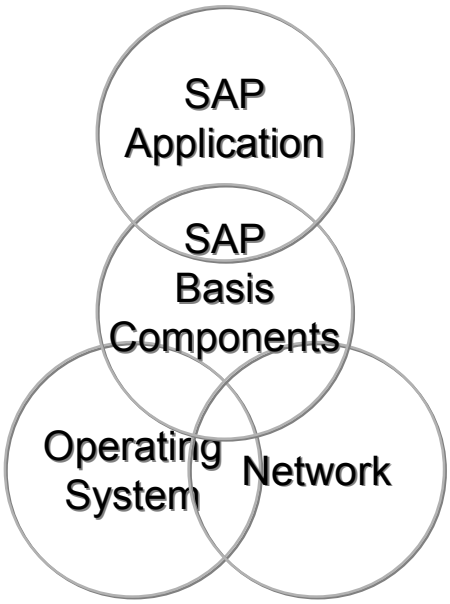

3.2 Important Aspects when Planning Your Network

3.2.1 Server Integration

Questions to be Answered:

- Does your SAP System consist of a single host or are there additional application servers?
- Does the database run on a separate host or together with an SAP instance?
- Where are the SAP servers located?

Recommendations and Tips:

- ❑ Ensure that all the servers in an SAP System are located centrally, and that there is no WAN between the servers.
- ❑ For SAP Systems with several application servers, you must determine if the servers for this SAP System (database server and application server) are to be connected to each other by an additional, unique and separate network segment (server network).
  See: Network Topology [Page 18]
- ❑ When choosing the host names for the servers, IP addresses, static routes, see Host Names of SAP Servers [Page 25].

3.2.2 Frontend Integration

Questions to be Answered:

- Where are the branch office SAP users located and how many are there?
- What is the network bandwidth required for the branch offices connected over a WAN?
- What is the procedure for installing and upgrading the frontend software (SAPgui)?
- Which additional standard applications (e-mail, Internet browser, emulation, Office) will you use on the frontends?

Recommendations and Tips:

- ❑ A stable TCP connection to the SAP Server is a prerequisite for frontends.
- ❑ If the frontends are connected to the SAP Server using a LAN, the existing bandwidth is usually sufficient.
- ❑ For frontends connected by a WAN, you need to determine the required bandwidth for each location as accurately as possible.
- ❑ Also consider additional frontend integration (printing, online documentation, installation and upgrade, mobile frontends...).
  See Network Integration of SAP Frontends [External].

3.2.3 Printer Integration

Questions to Be Answered:

- At which branch offices are the printers located and how many are there?
- Which method is used to access the printers from the SAP System?
- What is the bandwidth required for the expected printing volume at the branch offices?

Recommendations and Tips:

- ❑ You can connect printers either directly to the SAP Server or position them as network or frontend printers at any location on the network (depending on the access method). When you position printers, you must consider criteria such as availability, and printing speed and volume.
  See the *BC SAP Printing Guide - Planning and Installing the SAP Print Architecture.*
- ❑ If printers are connected over a network, the same criteria applies as for frontends. This means that for all branch offices connected by a WAN, you need to estimate the bandwidth requirements for the expected printing volume. See Network Load from Printing [Page 55].

3.2.4 Links Between Applications

Questions to Be Answered:

- Are there transports between different SAP Systems, such as a development system and production system?
- Is there a temporary data transfer or connection to legacy systems?
- Are there long-term connections to other systems?

Recommendations and Tips:

- ❑ You can connect the SAP System with other SAP Systems or external systems in very different ways. The basis for the connection between the two systems is usually a TCP connection. The direction in which the connection is set up depends on the application and its configuration. The connection can generally be made in both directions and include any application server of the SAP System.
- ❑ For connections with mainframe-based legacy systems, the SAP gateway also supports SNA LU6.2 connections.
- ❑ To estimate the required bandwidth you need to know the application sufficiently well. See: Network Load from CPI-C Data Transmission [Page 55] and Network Load from RFC Data Transmission [Page 55].

3.2.5 High Availability

Questions to Be Answered:

- How high must the availability be for the SAP System (server)?
- Will you be using special high availability products such as *ServiceGuard, HACMP OBServe, MSCS* or *Solstice*?
- How high must the availability be for the different locations or user groups?
- Which dependencies arise when linking to external applications and what availability has to be ensured for these links?

Recommendations and Tips:

- If you use high availability products for the SAP Server, you must follow certain rules for the network configuration. See *SAP R/3 in Switchover Environments*.

3.2.6 Security

Questions to Be Answered:

- What security requirements exist on the database server?
- What security requirements exist for the communication between the SAP servers?
- What security requirements are made on the different connections between the clients and the SAP servers?
  - ❖ Clients in the LAN
  - ❖ Clients in the WAN
- Will you be using the *Internet Transaction Server* (ITS)? If yes, which security measures are set up for accessing clients in the Intranet, Extranet or Internet?

Recommendations and Tips:

- There are different ways of protecting access and encryption when clients access the SAP server. See Network Security [Page 39].
- The connection between the SAP servers (database server and application server) is normally unprotected. Therefore, you must protect these servers using network tools (such as a separate network segment for the server communication). See Network Security [Page 39].

3.2.7 Data Backup

Questions to Be Answered:

- How is the data backup of the database server organized?
- How is the data backup of the application server organized?

Recommendations and Tips:

- When operating an SAP System, you need to continually back up the SAP database. If the data is backed up at a central point over the network for administrative reasons (for example, by a tape library), there is a heavy network load during the backup. If the backup occurs during productive operation of the SAP System (which includes background processing), you need to consider that in the server network small increases in the delay time and limitations in the bandwidth can significantly decrease the performance of the SAP System.
- If you back up the data over the network, use a separate network segment (*Storage Area Network,* SAN) set up only for the data backup. To do this, you must install an additional network adapter on the database server.

3.3 Network Topology

As already described in the section Basics of SAP Communication [Page 6], an SAP System communicates using server communication and access communication.

There is no standard solution for access communication (client access or access methods using external systems). You must find an appropriate network topology to access the SAP server depending on the type and number of clients, and particularly depending on the location of the system.

For server communication topology, you must first determine the location of the SAP server installation (database server and application server). The SAP System consists of one or several servers that are usually responsible for the entire enterprise and are located at a central point, the computer center. This is a centralized concept (as with mainframes) and is an advantage for maintenance and operation. Numerous clients that are widely spread out access these servers.

The concept of a "server farm" is well suited to SAP Systems.

Note the following basic principles:

- All servers are connected to each other directly with a high bandwidth and minimal delay times. Switches are well suited for this.
- If the servers are connected using shared media (for example, the Ethernet), there may easily be temporary network overloads, called collisions. Since this lowers performance, we recommend that you use an additional, separate network segment to connect the servers of an SAP System.
- If you connect the SAP Server with the campus or backbone network, you need routing functions. You can use routers or modern layer 3 switches (OSI Layer 3 = network layer).
- By creating redundant network paths, you can protect your system against failure.

If the network topology covers several SAP Systems, the server connections should only include those servers of one system only so that optimal performance is ensured for each SAP System.

3.3.1 Network Topology for Switched Networks

If the SAP servers are connected using switches, access to the individual servers has a high and constant bandwidth. The connection between two servers has a minimal delay time.

Therefore, you only need to connect the servers of an SAP System to each other with an additional network adapter and switch if there is extremely high throughput and bandwidth requirements (for example, if the throughput of the network adapter is not sufficient). (See SAP System 1).

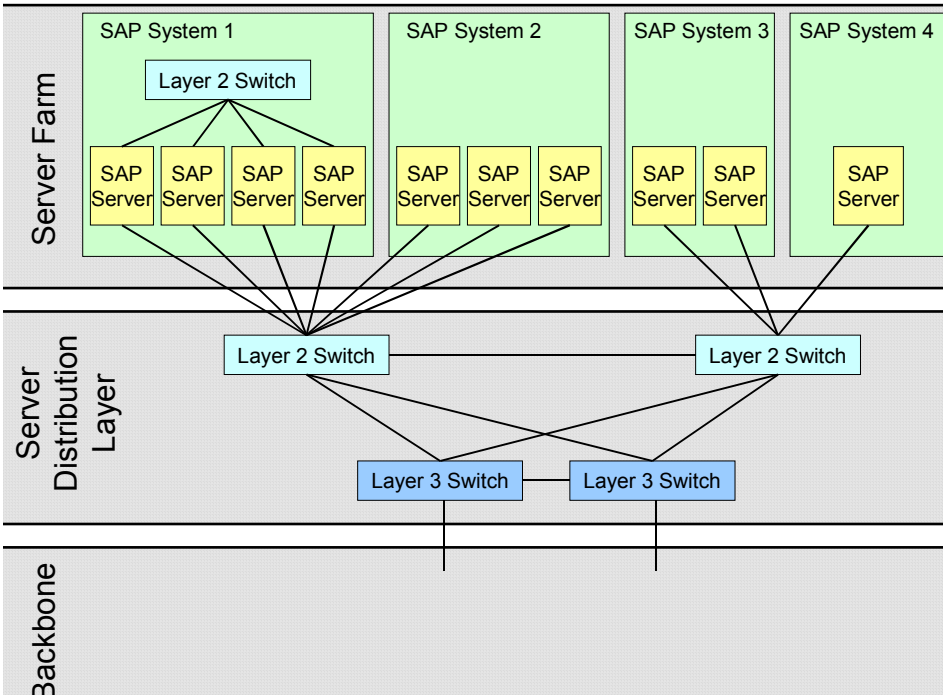

3.3.2 Network Topology for "Old Style" Shared Media Networks

For the SAP servers, you need a connection with a high bandwidth with minimum delay times. Shared media networks can contribute to this only partially. Communication between servers may lead to conflicts and network overloads. In this case, we recommend an additional network segment that connects the servers of an SAP System. We also recommend that you use a switch to connect the servers (see SAP System 1) rather than a shared media connection (see SAP Systems 2 and 3).

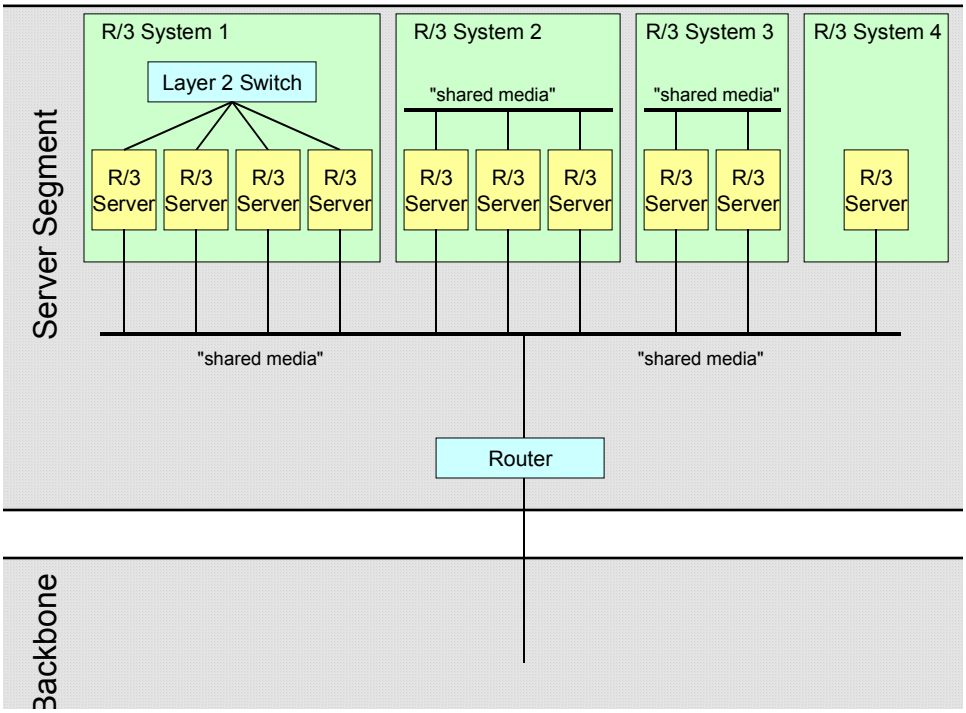
Server Segment
R/3 System 1
Layer 2 Switch
R/3 Server
R/3 Server
R/3 Server
R/3 Server
R/3 System 2
"shared media"
R/3 Server
R/3 Server
R/3 Server
R/3 System 3
"shared media"
R/3 Server
R/3 Server
R/3 System 4
R/3 Server
"shared media"
"shared media"
Router
Backbone

4 Network Integration of SAP Servers

4.1 Basic Principles of TCP/IP Configuration for SAP Servers

To transmit data between two hosts in one network, you need two **network adapters** (network interface card, NIC). One adapter sends a data packet and the other one receives it.

An **IP address** (IP = Internet protocol) is assigned to each network adapter. These numerical IP addresses are represented in the form *nnn.nnn.nnn.nnn*, where *nnn* is a number from 0 to 255 (for example, 192.168.1.10). An IP address functions as a unique ID for identifying the sender and the receiver in a network with multiple hosts.

The **IP name** is the name under which the network adapter is known in the network. An **IP name** is assigned to each IP address. This simplifies the use of IP addresses. This IP name is often mistakenly described as the host name.

The **host name** or computer name is a logical name for the host itself and is set in the operating system. The operating system knows its own host under its host name. You can query this name under UNIX or Windows NT by using the command `hostname`.

The IP name is mapped to the IP addresses using a name database. Here, it can be a local host file, or a name service, such as DNS (Domain Name System/Service) or WINS (Windows Internet Naming Service).

If the relevant host only has one network adapter, the host name and the IP name are normally identical. In this case, there is a unique relationship between the logical host name and the IP name (or IP address). This means that this host can be addressed by other computers using its host name.

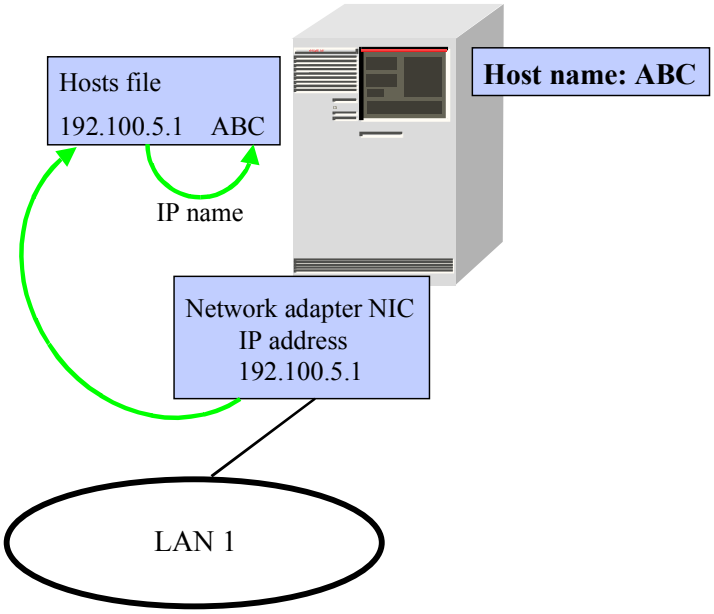

Terminology

SAP server	All hosts (computers) that belong to an SAP System (database server and application server)
Database server	Host on which the database service runs. If an SAP instance (central instance or dialog instance) also runs on the database server, the server must be handled like an application server when you configure it.
Application server	Host on which an SAP application service runs (central instance or dialog instance)
Server network	High speed network segment through which the SAP servers communicate with each other
Access network (frontend network)	All network segments through which the frontend is connected with the application servers
IP address	Network address in the Internet protocol (for example, 192.5.2.1), assigned to a network adapter
IP name	Logical name for an IP address. Is often equated with the host name, although the relationship of the host name and the IP name is not always unique.
Host name	Logical name of the host. It is normally assigned like the IP name of an IP address for the host. A host can have several IP addresses and IP names, but it can only have one host name.

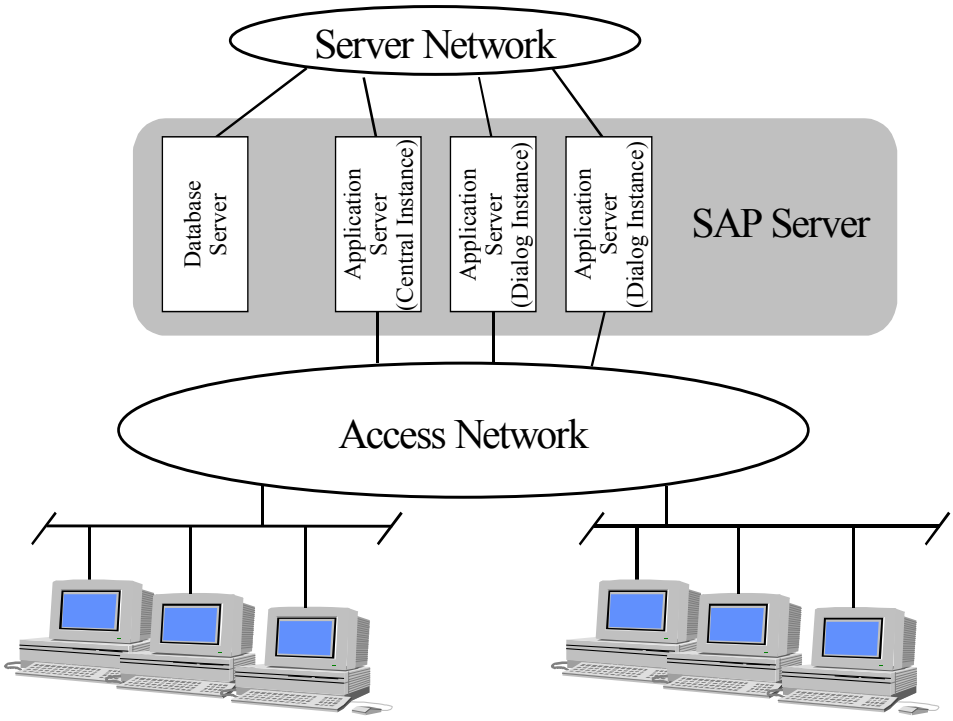

4.1.1 Multihomed Hosts

If a host has multiple network adapters (multiple NICs) or multiple IP addresses for each network adapter, it is called a **multihomed host**. In this case, the relationship between the IP addresses (with the assigned IP names) and the host name is no longer unique. The host name can be identical to the IP name for an IP address, or it can take on its own unique name. If a multihomed host is addressed by another host, one of the IP names (or the corresponding IP address) must be used. In the following graphic, LAN1 must use the IP name ABC to address the host ABC, whereas LAN2 must use the IP name XYZ to address the same host.

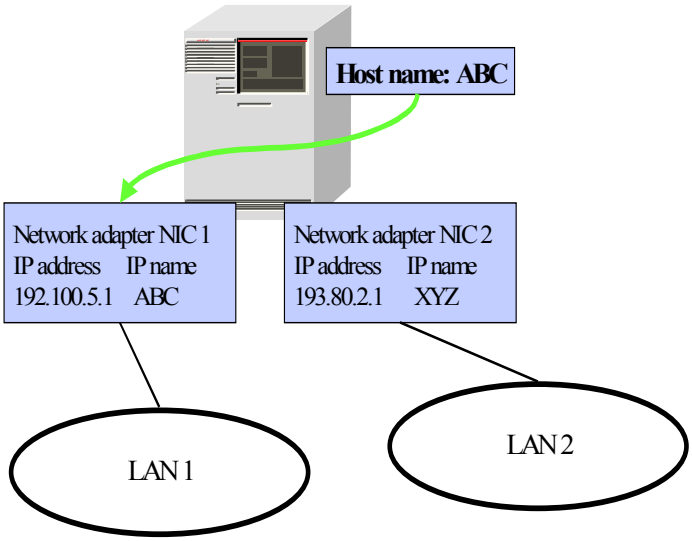

The name you choose by which a host is addressed, influences the route of the data traffic in the network. Application programs do not usually know the various IP names, or their accompanying LAN segments, but they do know the logical host name. Therefore, choosing the right host name is very important. If static IP routes are used, you can also access IP names and/or IP addresses that are not in your own subnet.

4.1.2 Using IP Routing to Choose the Route

If you want o use TCP/IP to send a data packet, the network layer (layer 3 in the OSI layer model or the **IP layer)** must determine if and how the target IP address can be accessed. The routing information necessary for this is stored in the routing table. You can set the entries in the routing table as being static (manually) or dynamic (using the routing daemon process). The entire route to the target address is not described in the routing table, instead, only the route up to the next node (next hop) to which the data packet is forwarded. There, a check is made in one of its own routing tables as to how the target address should be accessed. This is called an *indirect route*, since the route goes through at least one **gateway**. If the address can be directly accessed, this is called a *direct route*.

There are three ways of defining the target address as the routing entry:

- Complete IP address of the target host (**host route**)
- Only the network portion of the IP address (**network route**)
- Default route (**only one entry**)

Example of a Simple Route

In this example, the three hosts are connected together by network A and network B:

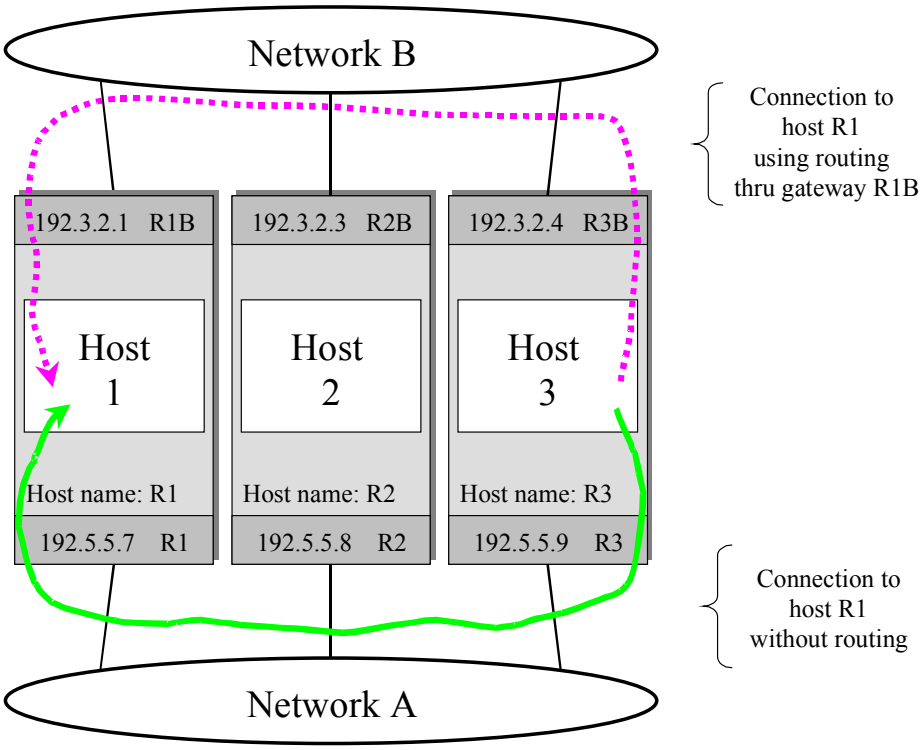

A host is usually addressed by its host name. If Host 3 makes a connection to Host 1, the target address 192.5.5.7 (R1) is addressed. By default, this connection goes over network A. If you want the connection to the same target address to run over network B instead, you must have a corresponding routing entry in Host 3. You may want to route the connection over network B, for example, to increase performance (network B is faster than network A). For communication to run in both directions over network B, there must also be a corresponding routing entry for Host 3 on Host 1.

Example of Routing With UNIX

The above example described an indirect host route. The target address is a complete IP address (a host route), and it is also accessed through a gateway using the IP name R1B (an indirect route). On UNIX hosts, you can display the routing table by using the command **`netstat -r`**. For indirect routes, the flag **G** is used, and for host routes, the flag **H** is used. To create an indirect route, you must specify in the command **`route add`** a *Metric* (hop count) greater than 0.

Action	Command
Create route on host 1:	`route add R3 R3B 1`
Create route on host 3:	`route add R1 R1B 1`
Display all static routes:	`netstat -r`
Test routes:	`traceroute`

Host 3:

Destination	Gateway	Flag
R1	R1B	UGH

Host 1:

Destination	Gateway	Flag
R3	R3B	UGH

Example of Routing With Windows NT

On Windows NT hosts, you can display the routing table by using the command `route print`. Indirect routes are indicated by a *Metric = 1* and are made defaults with the command `route add`.

Action	Command
Create route on host 1:	`route add R3 R3B 1`
Create route on host 3:	`route add R1 R1B 1`
Display all static routes:	`route print`
Test routes:	`tracert`

Host 3:

Network address	Gateway address	Metric
R1	R1B	1

Host 1:

Network address	Gateway address	Metric
R3	R3B	1

4.2 Host Names of SAP Servers

The host name of SAP servers is used in many places, for example, during installation, in the start scripts, and for communication between servers. You can choose this name if you observe certain rules, which are explained in the following section. You must follow these rules for your SAP System to operate smoothly and efficiently.

Read and take these rules into consideration **before** you install an SAP System.

As a prerequisite, you must have basic knowledge of the TCP/IP protocol.

See also: Basic Principles of TCP/IP Configuration for SAP Servers [Page 22]

4.2.1 Requirements for Host Names and IP Names

The host name must be assigned **before** you install the SAP System. SAP does not support a name change after installation.

If you use only one network card, the host name must match the IP name that is assigned to this network card.

See also: Configuring SAP Servers with Multiple Network cards [Page 32].

4.2.2 Rules for Host Names for SAP Servers

The host name of an SAP server is a text string that can consist of letters (a-z), numbers (0-9), and the minus symbol (-). The last character **cannot** be a minus symbol. The host name is not case-sensitive. You cannot use a period (.) as part of the host name; you can only use it in the full qualified domain name.

Up to and including Release 4.5, the host name can have a **maximum length of 8 characters**.

Except for the maximum character length, this definition corresponds to the Internet standards (standards RFC 952 and RFC 1123).

The full qualified domain name (for example, `r3app01.fiboflex.com`) can have a maximum of 60 characters.

4.3 Name and Address Resolution for SAP Servers

4.3.1 Overview

The SAP System uses IP names and IP addresses to identify SAP servers and external communication partners. Therefore, the correct configuration for resolving IP names and IP addresses is extremely important for the SAP System to function properly.

Read and take these rules into consideration **before** installing an SAP System.

This document explains the requirements of the SAP System for resolving IP names into IP addresses and vice versa. All the operating system platforms supported by SAP offer various mechanisms for name and address resolution (for example, the Hosts file and DNS). These are explained briefly in the second part with respect to the SAP System.

As a prerequisite, you must have basic knowledge of the TCP/IP protocol. (See also: Basic Principles of TCP/IP Configuration for SAP Servers [Page 22].)

4.3.2 Terminology

Translating an IP name into an IP address or the reverse, is called **resolution**. This can occur in two directions (see graphic). The most common case involves resolving an IP name into an IP address. For the SAP System, the opposite resolution is also very important (this is called **reverse lookup**).

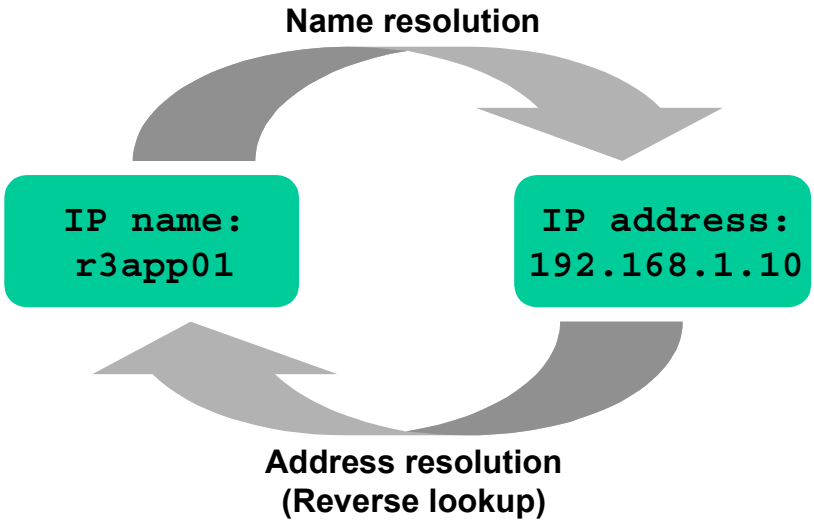

4.3.3 Requirements for the SAP System

Due to the client/server architecture of the SAP System, there are specific requirements for SAP servers and SAP clients when resolving IP names and IP addresses. This document describes the requirements for the SAP Releases up to and including 4.5.

SAP Servers

The sections Host Names for SAP Servers [Page 26] and Configuring SAP Servers with Multiple Network Cards [Page 32] list the rules for assigning host names, IP names, and IP addresses for SAP servers. Follow these rules to ensure smooth operation of the SAP System.

You must meet the following requirements for your SAP System to function properly.

- The IP names used to configure the application servers must be able to be resolved into an IP address by all the other application servers in the same SAP System. The assignment must be identical on all servers.
- Reverse lookup for these IP addresses must be possible and they must return the same IP name. Resolution in the sequence *IP name* → *IP address* → *IP name* must yield the identical name.
- The IP name may not be an alias. On Windows NT, the sequence of the network adapters must be configured correctly in the system.
- The name *localhost* must be able to be resolved (normally, the accompanying IP address is 127.0.0.1).

The following requirements also apply for communication functions:

- RFC or CPI-C client programs connect externally to the SAP System. In the SAP System, the IP address of the client host does not have to be resolved. This also applies for RFC server programs that register on the gateway or that start from a SAPgui. One exception is the gateway security functions, which can be activated in the file `secinfo` (for more information, see the SAP Library under BC - SAP Communication [External]).
- The SAP gateway may start RFC or CPI-C server programs on another host. If this is the case, reverse lookup from the IP address of the host where the server program is started to the corresponding IP name must be possible. If you want to use host names to configure the RFC destinations, you must ensure that the IP name can be resolved.
- You do not need additional name resolution to print on the frontend, since the existing connection between the application server and SAPgui is used. If you want to address a printer or print server using names, these names also have to be resolved on the SAP server. If you use one or more print servers, the names of the print servers do not have to be resolved by themselves.

In addition to the name resolutions necessary for the system to function, other reverse lookups occur that are not functionally necessary. This type of name resolution occurs, for example, so that the actual names are displayed in monitors and in the system log, instead of the numerical IP addresses.

- When the SAP System is started, a reverse lookup is performed for the IP addresses of all the network adapters that are built into the application server. The IP addresses of the network adapters from other application servers are also resolved under certain circumstances.
- The SAP System tries to resolve the IP addresses into IP names for all frontend computers that use SAPgui to log on to the system directly, and for all RFC and CPI-C client computers.

In these cases, the resolution does not have to return a valid result for the SAP System to function properly. However, if the configuration for resolving names is incorrect (in the operating system or in the name server), there may be long wait times due to timeouts, which can have a highly negative affect on SAP System performance. Therefore, ensure that the configuration is correct when using name resolution with DNS (Domain Name System).

Buffering Names and Addresses in the SAP Server

The SAP Network Interface (NI) is a part of the SAP System, and stores the result of the name resolution and reverse lookups in a special buffer. This means that the operating system only has to perform the name resolution once. Successful resolutions remain in the buffer until the SAP System is shut down, or until an overflow, or until the buffer is manually deleted (Transaction SM51). Manually deleting the buffer does not include the message server. For this reason, restart the SAP System to make the changes of IP names or IP addresses effective in the system.

If a resolution fails, this is noted for a duration of 10 minutes (as of SAP Release 4.0). This is because the name server may not be available. If another attempt is made to resolve the name after this period of time, the SAP System repeats the operating system call.

The advantage of buffering in the SAP System is that it increases the performance during name resolution and it only very rarely reaches a critical situation.

SAP Frontends (SAPgui)

To be able to use all of the SAP functions, the frontend computer must be able to resolve the names of the message and application servers into IP addresses.

During load balancing, the SAPgui receives a list of IP addresses from the message server. Resolution does not occur.

If the frontend computer cannot resolve the names of the SAP application servers, you can also start the SAPgui by using numerical IP addresses. However, the functions may become limited, for example, when you use certain desktop components.

External RFC and CPI-C Communication Components

An RFC or CPI-C client establishes a connection to the SAP System. To do this, it requires the IP address of an application server. You can enter this address directly to establish the connection. In this case, no resolution occurs. You can also use load balancing with logon groups for RFC. Here, the RFC client connects to the message server and receives a list of IP addresses from the server. No resolution occurs here.

Server programs are usually started by an SAP gateway, or by a SAPgui. Each SAP instance has a gateway. You can also operate an SAP gateway without an SAP instance (*standalone*). After the server program starts, it must establish a connection with the gateway within a specific period of time. You can start the program in different ways:

Start using a remote shell: On the command line, the server program receives the IP name of the gateway to which it should connect. This name must be able to be resolved.

Start using the SAPgui (only RFC): The command line communicates the information with which the SAPgui was started to the RFC program. It also communicates a SAProuter string is used. This means that the resolution is identical to the resolution with the SAPgui.

4.3.4 Naming Resolution Mechanisms

There are four common mechanisms for assigning names with addresses:

- *Hosts* file
- Domain Name System (DNS)
- Windows Internet Naming Service (WINS)
- Network Information Service (NIS, used to be called *Yellow Pages*)

The operating system determines which mechanisms are available, which are active by default, and how they are configured. The administrator can decide which of the available mechanisms are used. In most operating systems, you can configure several mechanisms with different priorities. These are used one after the other, if the first mechanism does not give you the result you want.

Choosing a Naming Resolution Mechanism

The configuration of the naming resolution mechanism on SAP servers must meet the following requirements:

- Robustness
- Flexibility
- Simplicity
- Performance
- Integration in an existing infrastructure

Hosts File

The *Hosts* file contains a list of IP addresses and the accompanying IP names. The resolution occurs by searching through the list. Under UNIX, the file is called `/etc/hosts`, under Windows NT it is called `\winnt\system32\drivers\etc\hosts`.

The *Hosts* file is a robust and simple solution, since all of the information for naming resolution is available on the computer. No network access to other computers is required to resolve names. The advantage is very high performance. However, maintaining the file requires some effort, especially if there are many hosts, or the hosts are divided among many separate networks.

When you maintain the *Hosts* file, you must enter the IP name that corresponds to the host name as the first name directly after the IP address, and not as an alias (second name).

NIS

The Network Information Service (NIS) is a distributed database system that replaces configuration files, which usually have multiple copies, with a central administration. Files such as `/etc/hosts`, `/etc/passwd`, and so on do not have to be administrated in each host. Instead, the information is retrieved as needed from the database.

If you use NIS to administrate the *Hosts* file, only the settings in NIS are valid. The local *Hosts* file is no longer used. This behavior depends, however, on the operating system.

NIS is a centrally administrated system with a master server. To improve availability and performance, you can set up additional slave servers that answer the client queries. Since access to an NIS server is required for each resolution, you must have high availability, short wait times, and you must carefully plan the NIS system for your system to work smoothly.

WINS

WINS (Windows Name Service) is suited for use in a pure Microsoft Windows environment and cannot be directly compared with the other mechanisms described here for naming resolution. One main difference is that no IP names, but NetBIOS names, are mapped to IP addresses. The advantages of WINS are its simplicity, robustness, and flexibility. The disadvantage is that WINS is only suited for smaller, flatly structured and homogeneous Microsoft networks.

DNS

The Domain Name System (DNS) can map complex, hierarchical structures with a large number of hosts. It lets you use efficient centralized and decentralized administration. For example, the entire Internet uses a single DNS infrastructure for naming resolution. DNS is the most robust of all methods of naming resolution, and is the most well suited for the future, for example, for potential integration with directory services. DNS is available on all server and frontend platforms supported by SAP.

With regards to availability and wait times, the same applies here as it does for NIS. With DNS, planning is particularly important, since configuring and maintaining DNS is complex and, therefore, prone to errors. The most dangerous errors are not those that lead to a total failure, rather those that lead to performance problems, because they remain undiscovered.

Tips for DNS Configuration

The effects of an incorrect DNS configuration are explained in the following example:

If a reverse lookup in the DNS is not maintained, and the name server in the corporate internal network is configured incorrectly so that it accesses the central name server in the Internet for unknown domains (this is the default setting for certain name servers), the name server waits for several minutes for a response from the inaccessible Internet name server. During this time, the work process from the operating system is stopped and the user cannot continue working.

Wait times can also occur due to connection problems between the application server and the name server, or between two name servers.

If you want to use DNS (you need some basic knowledge of DNS):

- You need at least one backup name server. DNS contains methods for replicating the configuration data. In all operating systems, backup name servers can be configured that are automatically used if the first server does not respond within a specific period of time.
- Ensure that there is high availability for the DNS server and for the network connection between the DNS servers and to the SAP servers.
- Ensure that all the domains in your company are entered in the DNS. To do this, you must create a zone file for all name domains (for example, `accounting.fiboflex.com`), and set up a name server for this zone as the primary name server. You also need a zone file for all network domains (for example, 192.168.1) that resolves IP addresses into IP names (reverse lookup).
- You do not have to enter all your corporate computers in the zone files, if the information is not needed for operating your SAP System. It is only important that the zones exist. The above sections have already described which name resolution is required to operate your SAP System. If, for example, you do not need naming resolution from client computers, the zone files for the network domain and naming domain of your frontend network can remain empty.
- Ensure that queries for internal names or addresses are never forwarded to external name servers (for example, in the Internet). This is guaranteed if root name servers only are contained in the DNS cache files that are located in your local network.

Conclusion

Deciding on which mechanisms you should use for naming resolution depends on your priorities. When making your decision, remember that the name resolution is a critical function in the SAP System that influences the system operation.

WINS is only used in pure Windows environments. In the future, it will not be supported in certain circumstances, and we therefore do not recommend it.

DNS is the most flexible and powerful solution, almost a standard. However, its use requires a high level of knowledge about implementation and maintenance. Buffering in the SAP System means that wait times are not critical, within certain limits. However, you must have high availability for the DNS server and the related network connections. If you already have a functioning and well-maintained DNS infrastructure, we recommend using DNS on the SAP servers and frontends.

Simplicity and robustness are the advantages of the *Hosts* file. It is therefore highly recommended for the SAP servers. Generally, the SAP server does not have to use the same mechanism as in the rest of your corporate network. However, maintaining and distributing the *Hosts* file may require a great deal of effort to administrate.

On UNIX servers, you can replace the *Hosts* file by NIS. Here, like DNS, you must pay attention to proper implementation and high availability of the NIS server to ensure that the SAP System operates properly.

4.4 Configuring SAP Servers with Multiple Network Cards

4.4.1 Assigning Server Services to Network Addresses

In a client/server environment, the clients use the network addresses or host names to access the server services. This means that a server service is assigned to a particular network address.

If a server service runs on a multihomed host, this service can be accessed through several network addresses. Since this scenario is no longer a unique assignment of a server service to one network address, this can lead to problems if the clients access it using different network addresses from different subnets.

Static Assignment

With static assignment, a client always uses the same network to address the server. This address is usually determined when installing the client.

Dynamic Assignment

If multiple identical server services are exist, static assignment of the clients to one server service is not ideal. In this case, dynamic assignment of the clients to the server services is better. For this dynamic assignment, criteria such as performance, load balancing, availability, and organizational assignment all play a part. A higher-level instance takes over the dynamic assignment of clients for a specific server service, or gives information about which network addresses can be used to access a specific server service.

Switchover Process

Switchover is a process where server applications are protected against failure by being forwarded to another standby host if one computer fails. For the switchover to remain transparent to the clients, the server must be able to be accessed by the clients in the same way after switching over to another host. To do this, a virtual IP address and a virtual IP name are assigned to the network adapter, in addition to

the normal IP address and IP name for each server application. If one host fails, this virtual IP address switches with the server application to the standby host. Therefore, the switchover of a server application from one host to another remains transparent for all clients that use the virtual IP name to access the server application.

The SAP System

In the SAP System, the message server knows the active SAP server services or instances, and the network addresses through which these SAP server services can be accessed. The system knows only one network address for each instance, the address that corresponds to the host name of the host where the SAP server service runs. This means that each client must access its server service through the host name. A client in this form is not only a SAPgui frontend, but also an SAP application server that requests a service from another SAP server. This accessibility over the host name is important if the SAP servers work with multiple network adapters and a separate server network.

For the SAP System to run optimally in a network, the following SAP guidelines for the network setup were formulated.

4.4.2 Network Setup for the SAP System

Network Setup Guidelines

When setting up your network configuration, follow these guidelines:

1. Each SAP application server must have at least one network adapter that all the other SAP application servers in the same SAP System, and all frontend computers and other communication partners can access. Configure your network to meet this requirement.

2. The IP name of this network adapter (according to point 1) must match the host name (which you can find out by using the command **`hostname`**). For high availability configurations with switchover, special rules apply (see *Switchover Environments*).

3. To configure the SAP System, the frontends and all communication interfaces only use the host name of the SAP application server. Do not use a different IP name, such as another network adapter, or an alias in a table or in a configuration file of the SAP System. This ensures that only that name is used whose accompanying IP address can be accessed by any communication partner.

4. If the SAP servers are interconnected by a separate, fast server network, static IP routing entries must ensure that the data traffic between the servers is sent through this network.

For network configurations that do not follow these guidelines, SAP Basis functions such as system monitoring, CPI-C, RFC communication and SAPlogon, will not function, even if the SAP System functions partially. In this case, SAP cannot guarantee the functioning of the SAP System. Further developments in future SAP Releases may lead to other problems if you did not follow the SAP guidelines

Procedure for Configuring Your Network

1. Connect the SAP application servers through the access network.

2. Set the host names for all SAP application servers to the IP name of the network adapter that is connected to the access network. To find the restrictions for choosing the host name and the mechanisms for assigning the host name (for example, Hosts file, DNS), see *Name and Address Resolution for R/3 Servers*.

3. Configure your computers according to the relevant guideline **before** installing the SAP System, since the host names from the installation procedure are used automatically for various parameters. SAP does not support a name change after installation.
4. Only this host name is used for the SAP servers when you configure your SAP System.
5. If the SAP servers are interconnected through a server network to distribute the network load, the data traffic between SAP servers must be routed over the server network using static IP routing entries.

Conclusions for Network Topology

To achieve a stable and, above all, easy-to-manage network, choose the most simple and clearly divided network topologies.

The following two sections contain recommendations on simplifying the following configurations:

- One network adapter for each SAP server
- Two network adapters for each SAP server

One Network Adapter for Each Server

Only one network adapter and only one IP address is used for each server for communication between the SAP components (communication between the SAP servers, and with the frontends).

In this case, the host name must be identical to the IP name that belongs to the IP address of the network adapter.

The network must allow for communication between all the SAP servers using these selected network adapters, either directly or using a router (see the first guideline for setting up your network).

In this case, no special routing entries are required.

Two Network Adapters for Each Server (Multihomed Host)

For larger SAP Systems, the data traffic must be divided into multiple network segments. The data traffic between the SAP database and application servers is a particularly heavy load. This server network must have as high a bandwidth possible, and must be separated from the remaining network traffic either physically or by using a router.

To do this, the SAP servers need a second network adapter that is only used for communication in the server network. This network adapter handles the data traffic between the SAP servers only; it does not handle communication to the SAPgui frontends, or data backups. The host name of the server must match the IP name of the network adapter used for communicating with the frontends.

The data traffic between the SAP servers (database server, central instance and dialog instances) is routed through the relevant static IP routes over the server network.

Two network adapters for each server are sufficient for most of your network communication needs.

You can use more than two network adapters for each server (for example, for data backup). However, you must follow the guidelines above for setting up your network.

4.5 Configuration Examples

4.5.1 Database (DB) With Central Instance (CI)

The following example shows the simplest variation of a distributed SAP System. The central system, consisting of a central instance and a database on one host, is the core of the SAP System. You can add as many additional application servers as you want. In these SAP Systems, the communication of the servers among themselves (DB with APP and APP with APP) makes up the majority of the network load. Therefore, a separate high speed network (the server network) is used for server communication.

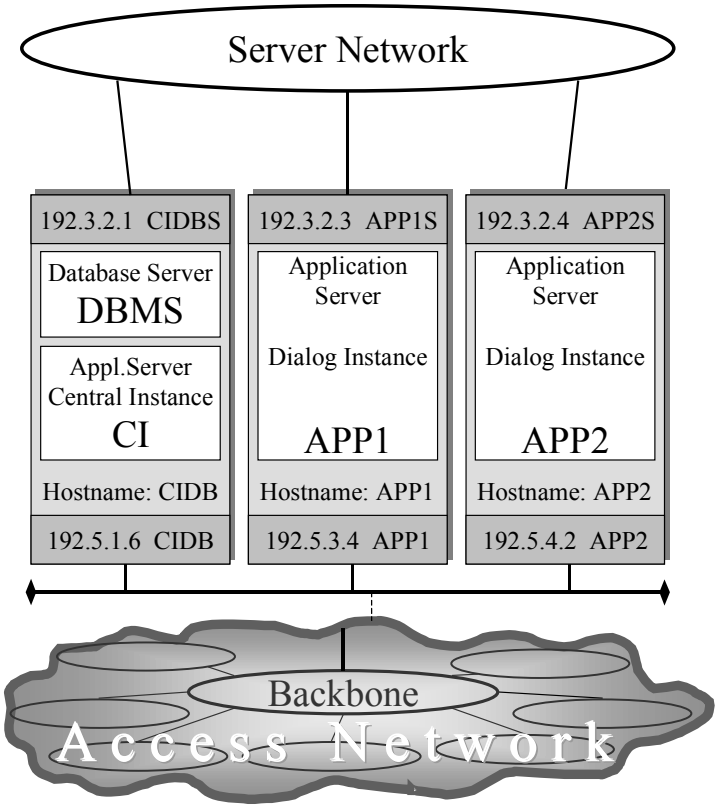

As described in guideline 3 for setting up your network, the host name is always used for accessing the SAP System. This applies to the connections of SAPgui to the central instance or to the application server, and also among the application servers, and for the application server to the database. To simplify the configuration, the host name must be identical to the name of the network adapter that the frontends can access over the access network. The servers also use this host name to address each other. To route network traffic between the servers over the server network, you have to use static IP routes on all the servers.

Example of routing entries (indirect host routes with UGH flags) on all servers:

CIDB		
Destination	Gateway	Flags
APP1	APP1S	UGH

APP1		
Destination	Gateway	Flags
CIDB	CIDBS	UGH

APP2		
Destination	Gateway	Flags
CIDB	CIDBS	UGH

4.5.2 Database (DB) Separate From the Central Instance (CI)

For large SAP Systems with high throughput requirements, you must distribute the central instance and the database over separate hosts. Here as well, you must ensure that the entire data traffic between the database and application servers passes over the server network. As you can see in the graphic, the server network is a separate subnet that only connects the SAP System servers with each other, and ensures that no external data traffic, and particularly no broadcasts, interfere with the data traffic between the SAP servers.

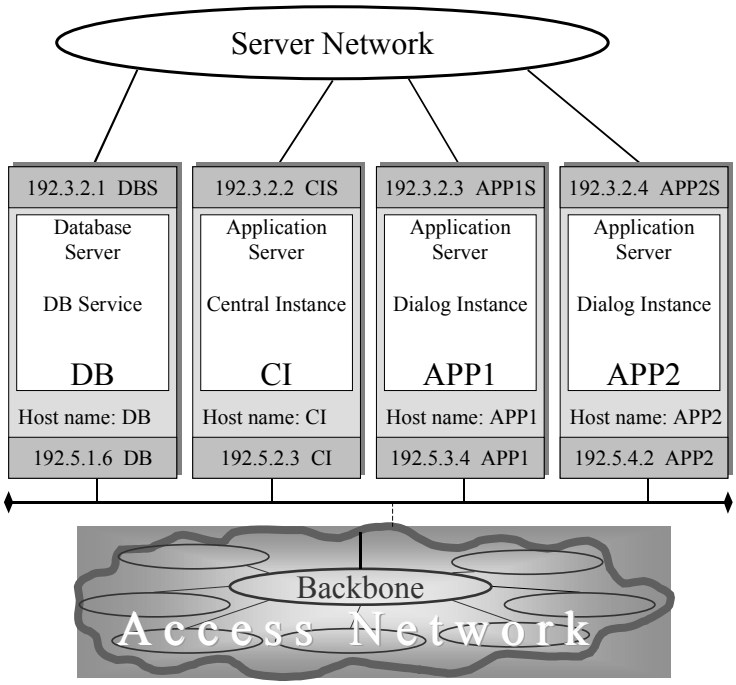

This configuration lets you meet guideline 2 most easily, because the host name of all SAP servers points to the network adapter that the frontends can access over the access network.

To route the data traffic among the servers over the server network when addressing using host names, the corresponding static IP routes must be set up on all servers.

Example of routing entries (indirect host routes with UGH flags) on all servers:

DB		
Destination	Gateway	Flags
CI	CIS	UGH
APP1	APP1S	UGH
APP2	APP2S	UGH

C1		
Destination	Gateway	Flags
DB	DBS	UGH
APP1	APP1S	UGH
APP2		UGH

APP1		
Destination	Gateway	Flags
DB	DBS	UGH
CI	CIS	UGH
APP2	APP2S	UGH

The same applies for additional application servers.

4.5.3 Isolated Database Server

The following example shows a database server that is only connected to the server network (isolated). No SAP applications run on this host (frontends have no access). The advantage of this configuration is the physical isolation of the database server. This server, on which all the data is located and which is essential for the SAP System to operate, cannot be accessed by each member in the network. This prevents unwanted access to the data, maintains performance and the ability of this important host to continue running in a simple and clear way. This increases the security and availability of the entire SAP System.

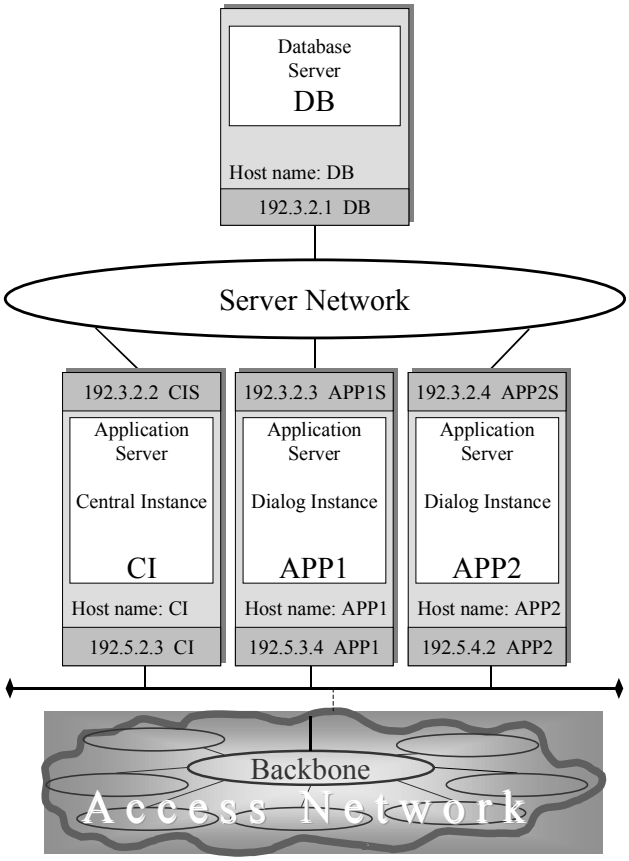

In this configuration, you can access the database from other SAP servers directly (without routing) through its host name. Since the database server only has one network adapter, this ensures that the communication with the database server is handled by the server network. To ensure that the data traffic between the SAP servers (central instance and application server) also occurs on the server network, you must set up the corresponding routes on these servers.

Example of routing entries (indirect host routes with flags UGH) on the database server:

DB		
Destination	Gateway	Flags
CI	CIS	UGH
APP1	APP1S	UGH
APP2	APP2S	UGH

Example of routing entries for a central instance and application server:

CI		
Destination	Gateway	Flags
APP1	APP1S	UGH
APP2	APP2S	UGH

APP1		
Destination	Gateway	Flags
CI	CIS	UGH
APP2	APP2S	UGH

APP2		
Destination	Gateway	Flags
CI	CIS	UGH
APP1	APP1S	UGH

4.5.4 Switchover Environments

To ensure high availability, various switchover products are offered for UNIX and NT platforms. There are two procedures for this:

- **Identity takeover**
  In the identity takeover procedure, a standby host completely takes over the identity and task of the failed host. In this case, no particular SAP configuration is necessary, since the standby host is identical to the failed host after the switchover.
- **Virtual IP address takeover**
  In the virtual IP address takeover, each network adapter also has a virtual IP address and a virtual IP name, in addition to the normal IP address and the accompanying IP name. During a switchover, another host takes over the failed virtual IP address and virtual IP name. For this virtual IP address takeover to function with SAP software, you may only use the virtual IP name in the configuration, which can be taken over by another host. The advantage of this procedure is that an empty standby host is not needed; instead the failed service (indicated by the virtual IP name) is taken over by another host in addition to its normal functions.

In the following example, the database server (DB) and the central instance (CI) are protected from each other. For example, if the central instance fails, the database server also takes over the functions of the central instance and also the virtual addresses VCI (access network) and VCIS (server network). For all the application servers and frontends involved, the central instance remains accessible through the names VCI or VCIS. However, the central instance does not run on the host with the host name DEF; instead it runs on the host with the host name ABC. Therefore, you must set the local host name that is valid for the SAP System by using the SAP profile parameter `SAPLOCALHOST`. You determine the name by which the database server can be accessed by using the parameter `SAPDBHOST`.

Do not use the following example as a standard solution for high-availability SAP Systems. It only serves as an example to explain the effects of an IP address takeover on the network configuration.

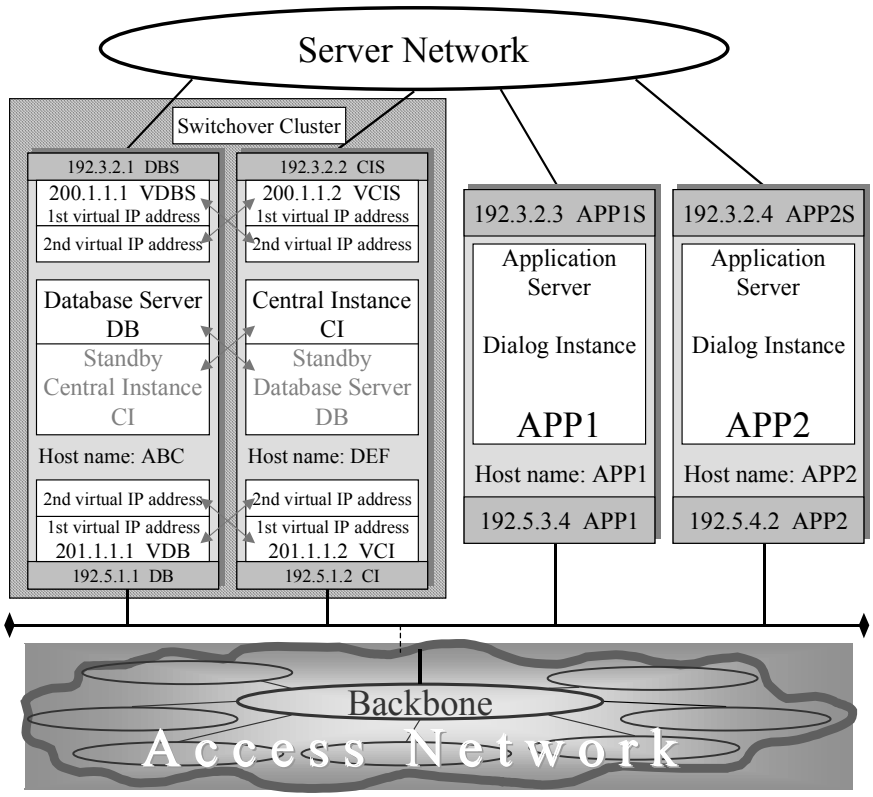

4.5.5 Frontend Configuration

In order for the frontends to access their dialog instances in the same way after the switchover (or, if using SAPlogon, the message server), you can use the virtual IP name only. This name must exist in the file `/etc/hosts` or in the DNS.

4.5.6 Server Configuration

In all profile parameters containing the host names of servers, you may only use the virtual IP name. Also, the instance names contain host names that may only be virtual IP names.

Example of the Most Important Parameters:

`DEFAULT.PFL`	`SAPSYSTEMNAME = C11` (SID for example, `C11` and the instance number `00`)
`DEFAULT.PFL`	`SAPDBHOST = VDB`
`DEFAULT.PFL`	`rdisp/mshost = VCI`
`DEFAULT.PFL`	`rdisp/vbname = VCI_C11_00`
`DEFAULT.PFL`	`rdisp/enqname = VCI_C11_00`
`DEFAULT.PFL`	`rdisp/btcname = VCI_C11_00`
`Instance CI`	`SAPLOCALHOST = VCI`
`Instance CI`	`SAPLOCALHOSTFULL` = fully qualified domain name from VCI (only required for NT)
`Instance APP1`	No special entries

Example of Static IP Routes:

This configuration lets you meet guideline 2 most easily, because the host name of all application servers are identical to the network adapter that the frontends can access over the access network.

To route the data traffic in the server over server network when addressing using host names, you must set up the corresponding static IP routes on all servers. Note that with the hosts belonging to the switchover cluster, you must use the virtual IP addresses or the virtual IP names.

Example of Routing Entries (Indirect Host Routes with UGH Flags) on All Servers:

DB		
Destination	Gateway	Flags
VCI	VCIS	UGH
APP1	APP1S	UGH
APP2	APP2S	UGH

C1		
Destination	Gateway	Flags
VDB	VDBS	UGH
APP1	APP1S	UGH
APP2	APP2S	UGH

APP1		
Destination	Gateway	Flags
VDB	VDBS	UGH
VCI	VCIS	UGH
APP2	APP2S	UGH

Same for APP2

5 Examples of SAP Network Configurations

5.1 Overview

The following network configuration examples show how you can set up your network for small, medium, and large SAP Systems. All these configurations serve as examples and are not compulsory. However, if you set up different networks, you must still follow the SAP guidelines for network configuration.

5.2 Small SAP Systems

For smaller SAP Systems, the SAP server uses only one network adapter for communicating with the SAPgui frontends. Since the network load is very low through SAPgui frontends, you do not need to connect the SAP server using several, physically separate subnets (network adapters). Dividing up the rest of your network into subnets is possible and also useful. High data traffic only occurs locally on the SAP server between the database and the central instance.

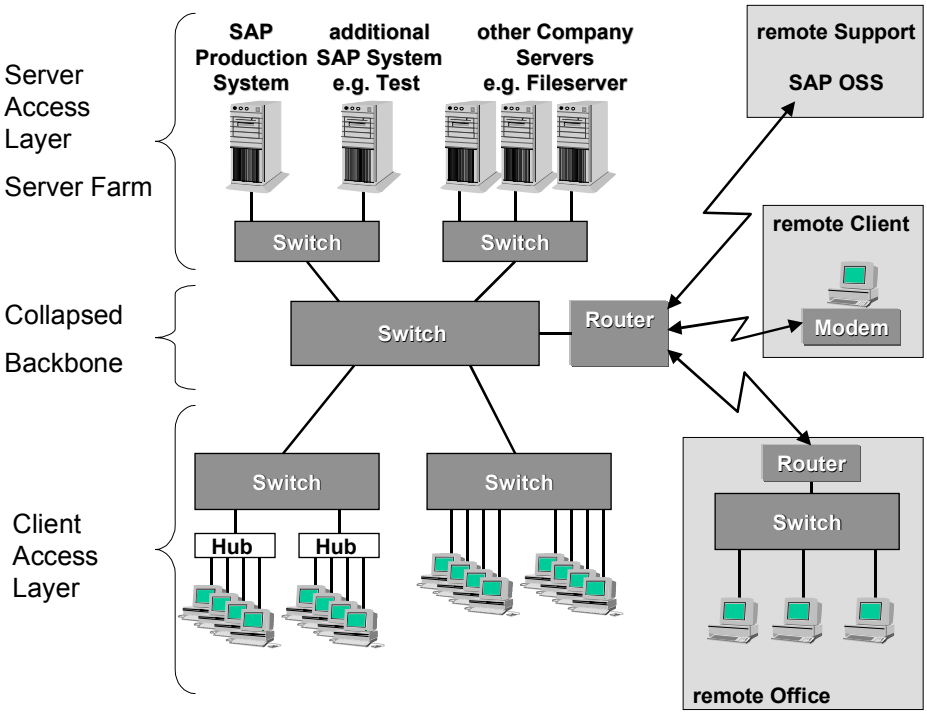

5.3 Midsize and Large SAP Systems

If your SAP System consists of several servers, a high network load occurs between the database server and the application servers. A high-speed switch (at least 100 MBit) is suitable for a simple and cost-effective connection of the servers.

You can use one or two network adapters for each server. The differences are explained below.

5.3.1 One Network Adapter for Each Server

The entire data traffic of a server runs over one network adapter only. Therefore, you need a network adapter with sufficient transfer speed (at least 100 MBit). If the servers connect using the switch, ensure that the entire bandwidth is available to all ports. The host name of the servers must match the IP name for the respective network adapter.

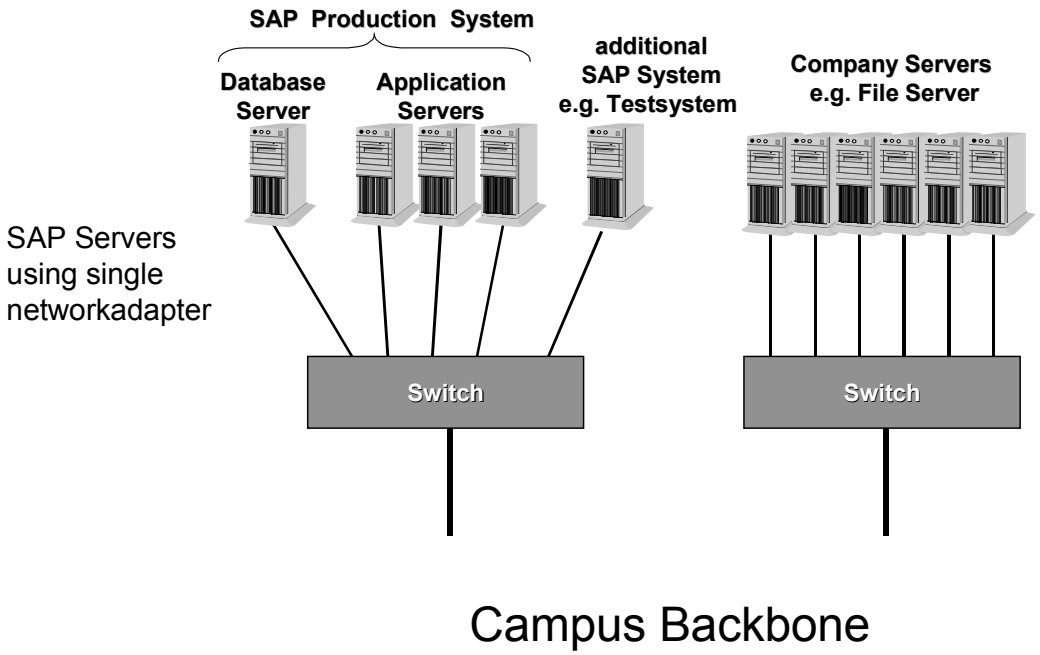

5.3.2 Two Network Adapters for Each Server (Separating Data Traffic)

In very large SAP Systems, the network adapter can become a bottleneck. In this case, you can install a second network adapter for each SAP server. It is only used for the communication between the SAP servers (database server, central instance and dialog instance) (server network).

This type of configuration also lets you separate the data traffic between the application servers and the database server from the rest of the data traffic. If you need to, you can install several network adapters on the database server.

Two possible configurations are described in the following section.

5.3.3 Connecting All SAP Servers to the Corporate Backbone

If an SAP instance also runs on the database server, you must connect this host to the backbone and the server network. When you assign the host name to an adapter, the same guidelines apply as with the application servers: The host name is assigned to the backbone connection.

In this configuration, routing entries must already exist on the application servers, ensuring that the database traffic flow is actually running over the server network.

When you are deciding if you want to install an additional SAP instance on the database server, consider performance and tuning aspects.

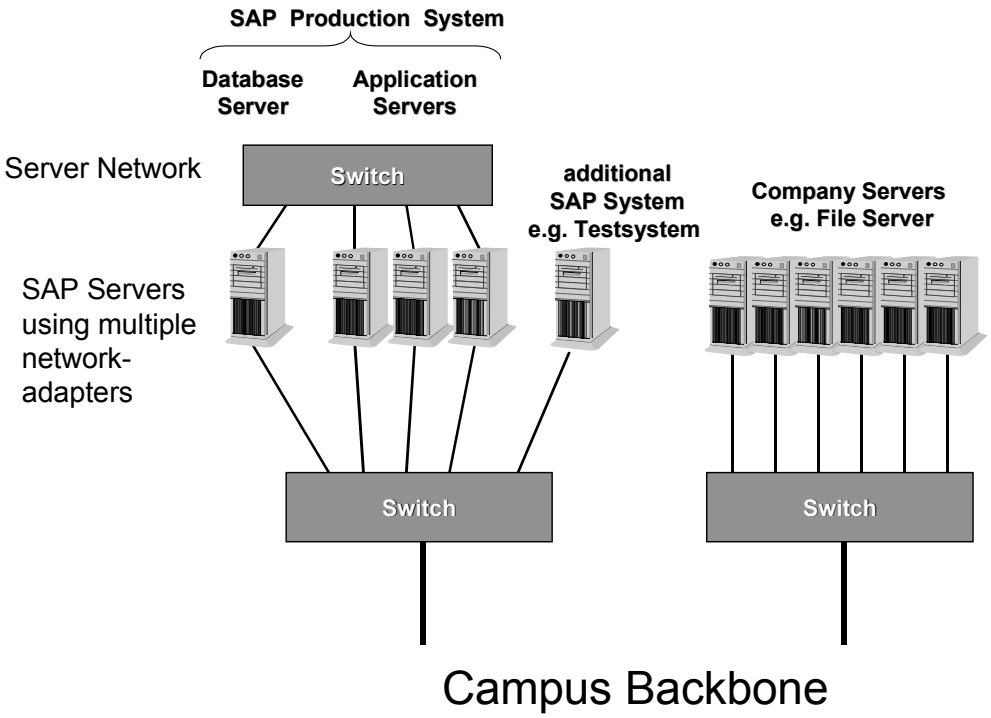

5.3.4 Connecting the Application Server to the Corporate Backbone

In this configuration, all the application servers are connected to the corporate backbone and to the server network. This complete isolation of the database server has advantages and disadvantages. Being separated from the rest of the corporate network guarantees high security against unwanted access to sensitive data, and against accesses that can hurt availability and performance on this important server. In this case, an SAP instance cannot run on the database server.

If a bottleneck occurs in the network adapter that connects the database server to the server network, you can establish an another connection to the server network with an additional network adapter.

For performance reasons, do not back up the data on the server, in particular the database server, over the server network. If you need to back up the data over the network, we recommend using a separate network for the backup (SAN = Storage Area Network). In this case, the database server must be equipped with an additional network adapter to connect to the SAN.

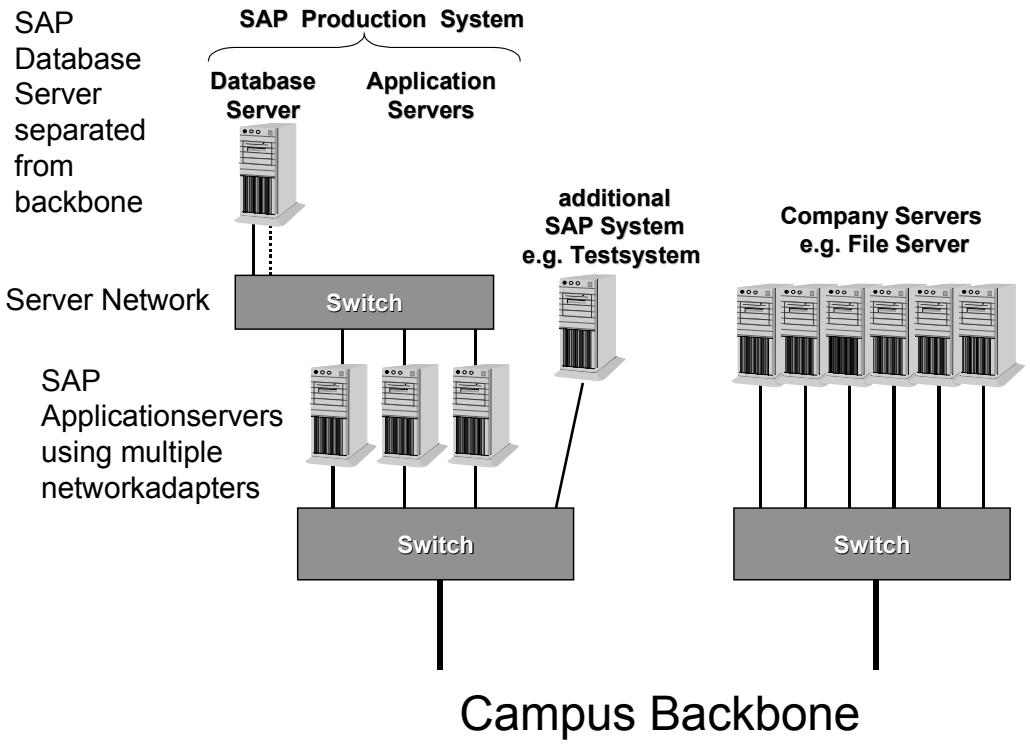
SAP
Database
Server
separated
from
backbone
SAP Production System
Database
Server
Application
Servers
additional
SAP System
e.g. Testsystem
Company Servers
e.g. File Server
Server Network
Switch
SAP
Applicationservers
using multiple
networkadapters
Switch
Switch
Campus Backbone

Appendix

Related Documentation

SAP Network and Infrastructure Design – SAP Internet Transaction Server (http://service.sap.com/network or mail to network@sap.com)

mySAP Workplace Network Integration Guide (http://service.sap.com/network or mail to network@sap.com)

SAP Network Integration Guide (http://service.sap.com/network or mail to network@sap.com)

ITS Administration Guide (http://help.sap.com/ → *SAP Library in HTML Format* → *Release 4.6C (English) Basis Components* → *Frontend Services* → *ITS*)

ITS Tuning Guide (http://www.saplabs.com/usa/index.htm?devarea/itsdoc.htm&2)

Platform and Technology Information Center (http://service.sap.com/platforms)

Smart Implementation (http://service.sap.com/smart)

R/3 Security Guide (http://service.sap.com/security → *Guidelines and Audits*)

Sizing (http://service.sap.com/sizing)

Sizing CRM (http://service.sap.com/sizing → *Media Center* → *Literature*)