Eliminate Passwords with Single Sign-on in your IBM i Environment

Thomas Barlen
Consulting IT Specialist
IBM Systems Lab Services & Training
http://www.ibm.com/systems/services/labservices/

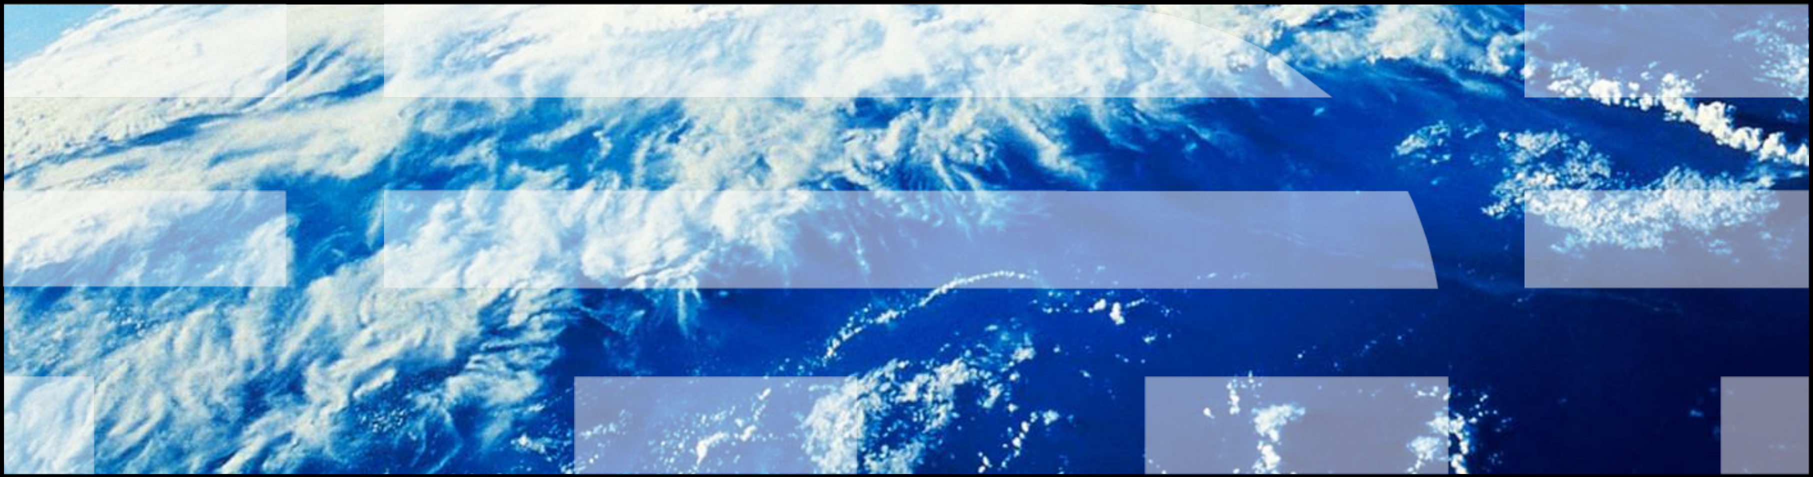

Notices

This information was developed for products and services offered in the U.S.A.

Note to U.S. Government Users Restricted Rights — Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM eServer™

AS/400®

i5/OS®

IBM®

iSeries

Redbooks (logo)™

IBM®

OS/400®

Redbooks

System i

System i5

IBM i®

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Acknowledgements

- This presentation was developed by Thomas Barlen, IBM Systems Lab Services & Training. Thomas is based in Germany, but works world-wide on mostly IBM i (i5/OS) and also AIX related security projects and presents at technical conferences.

- You can also engage Thomas for any kind of IBM i (i5/OS) related security (including but not limited to base security, object level access, object signing, IFS security, cryptography, security assessments, etc.) project or issue as well as cross-platform single signon projects. The best way to reach Thomas is by e-mail at barlen@de.ibm.com.

IBM Systems Lab Services & Training

IBM Systems Lab Services & Training can help you optimize the utilization of your data center and system solutions.

Lab Services has the knowledge and deep skills to support you through the entire information technology race. Focused on the delivery of new technologies and niche offerings, Lab Services collaborates with IBM Global Services and IBM Business Partners to provide complementary services that will help lead through the turns and curves to keep your business running at top speed.

The team consists of consultants and specialists from various geographies with extensive field experience and roots into the development labs.

What we offer

We apply the intellectual and technical capital of IBM development toward making sure IBM Systems products succeed, build loyalty and increase satisfaction. Our services include:

- Implementation, development and solution integration for IBM i™, AIX™, IBM System x™, and IBM System z™ technology.Storage services to simplify data migration to IBM System Storage™, enable and automate advanced hardware functions, implement virtualization technology for improving SAN-based storage and data management, and deploy copy services and disaster recovery solutions in an open systems environment.
- Application development, such as software components to accelerate solution deployment.
- IT optimization studies to help simplify, optimize, and reduce costs.
- Cross-platform virtualization, automation and systems management solutions.
- Data Center evaluation, consultation, and direction in power, packaging and cooling.
- Clusters support to serve a broad range of functions, from server/workload consolidation to high-performance parallel computing tasks.

Worldwide STG Lab Services Delivery Teams: a global team

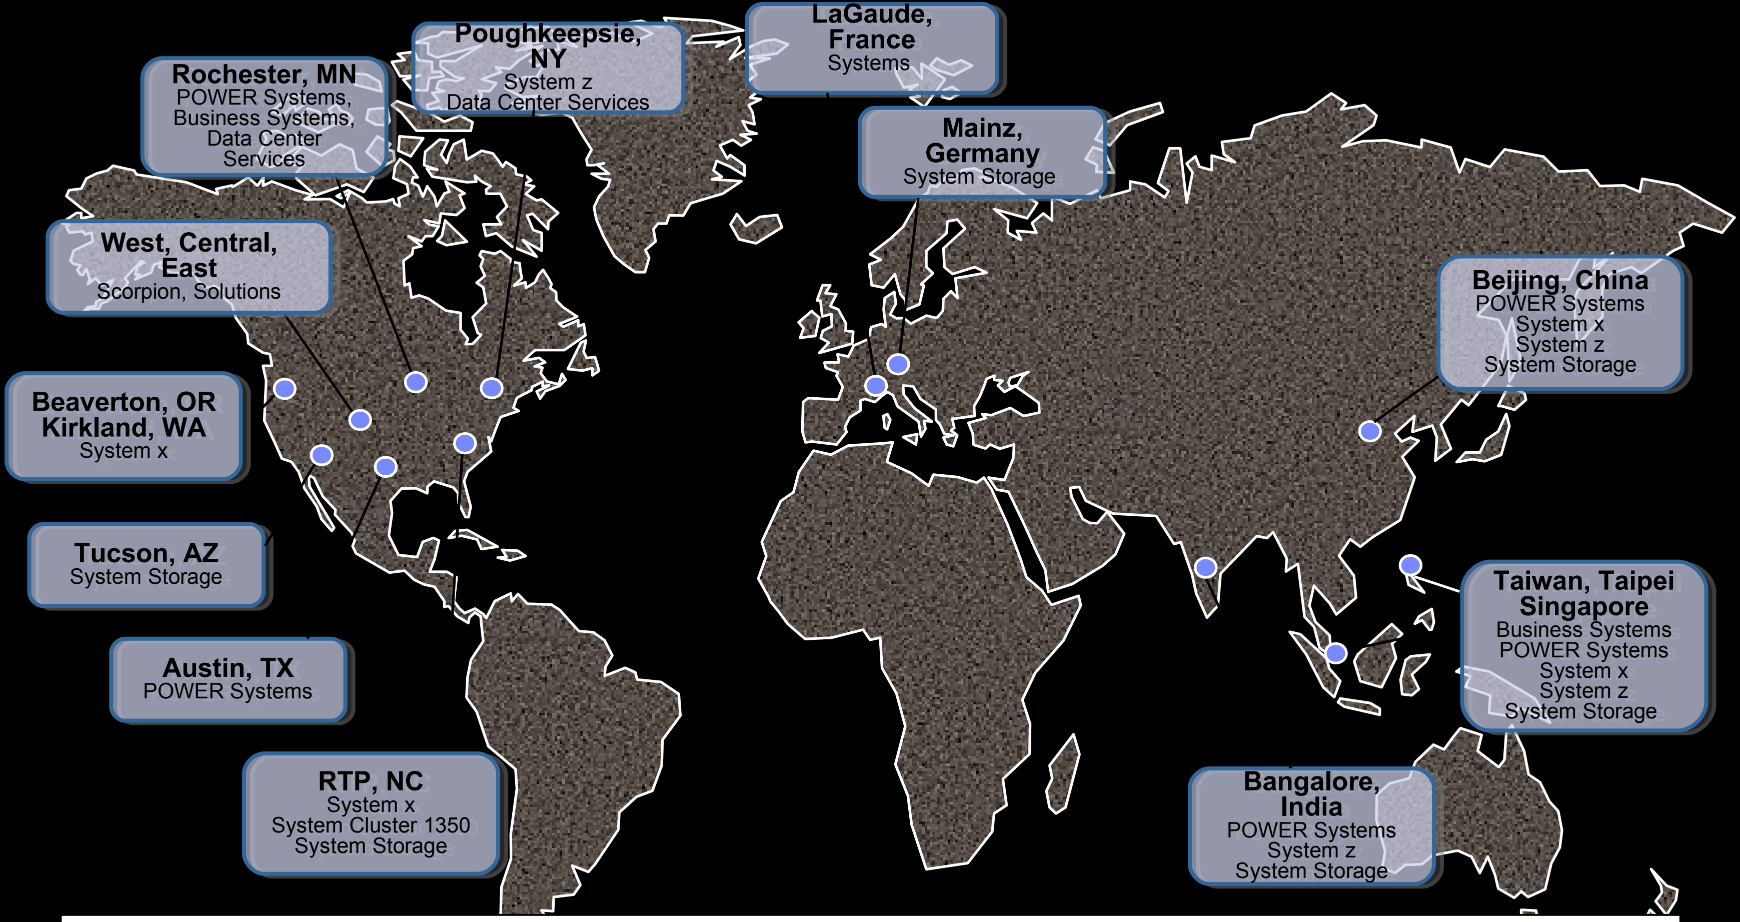

Our 1100 consultants reside in these centers as well as other cities around the world

IBM Systems Lab Services & Training Key Offerings ✓

Power Systems

- Copy Services
- HA / Business Continuity
- Performance
- **System and Application Security**
- **Security Assessments**
- **Security Education**
- **Single Sign-on (SSO)**
- **Cryptography**
- **Compliance**
- SOA
- AIX
- ✓ Power Care for High End POWER Systems
- ✓ Storage Migration for I5OS clients

Modular Systems

- System x Server & Storage Implementation & Skills Transfer
- System Management
- Virtualization, SCON, Migration
- High Availability
- Project Management (System Cluster 1350)
- ✓ iDataplex
- ✓ Blades Open Fabric Manager
- ✓ Virtualization Alternatives to VMware
- ✓ BCU Services

IT Optimization / Virtualization

- IT Systems Rationalization Study
- IT Systems Energy Rationalization Study
- ILM Assessments and Workshops
- Advanced Virtualization Rapid Assessment
- COBRA Rapid Assessment
- IT Systems Energy Efficiency Rapid Assessment
- Virtualization Rapid Assessment for UNIX
- Rapid Assessment for Linux on System z
- ✓ IT Current State Assessment
- ✓ Application IT Infrastructure Study
- ✓ Blue Cloud Offerings

Mainframe

- Linux on System z
- Java Performance Optimization
- System z SOA Services
- System z Currency and Migration Jumpstarts
- System z Performance Assessment for Parallel Sysplex middleware
- ✓ PCI Compliance Mitigation Services
- ✓ RACF Health Check
- ✓ DFSMS/DFSMShsm Health Assessment
- ✓ STP (Sysplex Timer Protocol) Planning and implementation

System Storage

- Technical Project Management
- Storage Consulting
- SAN Services
- SVC Services
- Advanced Copy Services
- ✓ Softek
- ✓ XIV
- ✓ N Series
- ✓ Novus
- ✓ Secure Data Erase
- ✓ Tape Encryption Key Management (TEKM)

Data Center Services

- Trends & Best Practices
- IT Systems Energy Efficiency Assessment
- Scorpion w/ Power
- Thermal Analysis
- Power Analysis
- Assessment of Cool Blue Rear Door Heat Exchanger (RDHX)
- ✓ MMT Thermal Offering
- ✓ IBM System Director Advanced Energy Manager Implementation Jumpstart
- ✓ EITA
- ✓ Data Center Planning Offerings

How to involve STG Lab Services

- Visit out Web site at http://www.ibm.com/systems/services/labservices/

United States [change]

Home | Solutions ▾ | Services ▾ | Products ▾ | Support & downloads ▾ | My IBM ▾ | Welcome Th

IBM Systems-related services

- IBM Systems Training
- IBM Systems events
- Executive Briefing Centers
- Design Centers
- High Availability Center of Competency
- Benchmark Centers
- Lab Services and Training
  - Systems Agenda Consulting
  - Systems Services
  - Systems Software Services
  - Cross Platform Consulting Services
  - Library
- IBM Solution Central Services

IBM Systems > IBM Systems-related services >

IBM Systems Lab Services and Training

Systems and Technology Group

Why IBM Systems Lab Services and Training

IBM has a vision for how we can improve the way the world works—how every pers business, organization, government, natural system and man-made system interact Each of these relationships represents a chance to do something better, more effici productively.

Agenda

- An introduction to single sign-on
- Possible solutions
- Technologies that enable Single Sign-on with password elimination in IBM i
- Implementation overview

Notes

Typical Environment Today

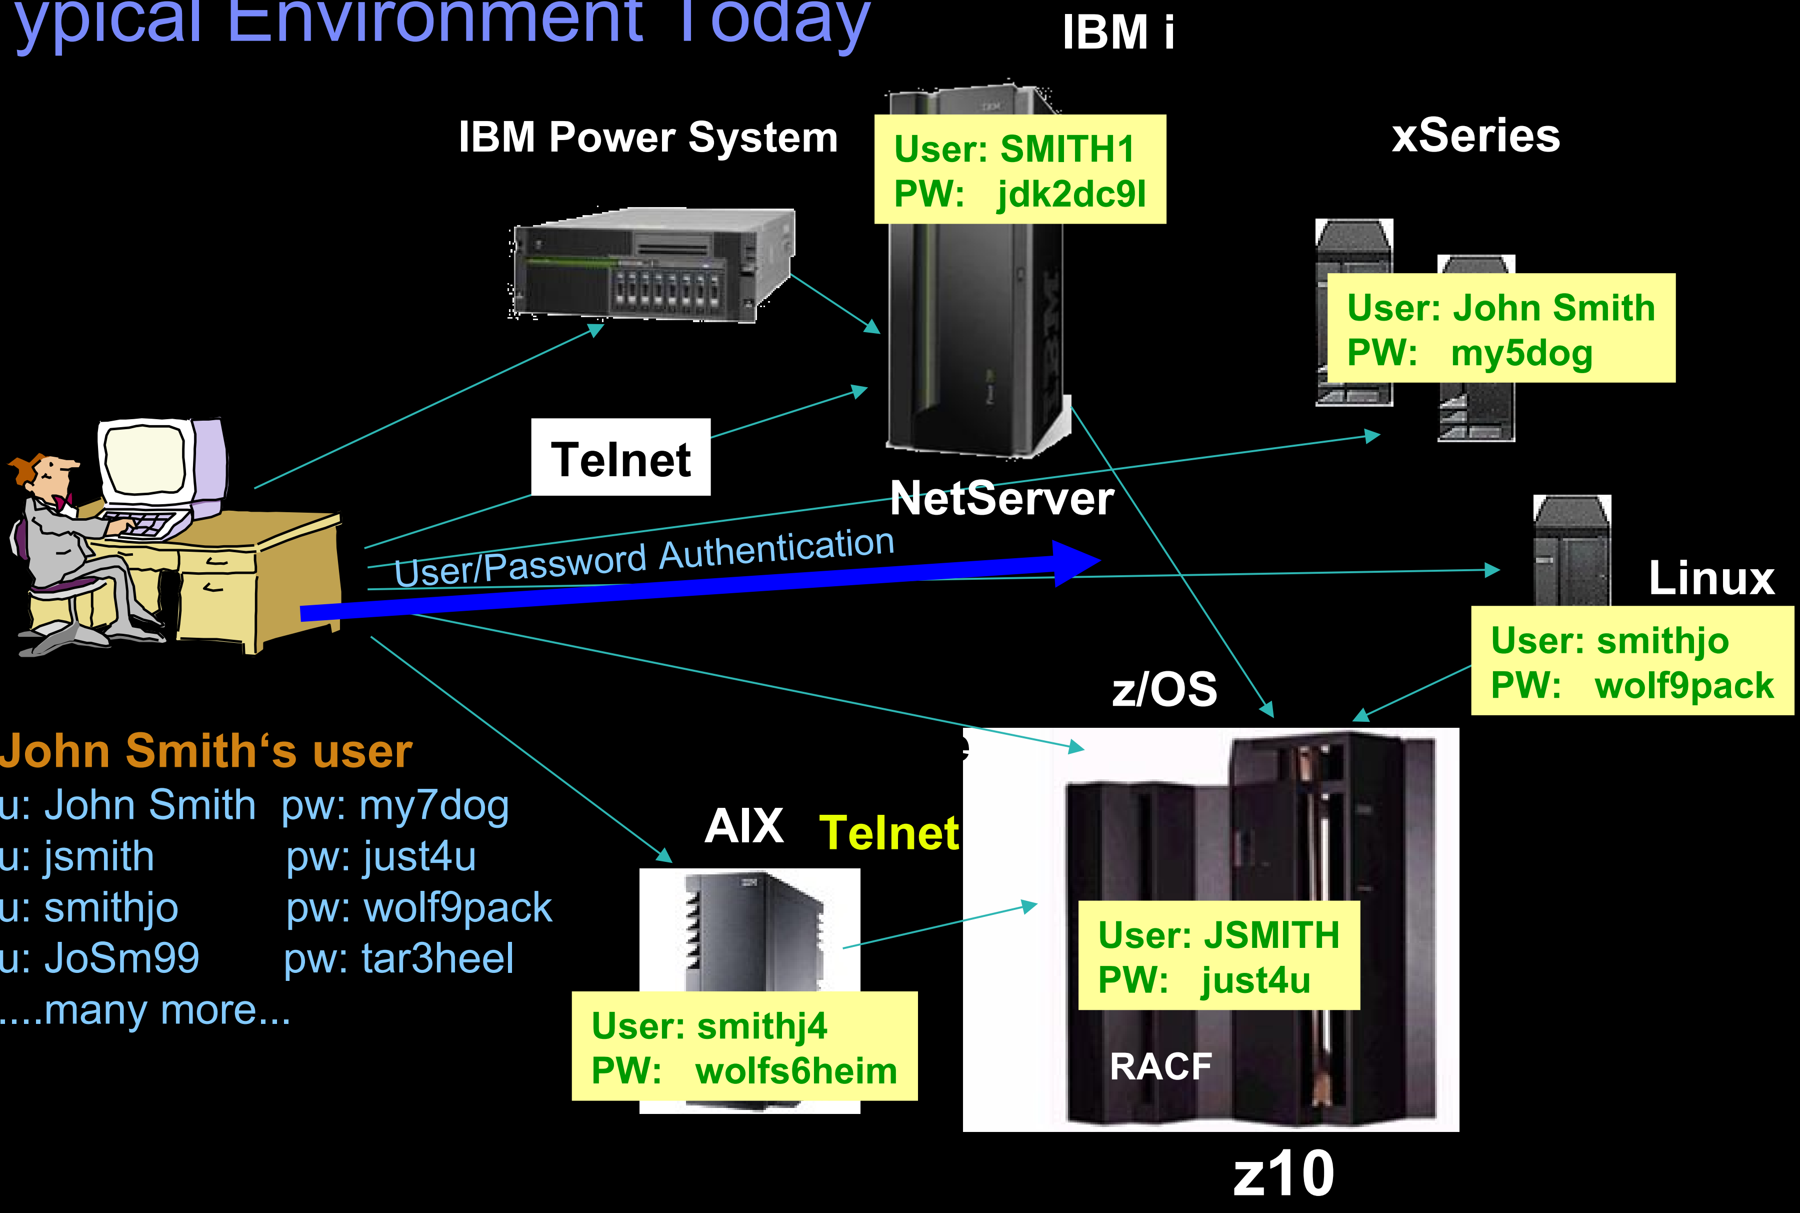

IBM i

User: SMITH1
PW: jdk2dc9l

IBM Power System

xSeries

User: John Smith
PW: my5dog

Telnet

NetServer

User/Password Authentication

Linux

User: smithjo
PW: wolf9pack

z/OS

AIX

Telnet

User: smithj4
PW: wolfs6heim

User: JSMITH
PW: just4u

RACF

z10

John Smith's user

u: John Smith pw: my7dog
u: jsmith pw: just4u
u: smithjo pw: wolf9pack
u: JoSm99 pw: tar3heel
....many more...

Notes: Typical Environment Today

- Each system has its own unique user registry, and most likely, its own rules for user IDs and passwords. Users end up with multiple user IDs and passwords. It is quite common that users try to simplify their own local environment by using the same password in multiple systems.
- As an application developer, you know that the customer data is spread out across many different types of systems. All of them having their own user registries and associated security semantics. Your only chance of providing a distributed application that works is to provide a new user registry for your application, despite the impact it will cause on administration.
- Cross platform distributed applications that span platforms often "agree" who a specific user is. When OS protected resources are accessed, the application projects (maps) the application view of a user into the OS view of the user. The back-end system is forced to trust the front-end servers.
- Passwords are often transmitted in the clear.

Agenda

- An introduction to single sign-on
- Possible solutions
- Technologies that enable Single Sign-on with password elimination in IBM i
- Implementation overview

Single Signon characteristics

- Sign on once to the network using, for example user ID and password
- Subsequent connection requests to application services and resources are authenticated without prompting for the user ID or password
- Taking different identities for various applications for a single entity into consideration is desirable

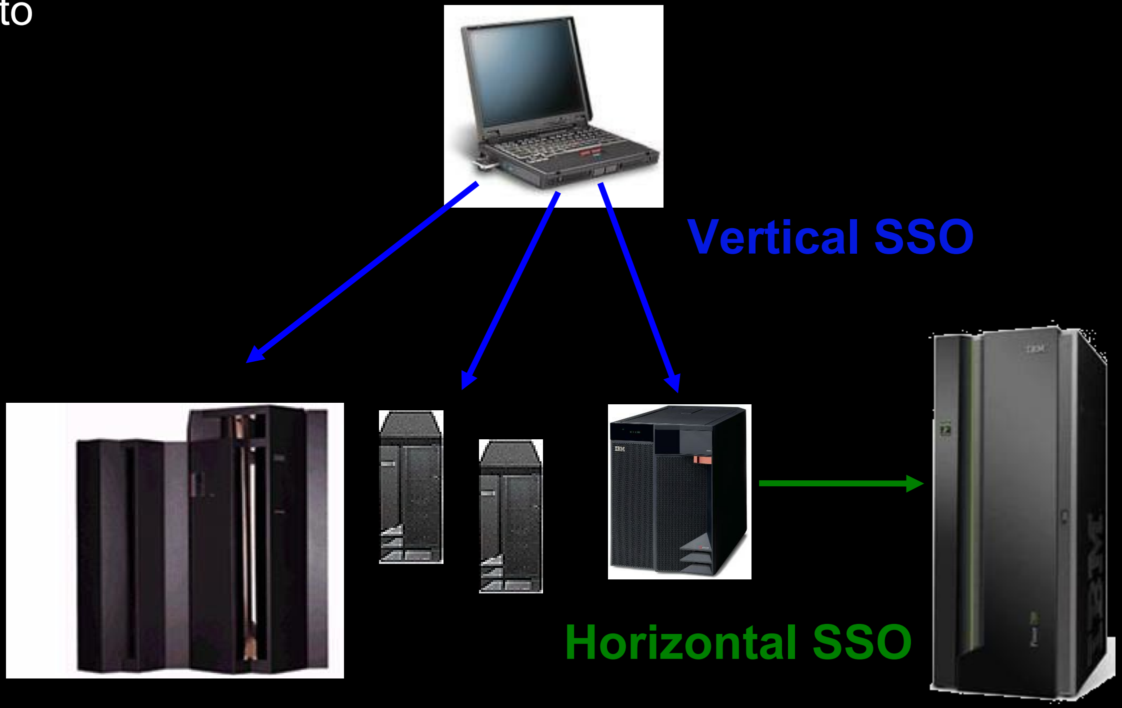

Notes: Single Signon characteristics

- The term single signon is often misinterpreted or confused with having a single user ID and password to sign on to a system. However, in most cases, users still have to sign on to each application or service individually. With a true SSO solution, a user signs on only once to the network (a central authentication service) and then accesses all participating services without re-entering a user ID or password. Many available SSO solutions, however, only offer SSO in a Web environment. It is desirable to have a SSO solution that works for both browser-accessible applications and local applications, such as Telnet or DB access.

- With SSO, we distinguish between horizontal and vertical SSO approaches:
  - Vertical SSO describes an approach where a client signs on from the client to each individual server using SSO.
  - Horizontal SSO involves a client signing on, for example, to a server application, which in turn connects to another server to access a database, signing on on behalf of the user (also with SSO).

Single Signon Solution using User / Password Authentication

- Pros
  - relatively simple to implement
  - covers basically every application signon that requires user and password
- Cons
  - users and passwords are stored centrally or decentralized
    - passwords are decryptable !!!
    - does not eliminate the need to manage passwords on all platforms
    - no multi-tier support, bad for audit trail

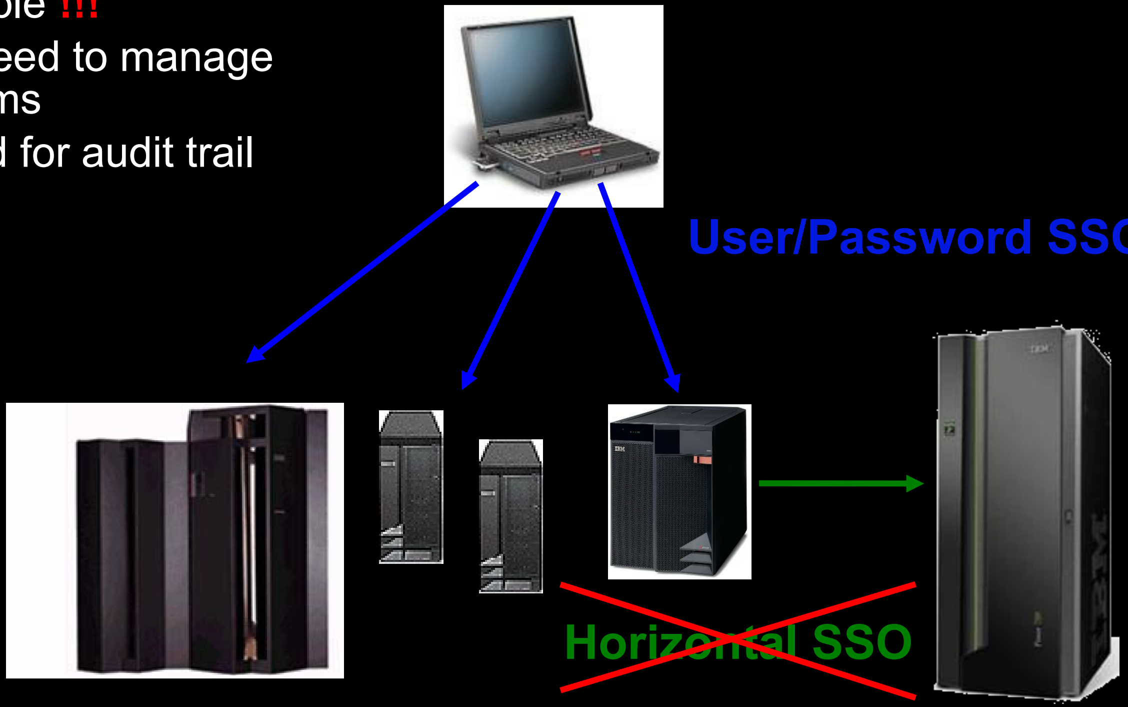

Single Signon Solution using Network Authentication

- Kerberos is an example of a widely used network authentication protocol
- Pros
  - eliminates the need to manage passwords on application systems
  - does not rely on passwords for authentication, it is ticket based
  - no passwords are stored in decryptable form
- Cons
  - requires support for every application (client and server, i.e. Telnet client and Telnet server)

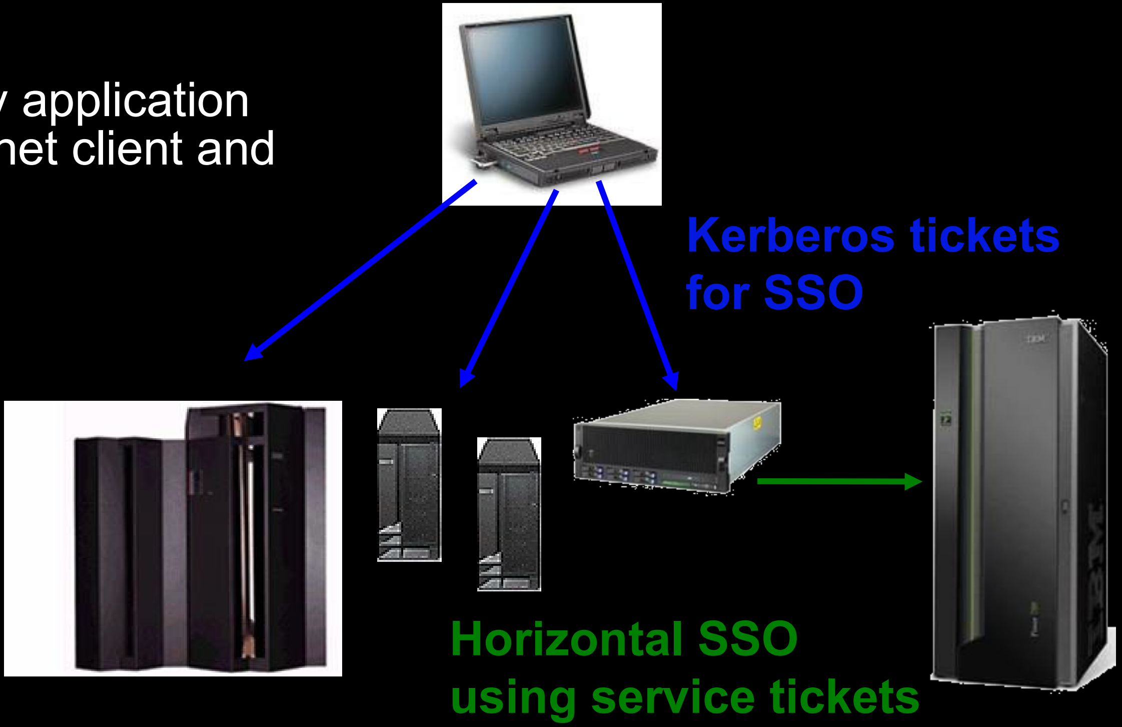

Kerberos tickets for SSO

Horizontal SSO using service tickets

Reducing the Costs of Managing a UserID

- Reducing the actual number of userIDs defined across the enterprise is not feasible today
  - Would require operating systems and applications to be re-written to exploit a centralized access control mechanism
    - The cost of doing this today is prohibitive and would eliminate much of the value add of OSes (and potentially applications)
  - Today’s implementations of centralized authentication mechanisms are advisory only. They do not enforce policy, they only define it.
    - Very useful for managing access control to “virtual resources”
  - Native access control mechanisms are already configured to protects many terabytes of data
    - Changing to a different access control mechanism would be prohibitive

Agenda

- An introduction to single sign-on
- Possible solutions
- Technologies that enable Single Sign-on with password elimination in IBM i
- Implementation overview

Enterprise Identity Mapping (EIM)

- Enterprise Identity Mapping (EIM) is a mechanism for mapping (associating) a person or entity to the appropriate user identities in various registries throughout
the enterprise
- EIM provides an infrastructure that lowers the expense for application developers to provide single signon solutions

EIM defined: Identity associations across user registries associated with OS platforms, applications, and middleware.

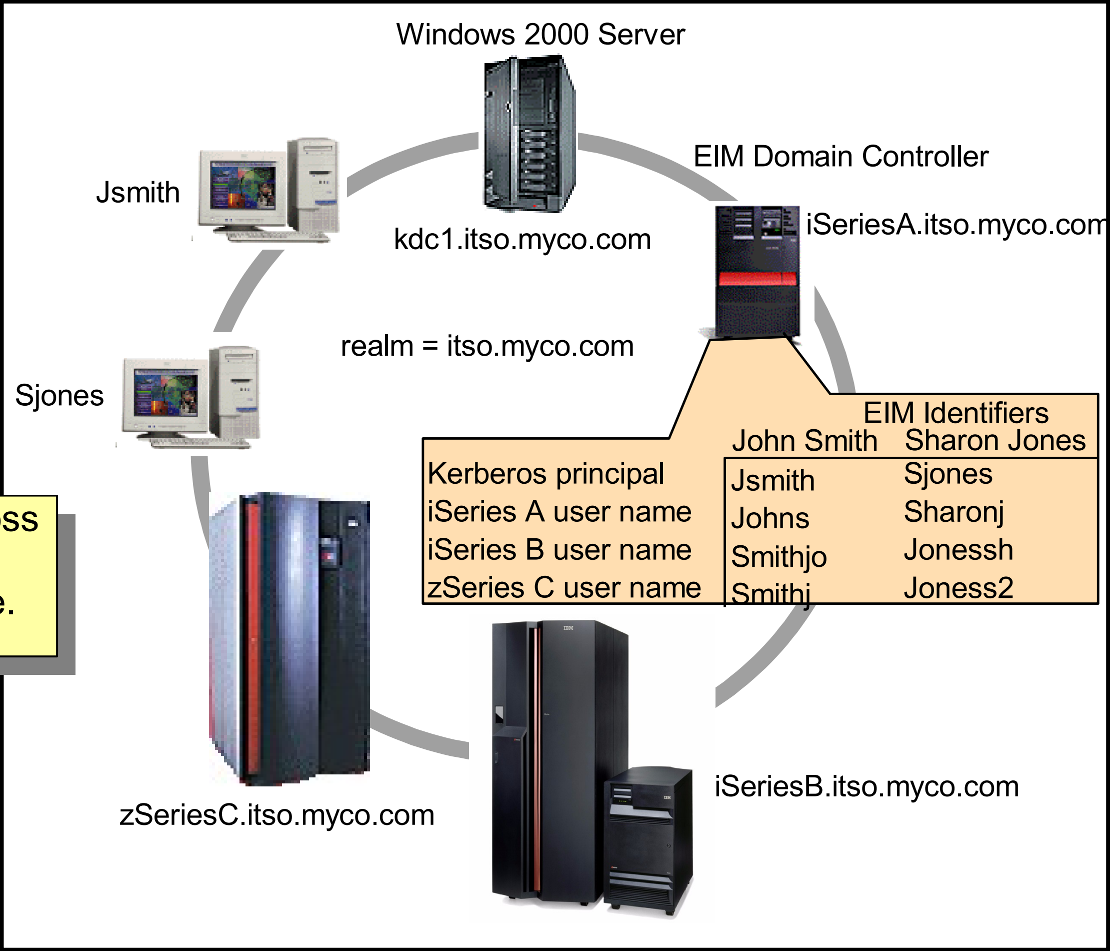

EIM Identifiers

	John Smith	Sharon Jones
Kerberos principal	Jsmith	Sjones
iSeries A user name	Johns	Sharonj
iSeries B user name	Smithjo	Jonessh
zSeries C user name	Smithj	Joness2

Notes: Enterprise Identity Mapping (EIM)

- EIM provides an infrastructure that lowers the expense for application developers to provide SSO solutions. IBM i (i5/OS) exploitation of EIM and Kerberos, along with exploitation by other IBM platforms and IBM software, provides SSO capabilities. This, in turn, provides users, administrators, and application developers the benefits of easier password and user identity management across multiple platforms—without changing the underlying security schema.

- EIM allows for operating system programmers and independent software vendors (ISVs) to independently implement support for a SSO environment without waiting for support from a specific product vendor.

- EIM is part of the IBM autonomic computing initiative that has a goal to give businesses the ability to manage systems and technology infrastructures that are hundreds of times more complex than those in existence today.

- The initiative represents the next stage of development under new tools. Self-managing servers are the ultimate in new tools for our customers. They're self-optimizing, self-configuring, self-healing, and self-protecting.

Weitere Technologien - LDAP

- LDAP Directory Server
  - EIM Domain Informationen werden im LDAP Verzeichnis abgelegt
  - IBM Directory Server ist Bestandteil von IBM i
- EIM Domain Daten
  - EIM Identifier
  - EIM Associations (Mappings)
  - Registries

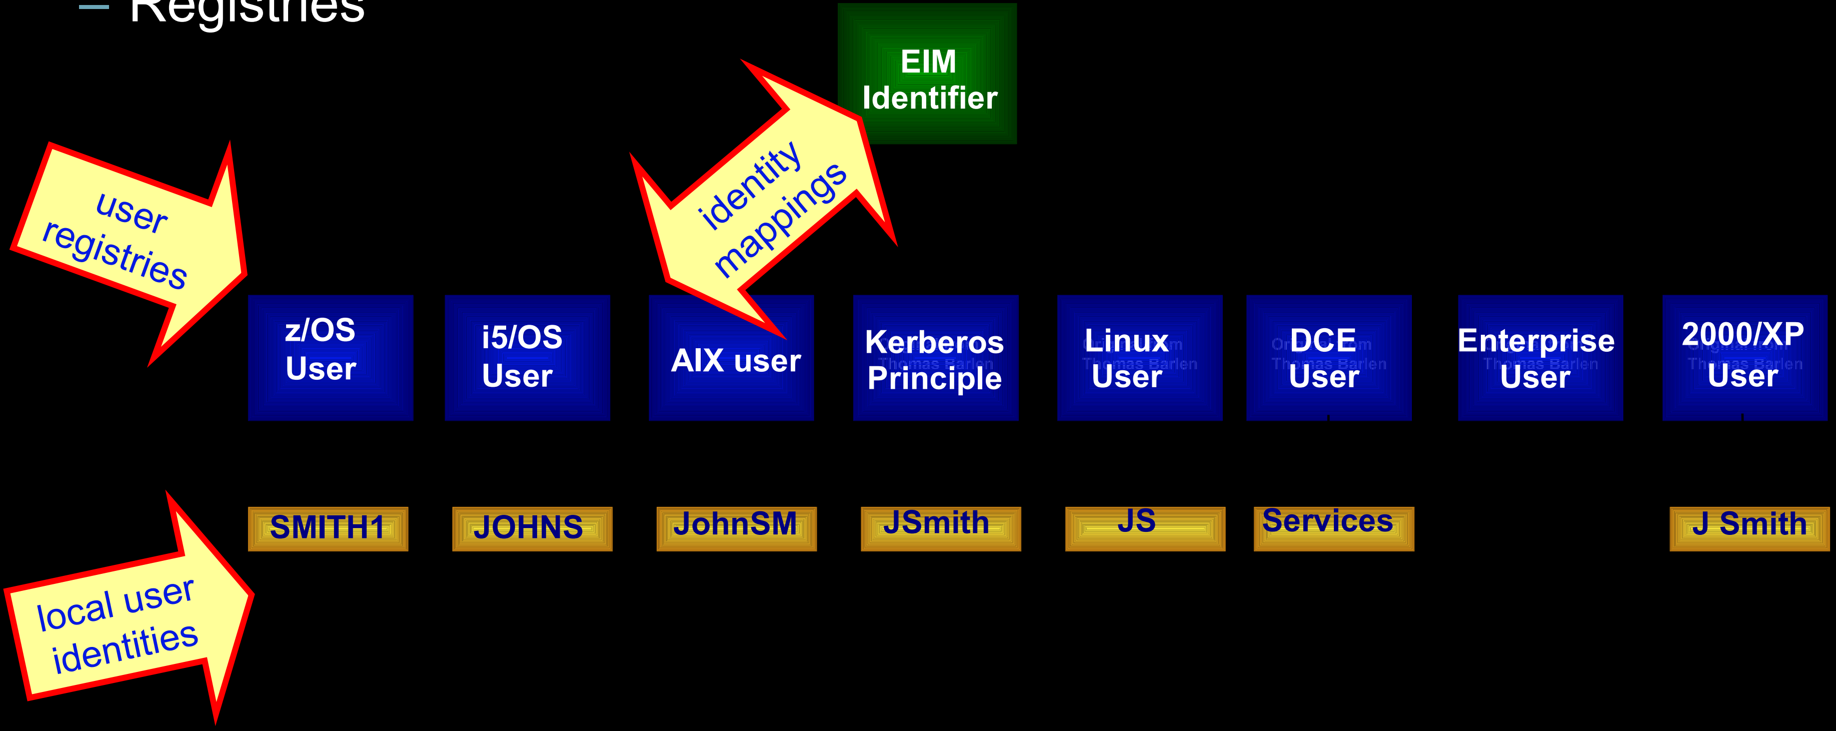

Notes: EIM Identifier

- An EIM identifier represents an actual person or entity in EIM. User identities for that person or entity can be associated with the EIM identifier. These identity mappings help to simplify the administrative task of keeping track of all of the user IDs that this person or entity may have within the enterprise.

- EIM identifiers can have a description, which can further define the person or entity it represents. You can also create aliases for the EIM identifiers, which can aid in locating a specific EIM identifier when performing a mapping lookup operation.

- Quite often different individuals within an enterprise share the same name. EIM identifier names must be unique within the EIM domain and can be confusing as to which individual the identifier belongs. Aliases allow the EIM administrator to have arbitrary and unique EIM identifier names, and to provide additional information about the individual to which the EIM identifier belongs. This information can also be used in a mapping lookup operation.

- For example, the EIM identifiers for two people named John S. Smith might be John S. Smith1 and John S. Smith2. The alias for John S. Smith1 could be John Samuel Smith and the alias for John S. Smith2 could be John Steven Smith.

- Each EIM identifier can have multiple aliases that can be used to identify which John S. Smith the EIM identifier is representing. Another alias might be added to each of the EIM identifiers for the two individuals that contains their department numbers.

Notes: EIM Domain Data

- Enterprise Identity Mapping (EIM) requires that the Directory Services (LDAP) server is configured with at least a basic configuration.
  If one does not exist, the EIM wizard configures one for you. From an EIM management point of view, you do not need to access the directory directly.

- But if you plan to use the directory for other functions, such as storing employee information, or configuring advanced functions, such as replication or SSL, you should first become familiar with the LDAP directory server. See "Plan your LDAP directory server" in the IBM i Information Center for planning information before you attempt to configure LDAP.

- Another excellent resource for iSeries Directory Services implementation and use is the IBM Redbook Implementation and Practical Use of LDAP on the IBM eServer iSeries Server, SG24-6193.

- The directory server is the container for the EIM domain and domain controller information, authorities, as well as access control to the information contained in EIM.

- For a production environment, we recommend that you configure the Directory Server to use SSL.

- **Do not** attempt to alter the EIM information without using the EIM APIs.

- The name space in a directory information tree (DIT) requires thorough planning. When setting up EIM using the Enterprise Identity Mapping setup wizard, you have the option to publish the EIM domain directly under the root (top) of the DIT or select an existing RDN within the DIT to publish the EIM domain underneath. If you configure an EIM domain on a new directory server and there is no intention of using the server for other purposes, you may want to publish the EIM domain under the root. Even though, it is recommended to set up a different DN, for example an entry of objectclass organization, under which the EIM domain will be published.

Authentication - Kerberos

- Kerberos is a network authentication protocol
- Designed to establish secure authentication from client to server (and vice versa) on an untrusted network
- NAS is built on the Kerberos Network Authentication Service (RFC1510)
  - Kerberos V5 is required
  - In IBM i, Kerberos is referred to as Network Authentication Servíce (NAS)
- Network Authentication Service (NAS) enables the operating system and applications to use Kerberos tickets for authentication instead of a user ID and password
- Applications can identify users and securely pass on the identity to other services
- Widespread throughout the industry, allows for interoperability between platforms
- Simplifies trust management

Notes: Kerberos

- The Kerberos system was designed and developed in the 1980s by the Massachusetts Institute of Technology (MIT), as part of the Athena project. The current version of Kerberos is Version 5, which is standardized in RFC 1510, The Kerberos Network Authentication Service (V5). For more details, see http://www.ietf.org/rfc/rfc1510.txt

- "Kerberos is freely available from MIT, under copyright permissions very similar to those used for the BSD operating system and the X Window System. MIT provides Kerberos in source form so that anyone who wants to use it may look over the code for themselves and assure themselves that the code is trustworthy. In addition, for those who prefer to rely on a professionally supported product, Kerberos is available as a product from many different vendors.

- In summary, Kerberos is a solution to your network security problems. It provides the tools of authentication and strong cryptography over the network to help you secure your information systems across your entire enterprise. We hope you find Kerberos as useful as it has been to us. At MIT, Kerberos has been invaluable to our Information/Technology architecture." Source: MIT

- Kerberos authentication itself does not automatically imply that the rest of the session is encrypted. However, Kerberos enables a secure exchange of encryption keys that could be used by a client program for session encryption. IBM i Access for Windows, for example, does not implement the encryption part of Kerberos. However, IBM i Access for Windows traffic can be encrypted by SSL instead.

Kerberos Environment

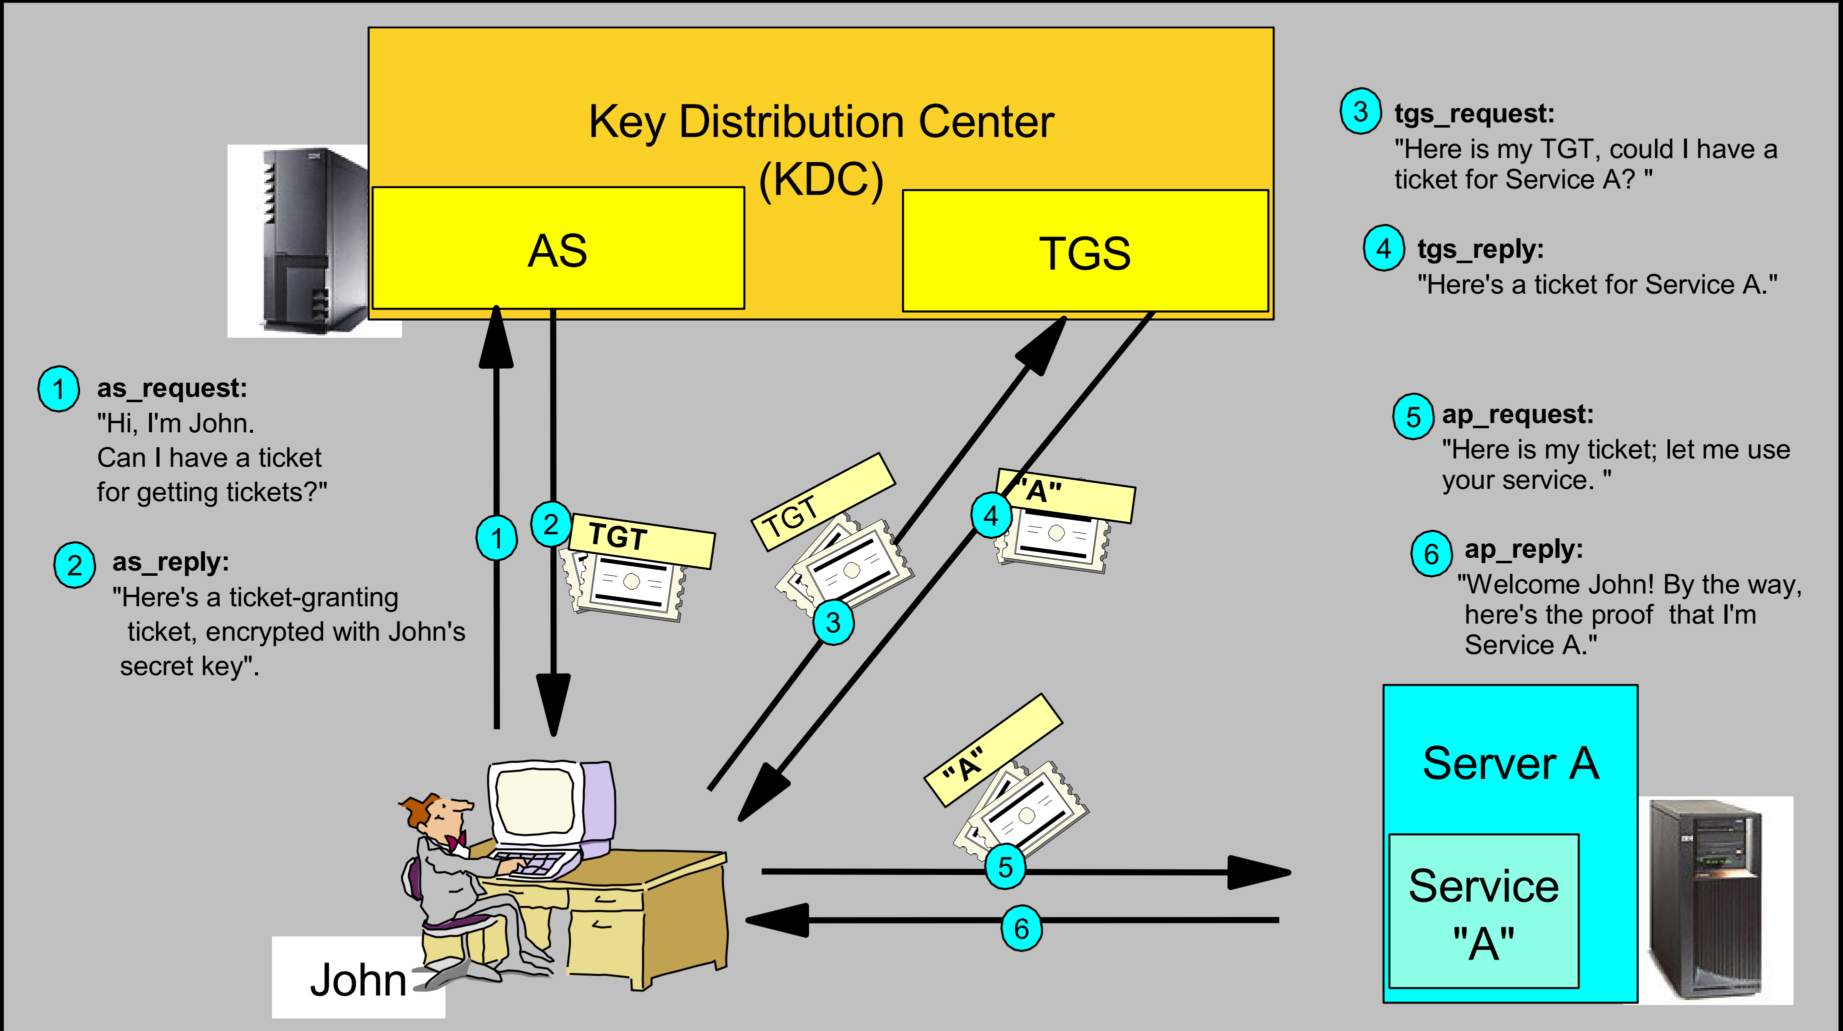

1 **as_request:**
"Hi, I'm John.
Can I have a ticket
for getting tickets?"

2 **as_reply:**
"Here's a ticket-granting ticket, encrypted with John's secret key".

3 **tgs_request:**
"Here is my TGT, could I have a ticket for Service A? "

4 **tgs_reply:**
"Here's a ticket for Service A."

5 **ap_request:**
"Here is my ticket; let me use your service. "

6 **ap_reply:**
"Welcome John! By the way, here's the proof that I'm Service A."

AS = Authentication Service, TGS = Ticket Granting Service

Notes: Kerberos Environment

- The Kerberos protocol consists of several sub-protocols (or exchanges). There are two methods by which a client can ask a Kerberos server for credentials. In the first approach, the client sends a clear text request for a ticket for the desired server to the Authentication Service (AS). The reply is sent encrypted in the client's secret key. Usually this request is for a ticket-granting ticket (TGT) that can later be used with the ticket-granting server (TGS). In the second method, the client sends a request to the TGS. The client sends the TGT to the TGS in the same manner as if it were contacting any other application server which requires Kerberos credentials. The reply is encrypted with the session key from the TGT.

- The client and server do not initially share an encryption key. Whenever a client authenticates itself to a new verifier it relies on the authentication server to generate a new encryption key and distribute it securely to both parties. This new encryption key is called a session key and the Kerberos ticket is used to distribute it to the verifier.

- The Kerberos ticket is a certificate issued by an authentication server, encrypted using the server key. Among other information, the ticket contains the random session key that will be used for authentication of the principal to the verifier, the name of the principal to whom the session key was issued, and an expiration time after which the session key is no longer valid. The ticket is not sent directly to the verifier, but is instead sent to the client who forwards it to the verifier as part of the application request. Because the ticket is encrypted with the server key, known only by the authentication server and intended verifier, it is not possible for the client to modify the ticket without detection.

- A Key Distribution Center (KDC) is a network service that provides tickets and temporary session keys. The KDC maintains a database of principals (users and services) and their associated secret keys. It is composed of the Authentication Server (AS) and the Ticket Granting Server (TGS). It is important that you use a secure machine to act as your KDC. If someone gained access to the KDC, your entire realm could be compromised.

Example Session

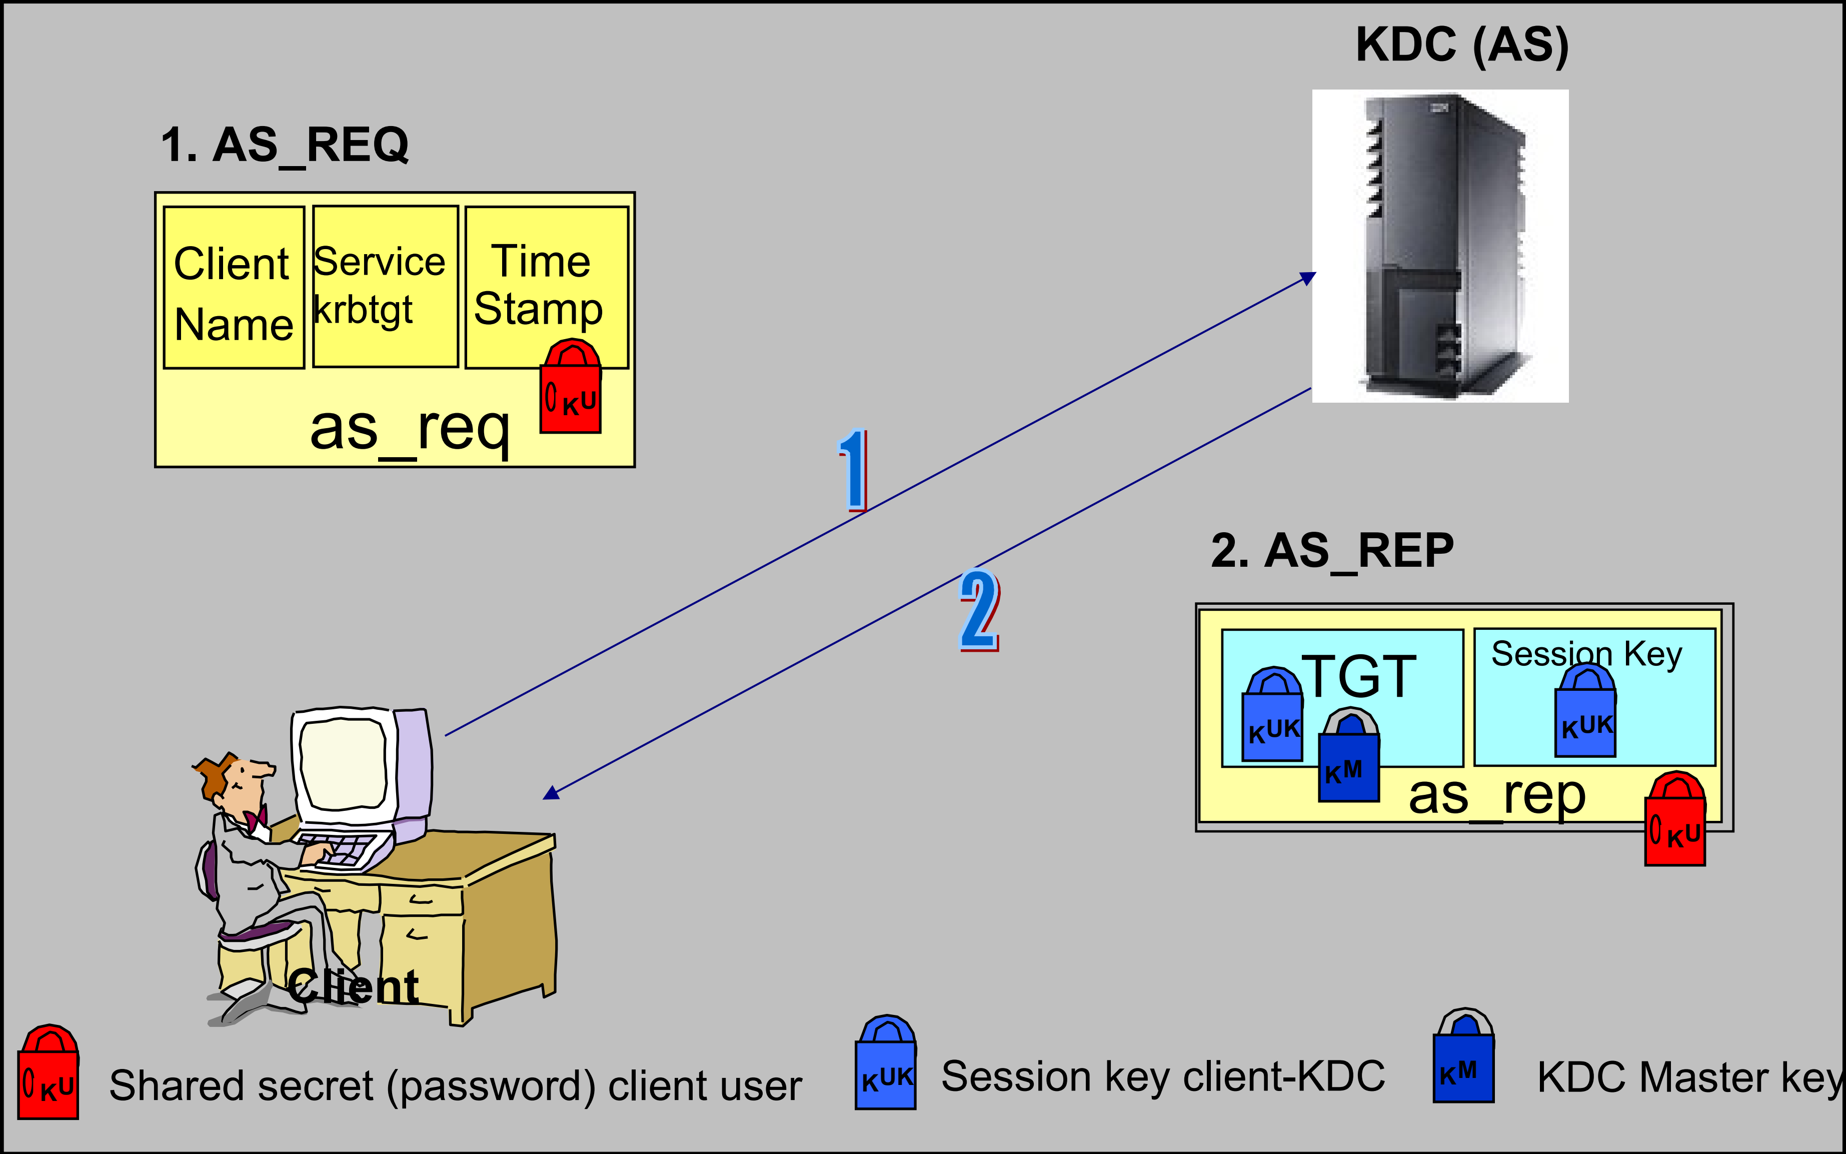

KDC (AS)

1. AS_REQ

Client Name

Service krbtgt

Time Stamp

as_req

K_U

1

2

2. AS_REP

TGT

K_{UK}

K_M

Session Key

K_{UK}

as_rep

K_U

Client

K_U Shared secret (password) client user

K_{UK} Session key client-KDC

K_M KDC Master key

Notes: Example Session

1. AS_REQ:

The client initiates a connection to the AS, requesting a TGT. Optionally, the server can require that the clients preauthenticate themselves by using the secret key* to encrypt a timestamp. The request sent contains the client's identity and the identity of the server** in clear text and the optional encrypted timestamp.

2. AS_REP:

The AS_REQ is compared with existing principals to retrieve the shared secret key.   A normal response is a Ticket Granting Ticket (TGT) and a Session Key, which will be used for further communication with the KDC. All are encrypted with the client's secret key. By using a TGT, the client does not have to use its own secret key every time a request is made for credentials to a new service. The TGT in the reply itself is encrypted with the KDC’s master key.
Usually the TGT has a lifetime of 8 to 10 hours.

*The secret key is derived from the password that the user enters the first time he signs in to the Kerberos service. In a Windows 2000 environment, the secret key is generated at the time of logging on to the Domain. Smartcards can also be used to increase the security level of the client and storing the secret key.

** The TGS server's identity is "krbtgt".

Example Session (cont‘d)

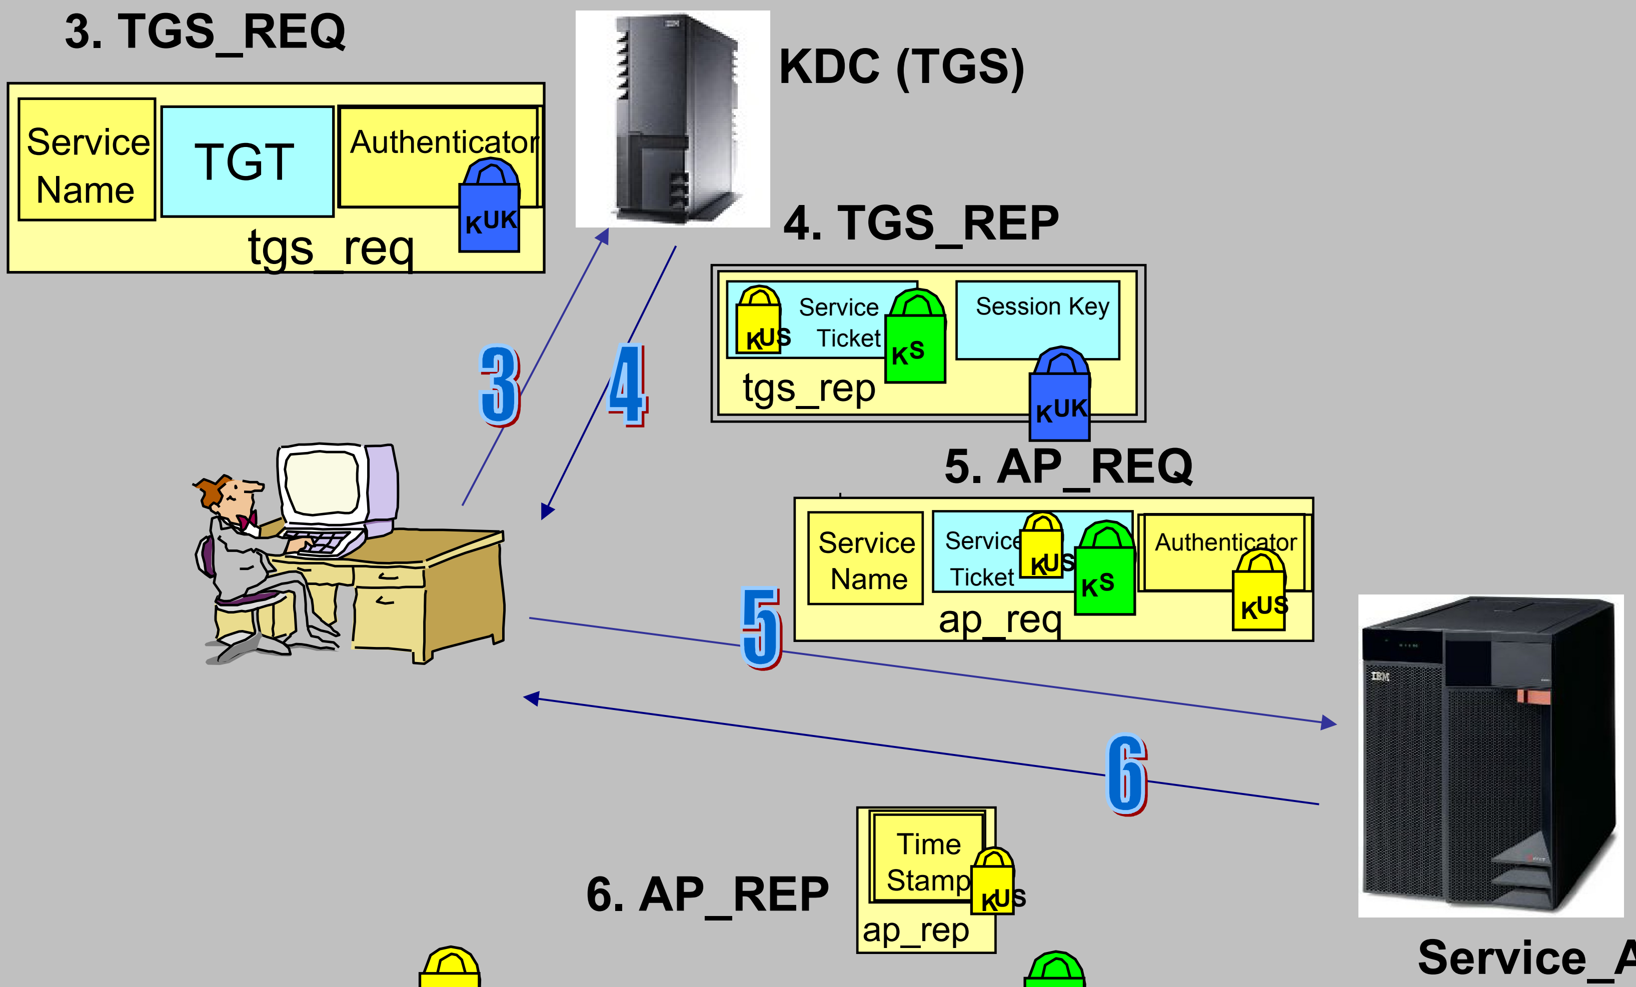

3. TGS_REQ

Service Name | TGT | Authenticator K^{UK}

tgs_req

KDC (TGS)

4. TGS_REP

Service Ticket K^{US} K^{S} | Session Key K^{UK}

tgs_rep

5. AP_REQ

Service Name | Service Ticket K^{US} K^{S} | Authenticator K^{US}

ap_req

6. AP_REP

Time Stamp K^{US}

ap_rep

Service_A

K^{UK} Session key client-KDC

K^{US} Session key client-Service_A

K^{S} Shared secret (password) Service_A

Notes: Example Session (cont’d)

These steps (3 through 5) are repeated for every new service requested.

3. TGS_REQ

When the client wants to initiate a connection with a service, the client first requests a service ticket from the ticket-granting server. This request consists of the service name, the TGT and an authenticator proving the identity of John. This transaction uses the session key the client received earlier from the AS_REP to encrypt the authenticator.

4. TGS_REP

The TGS responds with a service ticket for the requested service and a session key. This response is encrypted with the session key received earlier with the TGT. Except for the initial fields, the client is not able to decrypt the service ticket. The service ticket can only be used to be forwarded to the intended service. This is why the session key also is sent "outside" of the ticket for the client.

5. AP_REQ

The client can now forward the service ticket, along with an authenticator. After the server validates that the ticket came from the trusted third party, KDC, a session is established. The client used the session key to encrypt the authenticator, which the server can read once the ticket is decrypted with the server's shared secret.

6. AP_REP

Optionally, the client could require the server to authenticate itself by using the session key to encrypt a timestamp. This would prove that the server actually managed to decrypt the service ticket and used the session key for response.

Single Signon with EIM and Kerberos in a Windows Network

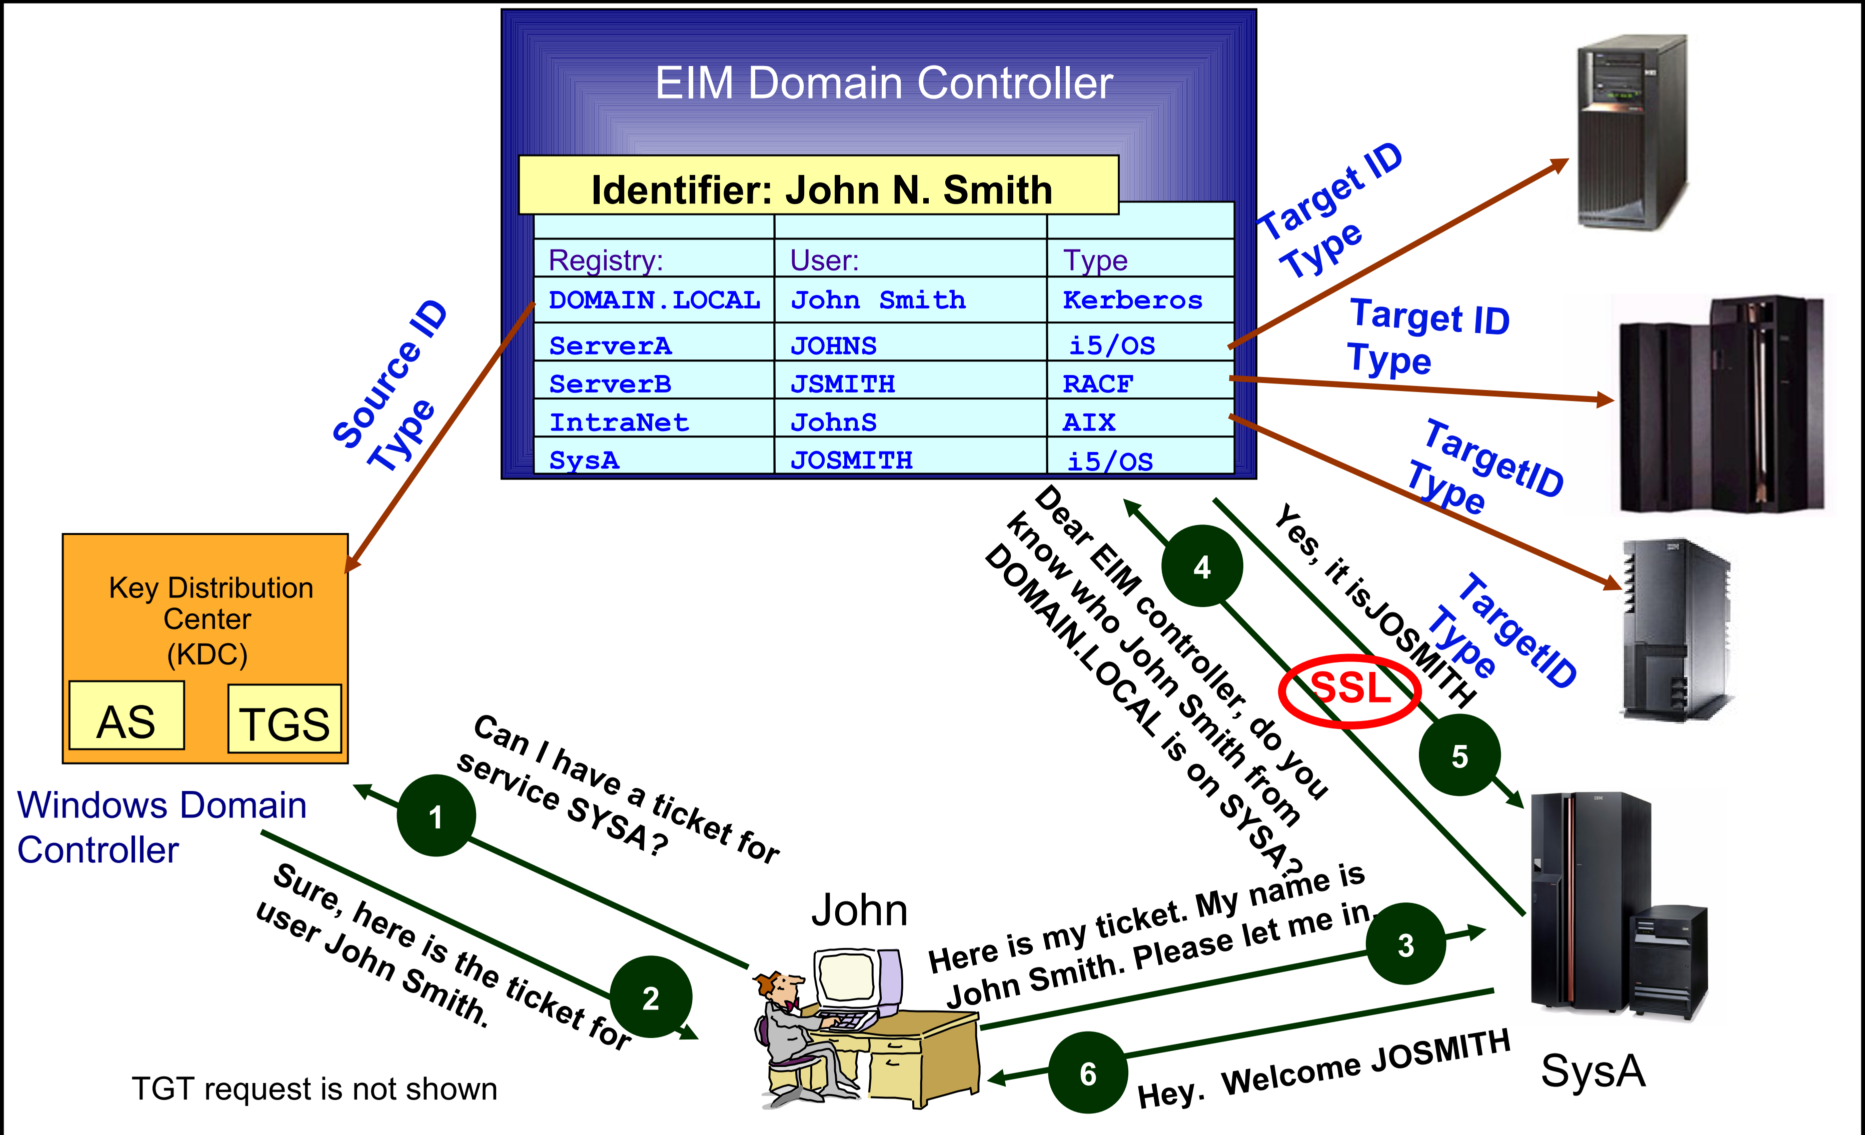

EIM Domain Controller

Identifier: John N. Smith

Registry:	User:	Type
DOMAIN.LOCAL	John Smith	Kerberos
ServerA	JOHNS	i5/OS
ServerB	JSMITH	RACF
IntraNet	JohnS	AIX
SysA	JOSMITH	i5/OS

1 Can I have a ticket for service SYSA?

2 Sure, here is the ticket for user John Smith.

3 Here is my ticket. My name is John Smith. Please let me in.

4 Dear EIM controller, do you know who John Smith from DOMAIN.LOCAL is on SYSA?

5 Yes, it isJOSMITH

6 Hey. Welcome JOSMITH

TGT request is not shown

Notes: EIM and Kerberos – Working Together

The following steps summarize how EIM and Kerberos are used for single sign-on assuming the client already has a TGT:

1.) Credentials for a service are requested from the TGS.

2.) A service ticket is returned for Sys_A.

3.) The client requests access to the service on SysA using the service ticket.

4.) Sys_A, which is capable of handling EIM requests, uses EIM APIs to forward the user identity to the EIM domain controller.

The EIM controller looks at the "source" user and registry to find an identifier in the EIM database.

5.) The EIM server returns the user ID for which that identifier has a "target" registry entry.

6.) Sys_A opens the connection for John Smith and lets him in as the IBM i (i5/OS) user JOSMITH, with appropriate authorizations.

Agenda

- An introduction to single sign-on
- Possible solutions
- Technologies that enable Single Sign-on with password elimination in IBM i
- Implementation overview

Prerequisites

- IBM i
  - Min. OS/400 V5R2 or i5/OS or IBM i (57xx-SS1)
    - Including Qshell interpreter (Option 30) and Host Servers (opt.12)
    - IBM i Access for Windows (57xx-XE1) / System i Access for Windows
- Client
  - Windows XP/Vista/7/8
  - IBM i Access for Windows (Version 5 Release 2 or higher)
  - IBM i Navigator including the "Network" and "Security" components (for administration)
  - Other clients that support Kerberos authentication
- KDC
  - Supporting Kerberos Version 5
  - IBM i KDC support added with i5/OS
    - with V5R4 and higher shipped with 57xx-NAE
  - Windows 2008 and Windows 2012 server with Active Directory
  - Linux KDC (MIT or Heimdal)

Notes: Prerequisites

- The previous chart lists all the prerequisites that need to be met by the server and client environments to implement SSO with Kerberos and EIM.

Kerberos and EIM-enabled applications

- Host servers (used by IBM i Access for Windows)
- Telnet server used by PC5250 from IBM i Access, WebSphere Host On-Demand V8, 5250 emulator in IBM i Access for Linux V1.8,
  IBM Personal Communications 5.9
- IBM i Telnet client (V7R2)
- QFileSrv.400
- Distributed Relational Database Architecture (DRDA), Open Database Connectivity (ODBC), Java Database Connectivity (JDBC)
- HTTP Server for IBM i (powered by Apache)
- Management Central
- Lightweight Directory Access Protocol (LDAP) Server (**Kerberos authentication only, no EIM involved**)
- Windows Integration
- FTP Client and Server (V7R2)
- NetServer
- IBM WebSphere Application Server V6.1 (only Kerberos for WebSphere authentication and Identity Tokens and EIM to backend IBM i)
- Network File System (NFS)

Notes: Kerberos and EIM-enabled applications

- IBM i / i5/OS client and server applications that are currently enabled for SSO are:
  - IBM i Host Servers (57xx-SS1 Option 12): Currently used by IBM i Access for Windows and IBM System i Navigator.
  - Telnet server: Currently used by PC5250 and IBM WebSphere Host On-Demand Version 8 and higher: Web Express Logon feature. The 5250 emulation of IBM i Access for Linux V1.8 also supports Single Signon with Kerberos.
  - Open Database Connectivity (ODBC): Allows SSO access to IBM i databases through ODBC.
  - Java Database Connectivity (JDBC): Allows SSO access to IBM i databases through ODBC.
  - Distributed Relational Database Architecture (DRDA): Allows SSO access to IBM i databases through ODBC.
  - QFileSrv.400
  - LDAP Server: Supports Kerberos authentication only. EIM is not used during the authentication process.

- The following applications were enabled for EIM, Kerberos, or both in V5R3:
  - Management Central for authentication between endpoint systems and the central system.
  - Windows Integration for user enrollment and for submitting network server commands.
  - HTTP Server for iSeries (powered by Apache) when using Microsoft's Internet Explorer 5.0 or later. This support was also added to V5R2 via the HTTP group PTF.
  - The V5R3 enhancement of storing user certificates in LDAP servers also provides the ability for OS/400 applications, such as the FTP server, to use EIM for lookup operation of a target association. This function only pertains to OS/400 applications using digital certificates for client authentication. It is not related to Kerberos at all.

- In IBM i V6R1, the Network File System (NFS) service has also been enabled for Kerberos.
- With IBM WebSphere Application Server V6.1, a SPENGO Trust Association Interceptor (TAI) is shipped with the product. It provides Kerberos authentication support for Web-based browser applications. The authentication is completed when the TAI verified the ticket, extracted the user principal name, and checked that the user exists in the configured WebSphere user registry. The WebSphere application will then work with the Kerberos user name. In addition, IBM i is shipped with a J2C connector that allows you to do single signon authentication from the Web application to an IBM i backend through the Java Toolbox (JTOpen Toolbox). In this case, ID Tokens will be used for authentication from WebSphere to IBM i and EIM for mapping the WebSphere user to the local IBM i user profile.

Implementation Overview

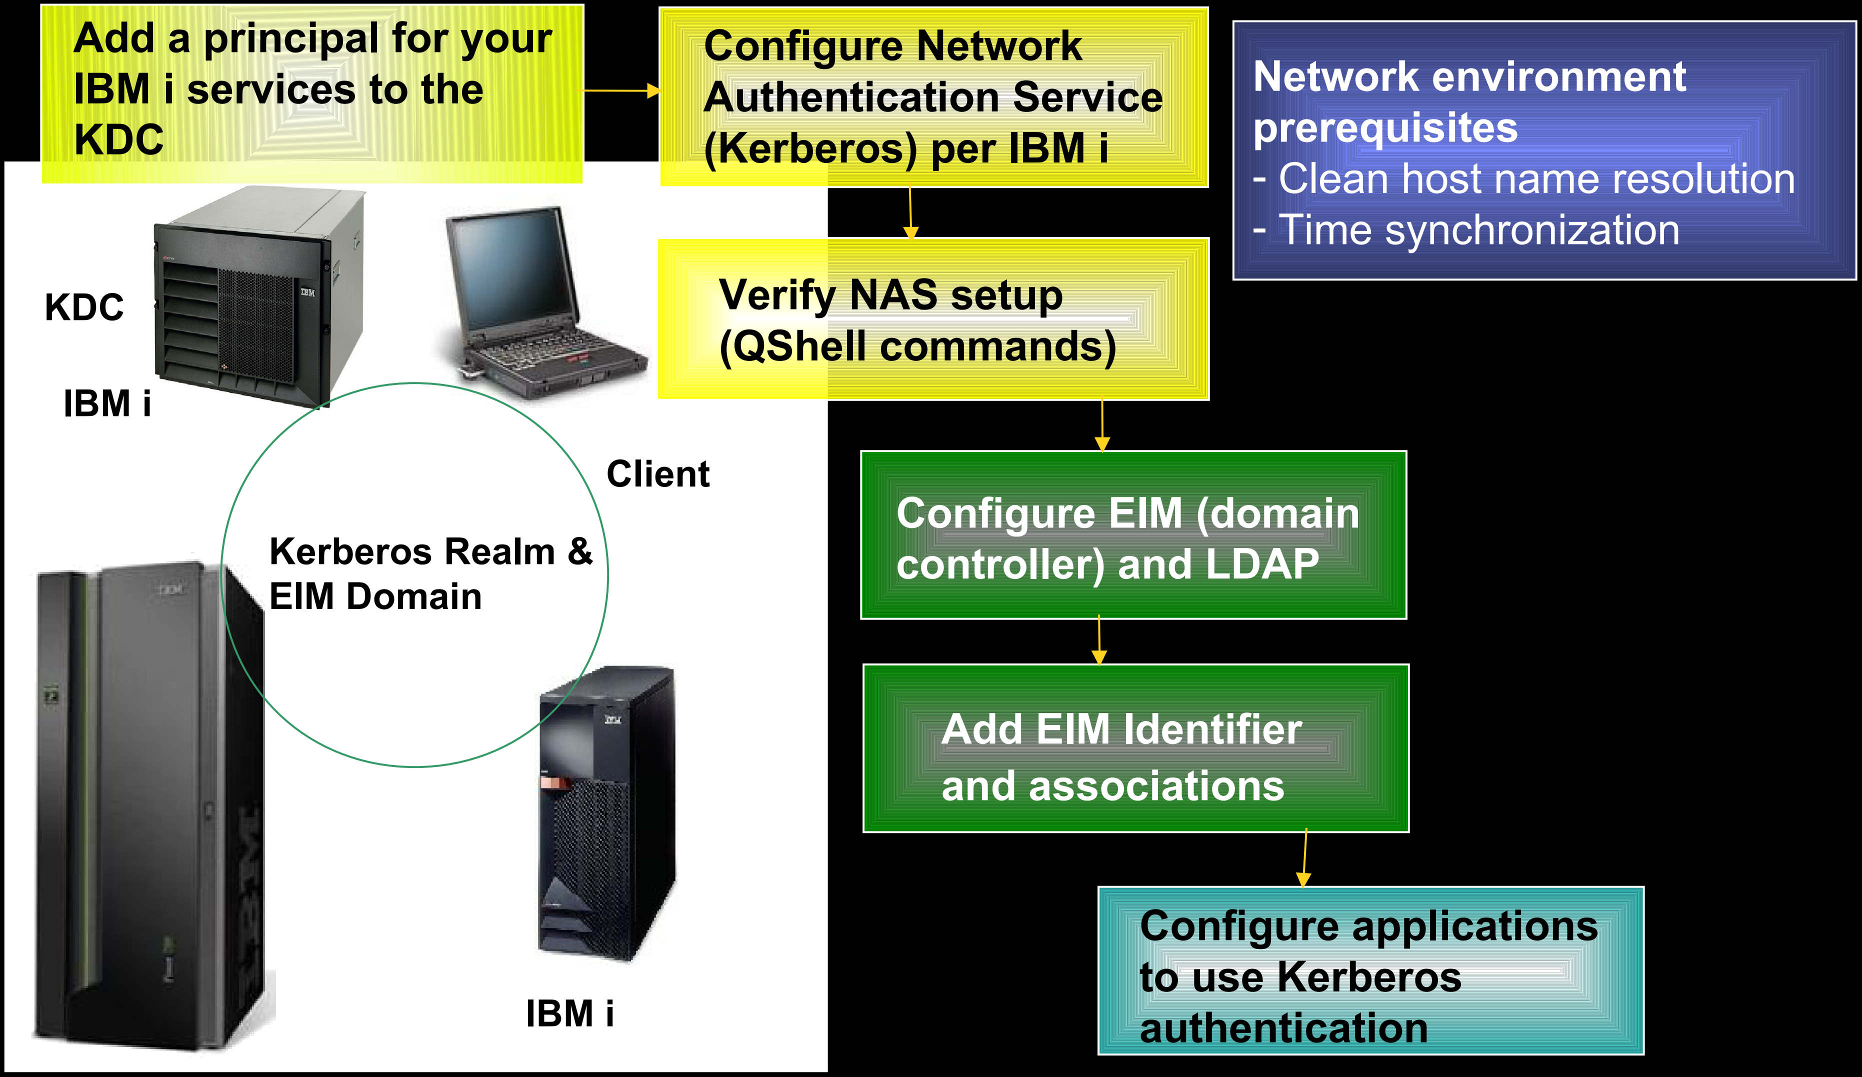

Add a principal for your IBM i services to the KDC

Configure Network Authentication Service (Kerberos) per IBM i

Verify NAS setup (QShell commands)

Configure EIM (domain controller) and LDAP

Add EIM Identifier and associations

Configure applications to use Kerberos authentication

Network environment prerequisites
- Clean host name resolution
- Time synchronization

Notes: Implementation Overview

- The previous chart summarizes the implementation steps for setting up an SSO environment with Windows and iSeries servers. The prerequisites in the network environment that need to be met is a clean DNS name resolution and a time synchronization between all participating systems.

DNS Resolution

- Clean IP name resolution is important with Kerberos
- Before setting up Kerberos, all IP addresses of services in a network should be resolved to the same host name
- Problems can arise when using different DNS servers or host tables
- In IBM i, the host name to which the IBM i IP address resolves in a network has to be added as the first entry in the host table

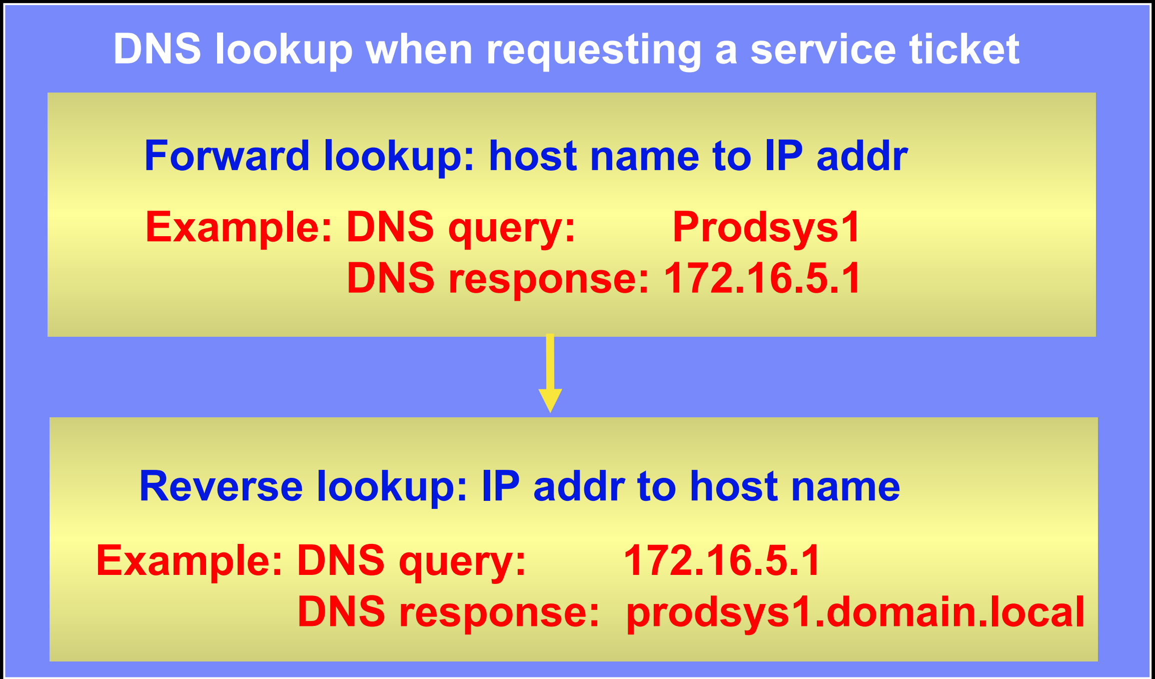

DNS lookup when requesting a service ticket

Forward lookup: host name to IP addr

Example: DNS query: Prodsys1
DNS response: 172.16.5.1

Reverse lookup: IP addr to host name

Example: DNS query: 172.16.5.1
DNS response: prodsys1.domain.local

Example SPN: krbsvr400/*prodsys1.domain.local*@DOMAIN.LOCAL

Notes: DNS Resolution

- All IP addresses of kerberized services should resolve to the same IP host name in a network. Typically a client application that wants to authenticate with Kerberos to service, performs the following steps when requesting a service ticket:
  - If, for example, IBM System i Navigator has a connection configured to Prodsys1, the application first does a forward DNS lookup for host name Prodsys1.
  - When the client application receives the DNS response, for example, the IP address 172.16.5.1, it performs a reverse lookup for that IP address. Typically, a DNS returns then the fully qualified domain name, such as prodsys1.test.com.
- When IP addresses get resolved to different host names (different DNS servers or local host tables) a different service principal will be created for the service ticket request. This can cause Kerberos authentication problems. Therefore, it is important to have a consistent name resolution across the network.
- For IBM i, the host table must contain the network resolved host name in the first position of the host table entry. For example, if the short name of the IBM i partition is prodsys1 and the fully qualified host name is prodsys1.test.com, which is the DNS resolved name, the host table should be:

```
      Internet          Host
   Opt   Address          Name

         172.16.5.111     prodsys1.test.com
                          prodsys1
```

- Proper name resolution can also be tested from the client side by using the following command:

```
ping –a 172.16.5.1
```

Configuring NAS for IBM i

- Network Authentication Services (NAS) can be easily configured via the NAS wizard.

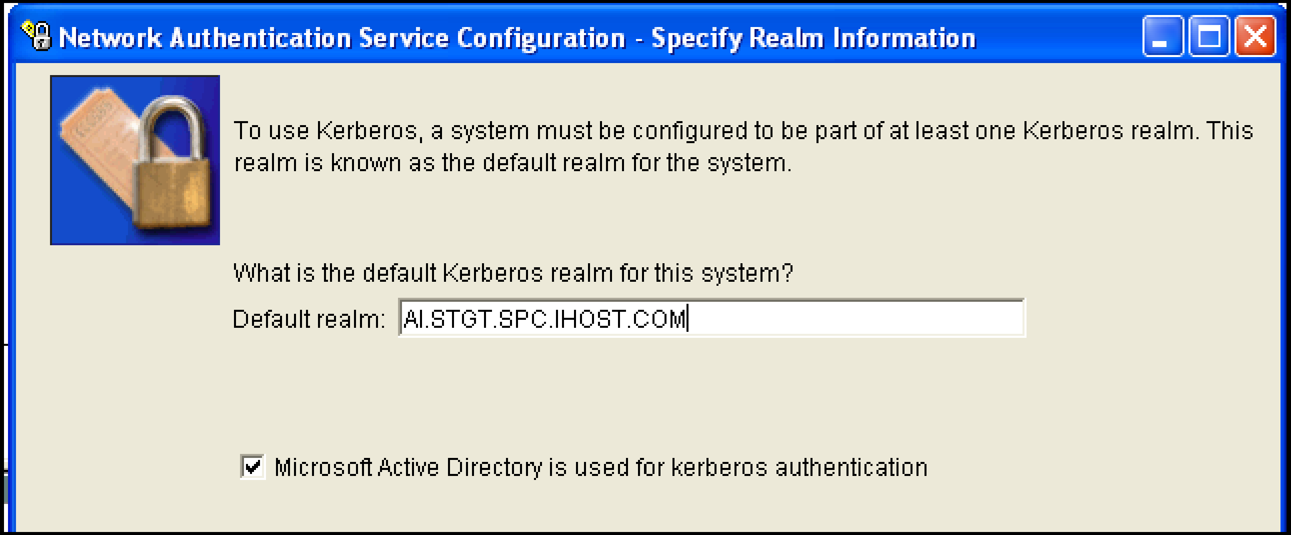

Network Authentication Service Configuration - Specify Realm Information

To use Kerberos, a system must be configured to be part of at least one Kerberos realm. This realm is known as the default realm for the system.

What is the default Kerberos realm for this system?

Default realm: AI.STGT.SPC.IHOST.COM

☑ Microsoft Active Directory is used for kerberos authentication

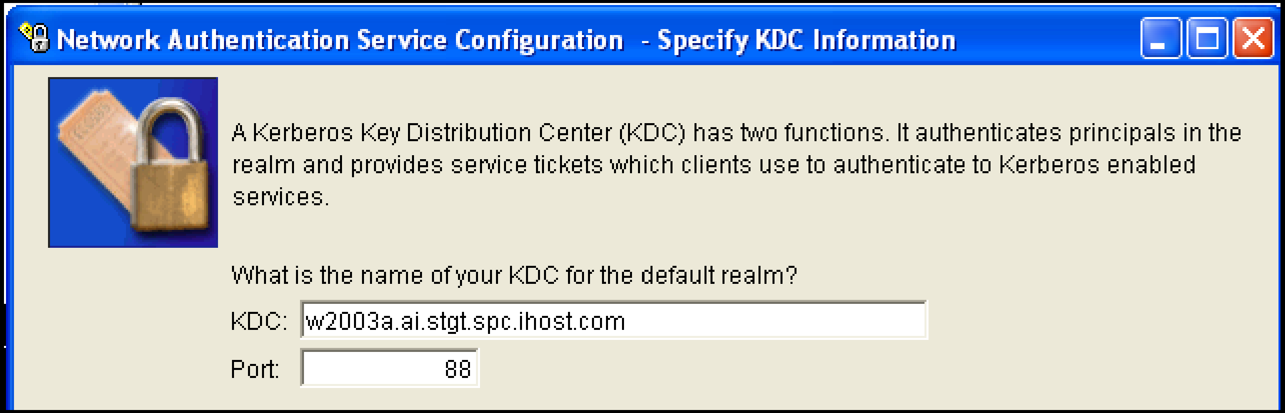

Network Authentication Service Configuration - Specify KDC Information

A Kerberos Key Distribution Center (KDC) has two functions. It authenticates principals in the realm and provides service tickets which clients use to authenticate to Kerberos enabled services.

What is the name of your KDC for the default realm?

KDC: w2003a.ai.stgt.spc.ihost.com

Port: 88

Notes: Configuring NAS for IBM i

- The Network Authentication Service (NAS) or Kerberos configuration can be easily performed by the IBM System i Navigator NAS wizard. An administrator need to provide the following information when running the wizard:
  - Kerberos realm name
  - Host name or IP address and port of the KDC
  - Host name or IP address and port of the password server (used for remote password change)
  - Selection of IBM i services to be configured for Kerberos along with a password. This password is used to generate a shared secret that is needed for Kerberos authentication. The password must match the password that is specified for the service principal account on the KDC.

Configuring NAS for IBM i (cont'd)

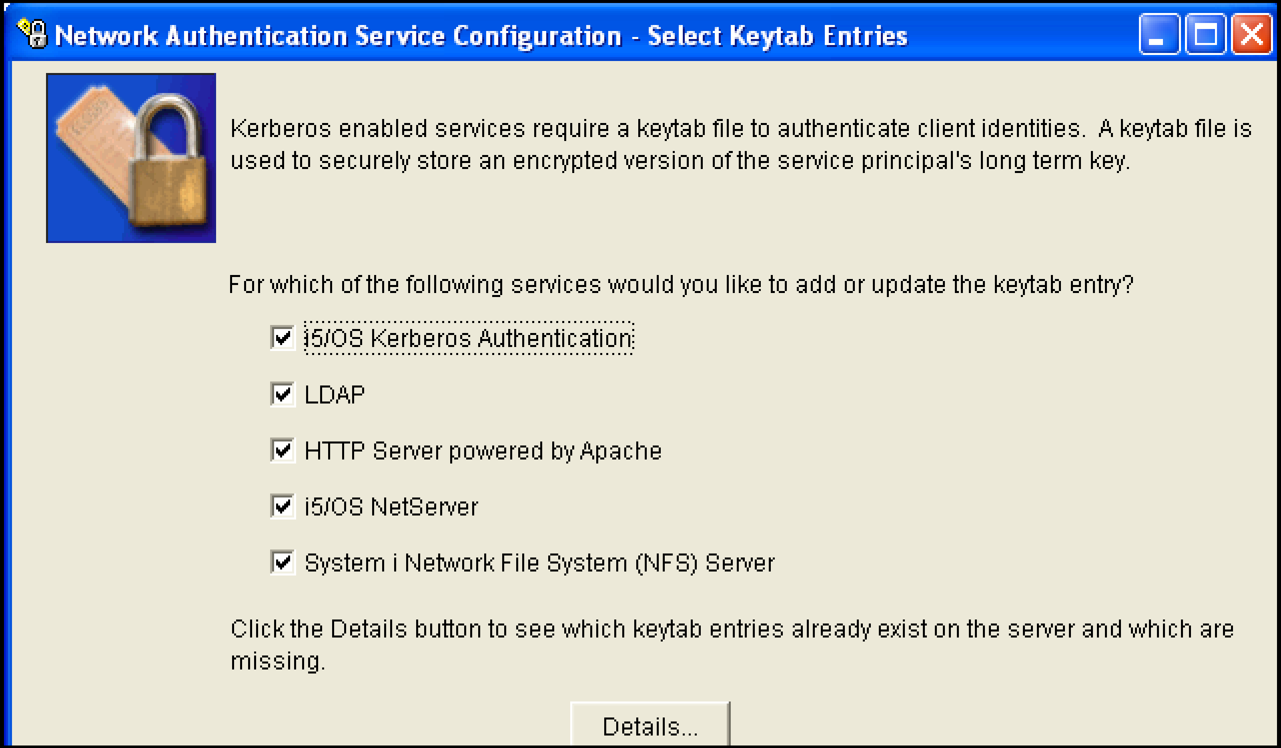

Network Authentication Service Configuration - Select Keytab Entries

Kerberos enabled services require a keytab file to authenticate client identities. A keytab file is used to securely store an encrypted version of the service principal's long term key.

For which of the following services would you like to add or update the keytab entry?

☑ i5/OS Kerberos Authentication

☑ LDAP

☑ HTTP Server powered by Apache

☑ i5/OS NetServer

☑ System i Network File System (NFS) Server

Click the Details button to see which keytab entries already exist on the server and which are missing.

Select the services that you want to include in the single sign-on environment

Notes: Configuring NAS for IBM i (cont'd)

The wizard lets you select different services that are enabled for Kerberos and that can participate in a single sign-on environment. You need to select the services that you want to include in your SSO environment.

Commonly used services are i5/OS Kerberos Authentication and i5/OS NetServer.

Kerberos Key Table

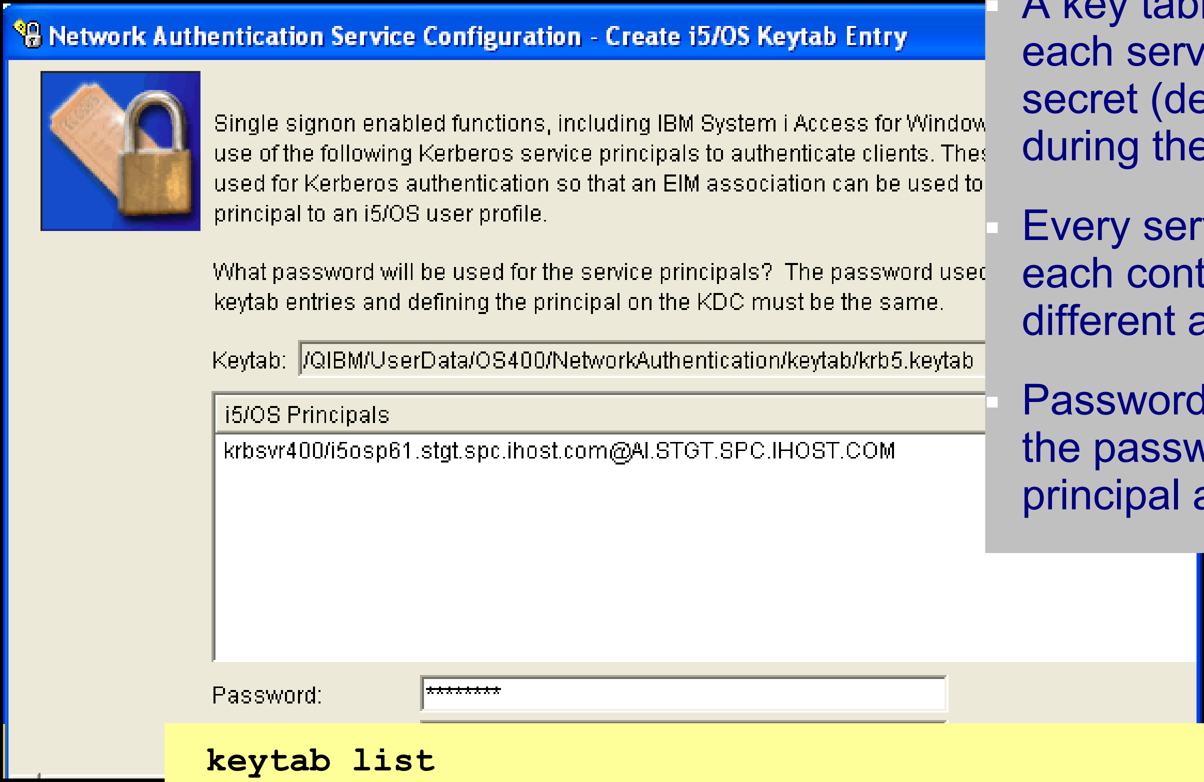

Network Authentication Service Configuration - Create i5/OS Keytab Entry

Single signon enabled functions, including IBM System i Access for Window
use of the following Kerberos service principals to authenticate clients. Thes
used for Kerberos authentication so that an EIM association can be used to
principal to an i5/OS user profile.

What password will be used for the service principals? The password used
keytab entries and defining the principal on the KDC must be the same.

Keytab: /QIBM/UserData/OS400/NetworkAuthentication/keytab/krb5.keytab

i5/OS Principals
krbsvr400/i5osp61.stgt.spc.ihost.com@AI.STGT.SPC.IHOST.COM

Password: *********

- A key table (keytab) contains entries for each service principal along with a shared secret (derived from the password entered during the NAS configuration)
- Every service principal has multiple entries, each containing a secret key created with a different algorithm
- Password in the wizard must match the password of the KDC service principal account

```
keytab list
 Key table: /QIBM/UserData/OS400/NetworkAuthentication/keytab/krb5.keytab
Principal: krbsvr400/i5osp61.stgt.spc.ihost.com@AI.STGT.SPC.IHOST.COM
  Key version: 1
  Key type: 56-bit DES
  Entry timestamp: 2011/08/12-10:34:12
Principal: krbsvr400/i5osp61.stgt.spc.ihost.com@AI.STGT.SPC.IHOST.COM
  Key version: 1
  Key type: 56-bit DES using key derivation
  Entry timestamp: 2011/08/12-10:34:13
```

Notes: Kerberos Key Table

- The best way to explain the purpose of the key table is by describing the authentication process of a user. When a user logs in to his Windows client, he enters a user ID and a password. From the entered password, the client generates a shared secret that is used to encrypt information (time) during authentication. The KDC on the other hand has also access to the shared secret (stored with the user account) and can decrypt the information.
  When a service ticket is issued by the KDC, it is also encrypted by a shared secret of the service principals account. However, the application service needs the corresponding shared secret to decrypt the service ticket. Since there is nobody at the system who enters a password every time an authentication request is received, the shared secrets are stored in a key table file.

- A key table contains entries for every service principal. For a single service, the key table contains multiple entries. They are distinguished by the shared secret. Each shared secret was generated by a different algorithm, such as DES 56 bit, 3DES, etc.

- Note that RC4 and AES encryption support has been added for V5R4 and higher in August 2011.

Adding service principals to the KDC – Windows AD

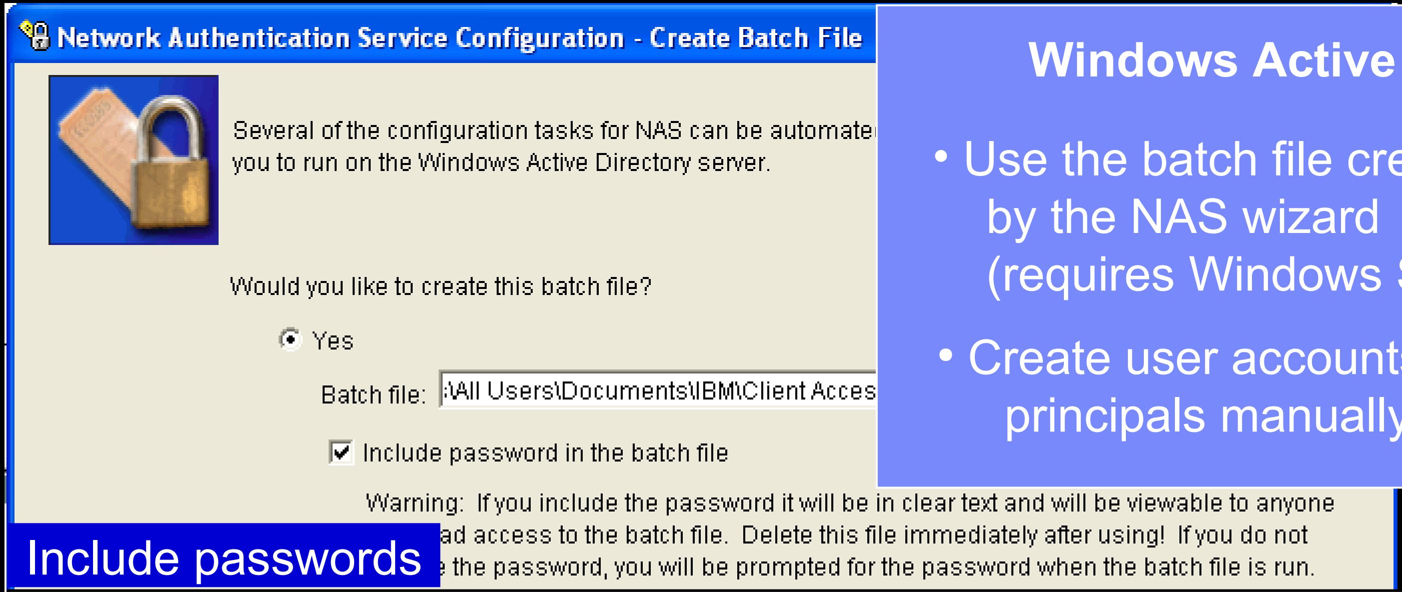

Network Authentication Service Configuration - Create Batch File

Several of the configuration tasks for NAS can be automate
you to run on the Windows Active Directory server.

Would you like to create this batch file?

● Yes

Batch file: :\All Users\Documents\IBM\Client Acces

☑ Include password in the batch file

Warning: If you include the password it will be in clear text and will be viewable to anyone
ad access to the batch file. Delete this file immediately after using! If you do not
e the password, you will be prompted for the password when the batch file is run.

Windows Active Directory

- Use the batch file created by the NAS wizard (requires Windows Support Tools)
- Create user accounts and map principals manually

Include passwords

```
DSADD user cn=i5osp61_1_krbsvr400,cn=users,dc=AI,dc=STGT,dc=SPC,dc=IHOST,dc=COM
  -pwd mysecret -display i5osp61_1_krbsvr400
KTPASS -MAPUSER i5osp61_1_krbsvr400
  -PRINC krbsvr400/i5osp61.stgt.spc.ihost.com@AI.STGT.SPC.IHOST.COM
  -PASS mysecret -mapop set -ptype KRB5_NT_PRINCIPAL
```

No passwords included

```
DSADD user cn=i5osp61_1_krbsvr400,cn=users,dc=AI,dc=STGT,dc=SPC,dc=IHOST,dc=COM
  -pwd * -display i5osp61_1_krbsvr400
KTPASS -MAPUSER i5osp61_1_krbsvr400
  -PRINC krbsvr400/i5osp61.stgt.spc.ihost.com@AI.STGT.SPC.IHOST.COM
  -PASS * -mapop set -ptype KRB5_NT_PRINCIPAL
```

Notes: Adding service principals to the KDC

- Every kerberized service has a service name. When a client requests a service ticket for a specific service from the KDC, it requests that ticket by specifying the service name. Typically, kerberized services have a service principal in the following format:
  `krbsvr400/fully_qualified_hostname@KERBEROS.REALM`

- i5/OS, for example, has a single service name for all host servers, DDM/DRDA, Telnet, and QFileSvr.400. It is krbsvr400. The NetServer, HTTP server, NFS server, and LDAP server have different service names.

- The service principal names need to be registered with the KDC. The approach of adding a service principal depends on the platform the KDC runs on.

- For Windows Active Directory, a user account needs to be created first. After the user has been created, the service principal needs to be mapped to the user account via the ktpass command. This command is part of the Windows Support Tools.

- For the i5/OS KDC, service principals are added through the kadmin utility in the PASE environment. The addprinc utility command adds a service principal to the KDC. In this case, no separate user account is required.

- The password of the Windows user account or i5/OS KDC service principal must match the password that was specified when running the NAS wizard.

Verifying the IBM i NAS Setup

- Verification is not mandatory, but strongly recommended
- Verification ensures:
  - Synchronized time between KDC and IBM i and correct DNS name resolution for IBM i
  - Service principals are correctly registered in keytab file and KDC along with their passwords
- To verify the IBM i NAS setup, you need to:
  - Sign on to IBM i with a user profile that has a home directory (required for credentials cache)
  - Start a QShell session
  - Display the keytab file with the `keytab list` command
  - Obtain an initial ticket (TGT) for a service with the `kinit` command
  - Display the TGT with the `klist` command

```
 > kinit -k krbsvr400/i5osp61.stgt.spc.ihost.com@AI.STGT.SPC.IHOST.COM

   $
 > klist

    Ticket cache:
FILE:/QIBM/USERDATA/OS400/NETWORKAUTHENTICATION/creds/krbcred_0dfdbd13
    Default principal: krbsvr400/i5osp61.stgt.spc.ihost.com@AI.STGT.SPC.IHOST.COM

  Server: krbtgt/AI.STGT.SPC.IHOST.COM@AI.STGT.SPC.IHOST.COM

     Valid 2014/08/19-11:10:38 to 2014/08/19-21:10:38
```

Notes: Verifying the IBM i NAS Setup

- These steps are not required for the Network authentication to work. However, by performing these steps, you confirm that the Kerberos environment is working correctly.
  Note: The user performing these steps must have a home directory in the IFS. The home directory stores the krb5ccname file, containing the link to the credential cache.

- > keytab list
  - This lists the current keys in the Kerberos key table. If the wizard completed correctly and made contact with the KDC, it should now contain three entries for the krbsvr400 principal (at different encryption levels). If the principal name of the krbsvr400 service displays a wrong host name, verify that the host table on the PC you are performing the configuration on has the correct entries.

- > kinit -k krbsvr400/i5osp61.stgt.spc.ihost.com@AI.STGT.SPC.IHOST.COM
  - This requests a TGT from the KDC. This should complete with out error and return the prompt.*

- > klist
  - This lists the tickets in the ticket cache and should display the newly received ticket from the KDC.

- * Some errors that could occur at the kinit stage:

  Unable to obtain initial credentials.

Status 0x96c73a06 - Client principal is not found in security registry.	The krbsvr400 principal had been misspelled.
Status 0x96c73a25 - Time differential exceeds maximum clock skew.	The KDC was using daylight savings time.
Status 0x96c73a9a - Unable to locate security server.	Realm name resolving incorrectly. Check case sensitivity.

Setting up EIM: LDAP default setup environment

- IBM i Navigator provides a wizard to set up EIM
  - Wizard defaults to store EIM domain data under root of the LDAP directory
  - The wizard also defaults to use the LDAP administrator as the IBM i system user

Easy to set up, but not a good practice

Prevents good directory information tree structure and provides administrator access to IBM i partition

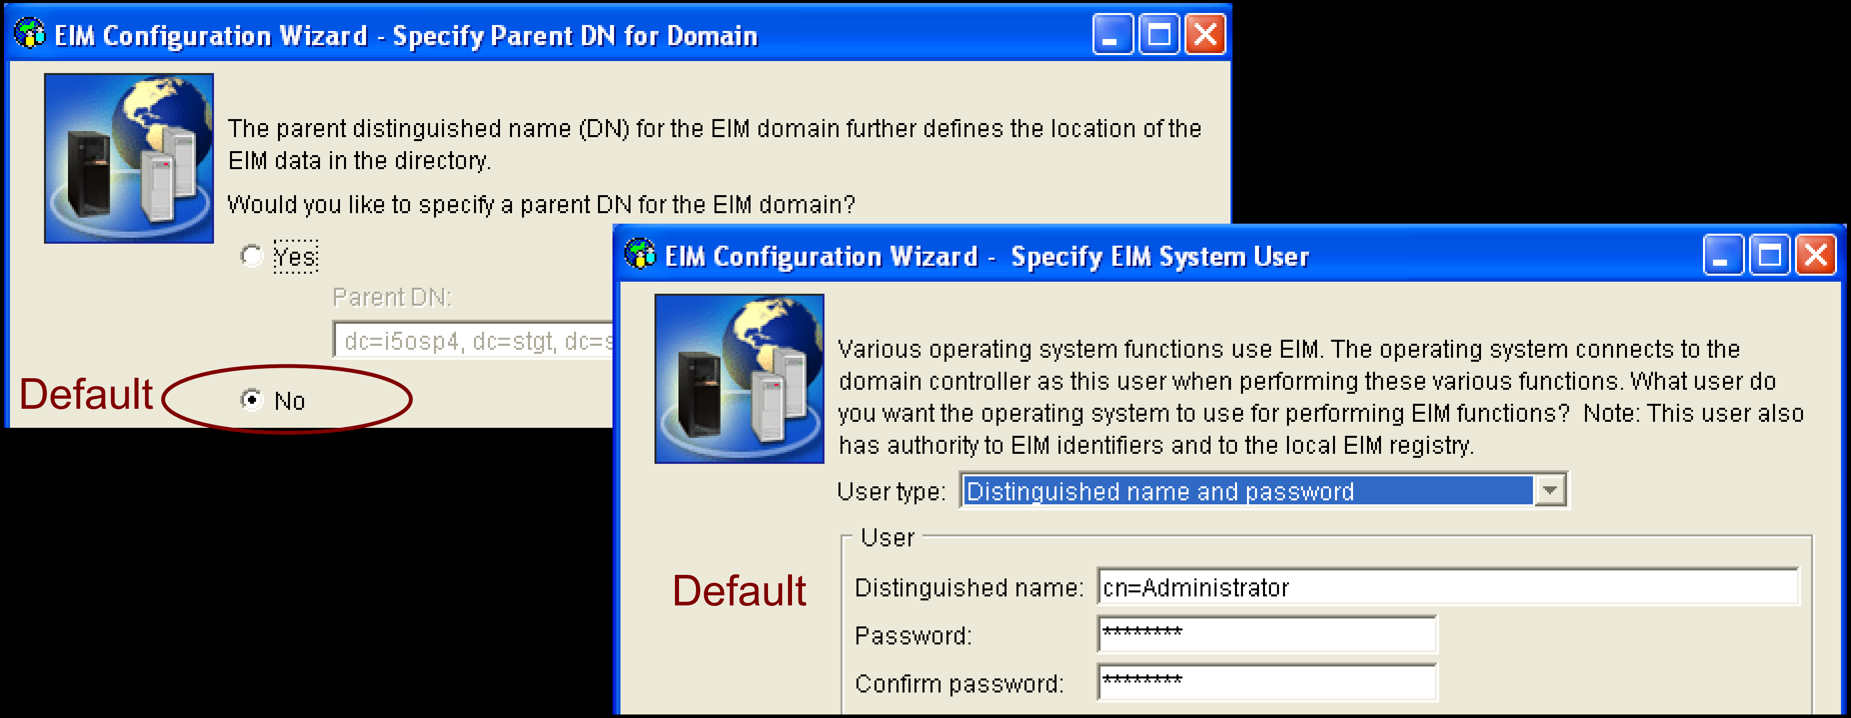

EIM Configuration Wizard - Specify Parent DN for Domain

The parent distinguished name (DN) for the EIM domain further defines the location of the EIM data in the directory.

Would you like to specify a parent DN for the EIM domain?

○ Yes

Parent DN: dc=i5osp4, dc=stgt, dc=s

Default ● No

EIM Configuration Wizard - Specify EIM System User

Various operating system functions use EIM. The operating system connects to the domain controller as this user when performing these various functions. What user do you want the operating system to use for performing EIM functions? Note: This user also has authority to EIM identifiers and to the local EIM registry.

User type: Distinguished name and password

User

Default

Distinguished name: cn=Administrator

Password: ********

Confirm password: ********

Recommended LDAP directory server setup for EIM

- Enter a new suffix for EIM to the LDAP server properties
  - i.e. add an organization (o) entry to hold all EIM related data

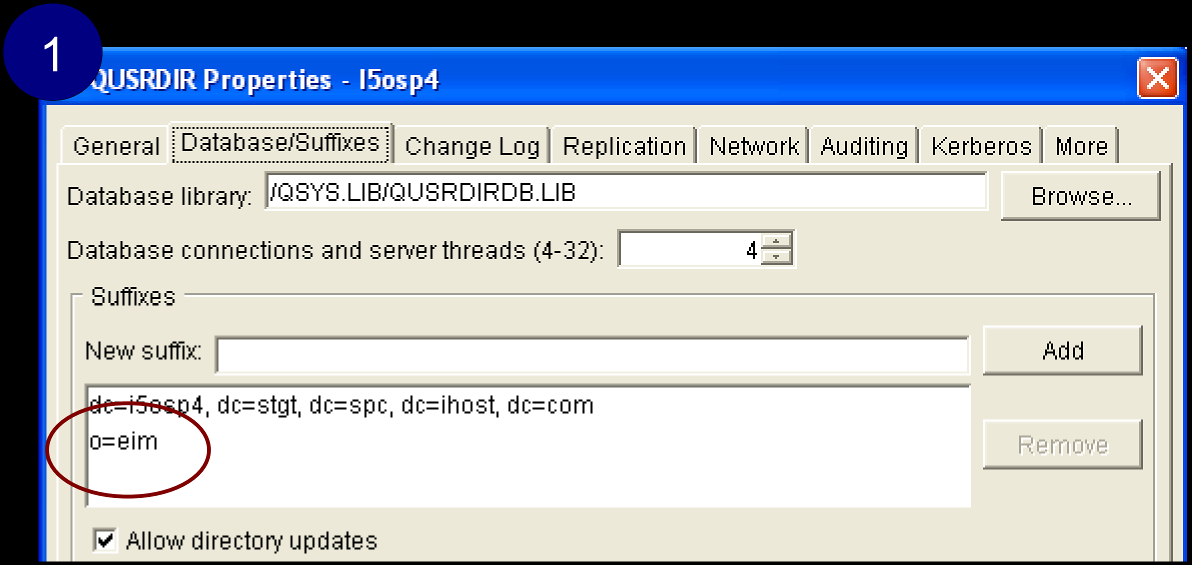

1

QUSRDIR Properties - I5osp4

Database library: /QSYS.LIB/QUSRDIRDB.LIB

Database connections and server threads (4-32): 4

Suffixes

New suffix:

dc=i5osp4, dc=stgt, dc=spc, dc=ihost, dc=com
o=eim

☑ Allow directory updates

Recommended LDAP directory server setup for EIM

Add the suffix distinguished name (DN) as an entry to the LDAP directory

You can either use the IBM Tivoli Directory Server Web Administration tool

OR

Add the entry via the ldapadd utility and an LDIF file (next page...)

2

IBM Tivoli Directory Server Web Administration Tool

Allows you to administer the IBM Tivoli Directory Server for i (Located in Network)

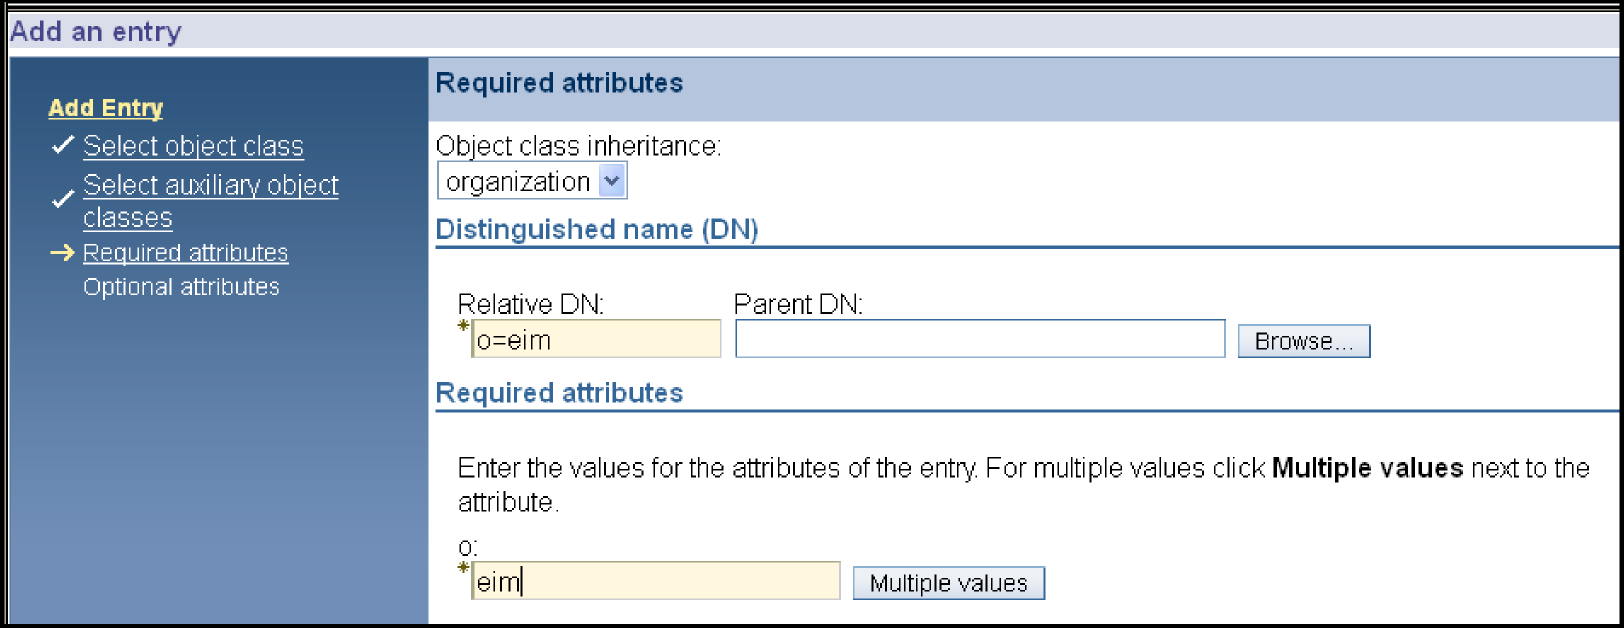

Required attributes

Object class inheritance: organization

Distinguished name (DN)

Relative DN: * o=eim

Parent DN:

Required attributes

Enter the values for the attributes of the entry. For multiple values click **Multiple values** next to the attribute.

o: * eim

Recommended LDAP directory server setup for EIM

Using Qshell with an LDIF file

1. Create an LDIF file with the DN information

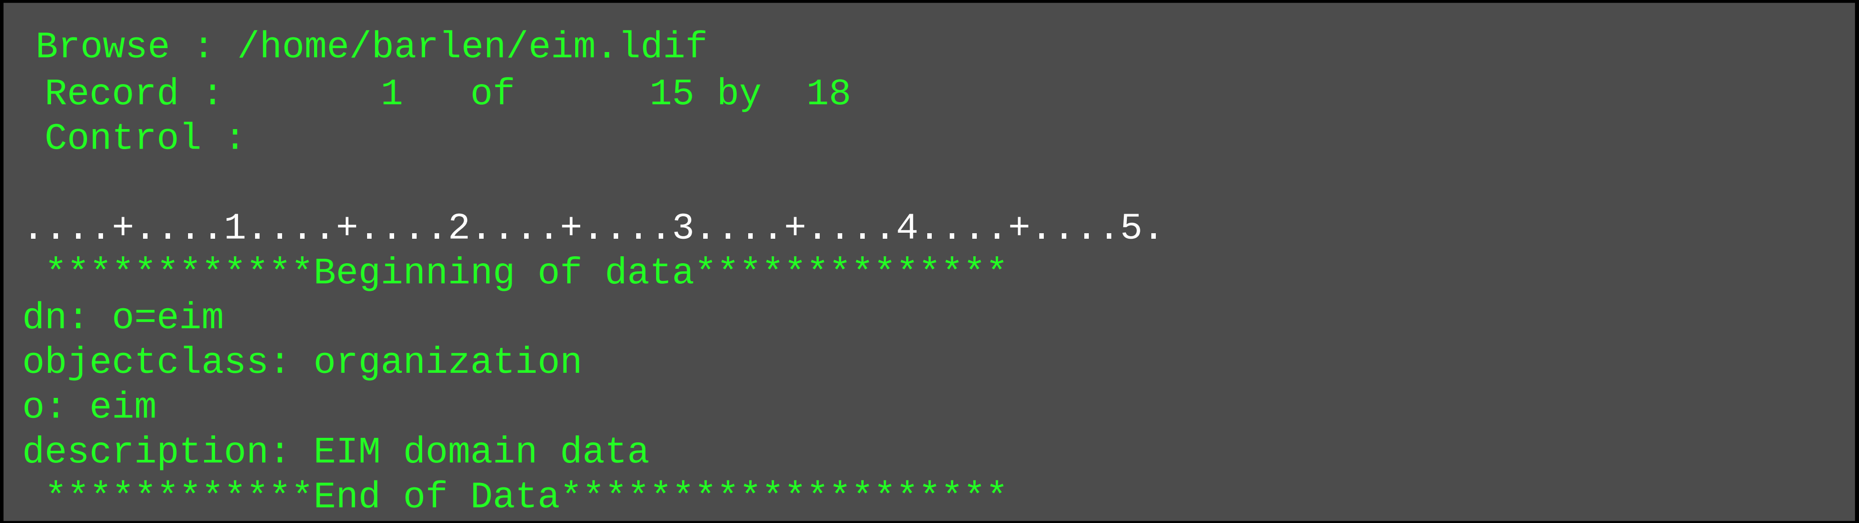

```
 Browse : /home/barlen/eim.ldif
  Record :         1   of       15 by   18
  Control :

....+....1....+....2....+....3....+....4....+....5.
 ************Beginning of data**************
dn: o=eim
objectclass: organization
o: eim
description: EIM domain data
 ************End of Data********************
```

2. Add the entry to the LDAP directory in QShell

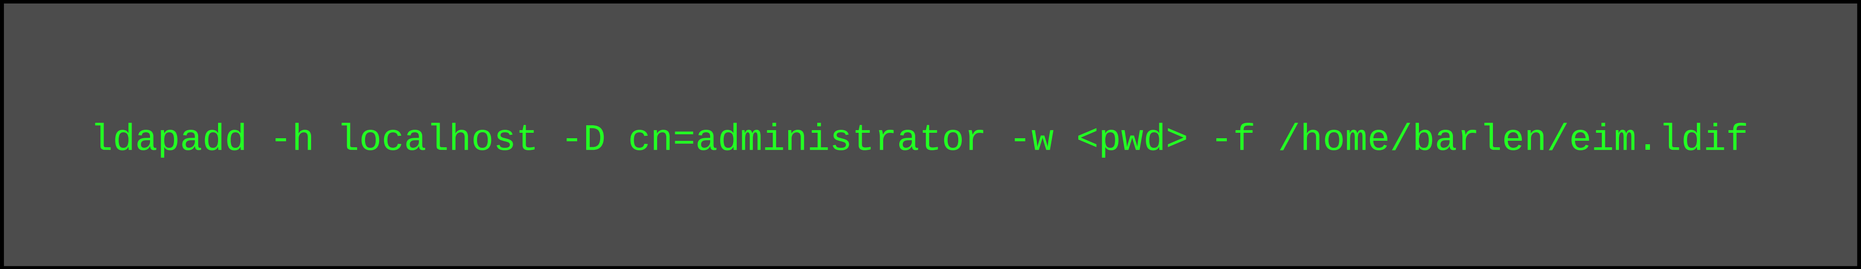

```
ldapadd -h localhost -D cn=administrator -w <pwd> -f /home/barlen/eim.ldif
```

Recommended LDAP directory server setup for EIM

Add a user entry to be used as the IBM i system user for EIM

1. Create an LDIF file with the DN information

```
 ....+....1....+....2....+....3....+....4....+....5.
 ************Beginning of data**************
dn: cn=eimlookup,o=eim
objectclass: inetOrgPerson
objectclass: ePerson
cn: eimlookup
sn: EIM
description: EIM system user
uid: eimlookup
userPassword: eim4all
 ************End of Data********************
```

2. Add the entry to the LDAP directory in QShell

```
ldapadd -h localhost -D cn=administrator -w <pwd> -f /home/barlen/eimuser.ldif
```

Advantage of having a separate user is to restrict access to only EIM data and being able to change the administrator password without impact to EIM

New EIM LDAP environment

LDAP server directory structure

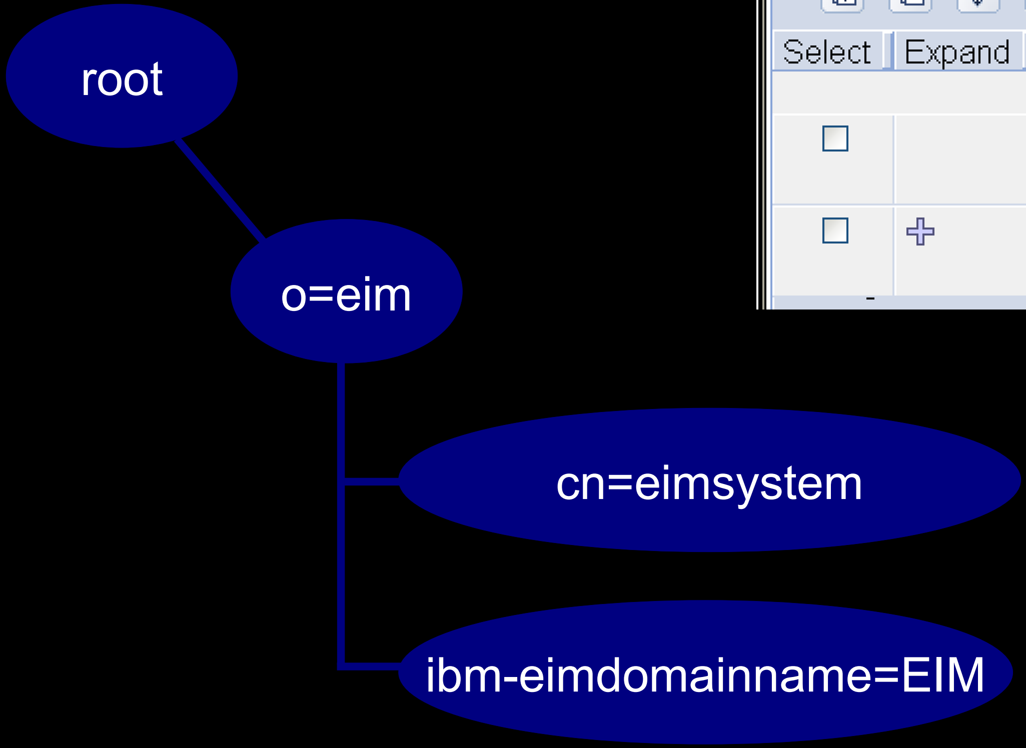

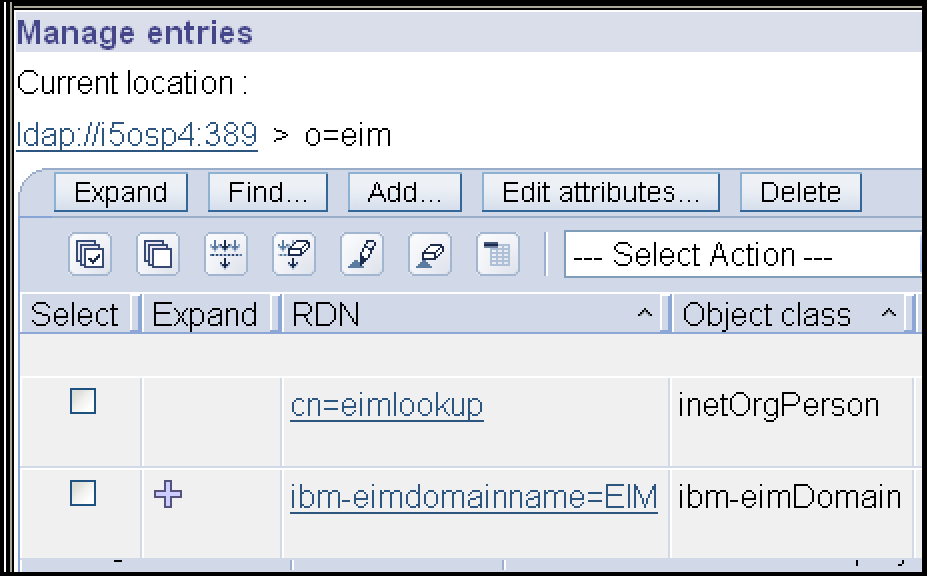

Current location :

ldap://i5osp4:389 > o=eim

Select	Expand	RDN	Object class
☐		cn=eimlookup	inetOrgPerson
☐	+	ibm-eimdomainname=EIM	ibm-eimDomain

Setting up an EIM Domain Controller

- EIM wizard simplifies setup tasks
  - Creates LDAP directory entries and IBM i environment
  - IBM System i Navigator wizard can create domain controller locally or remote

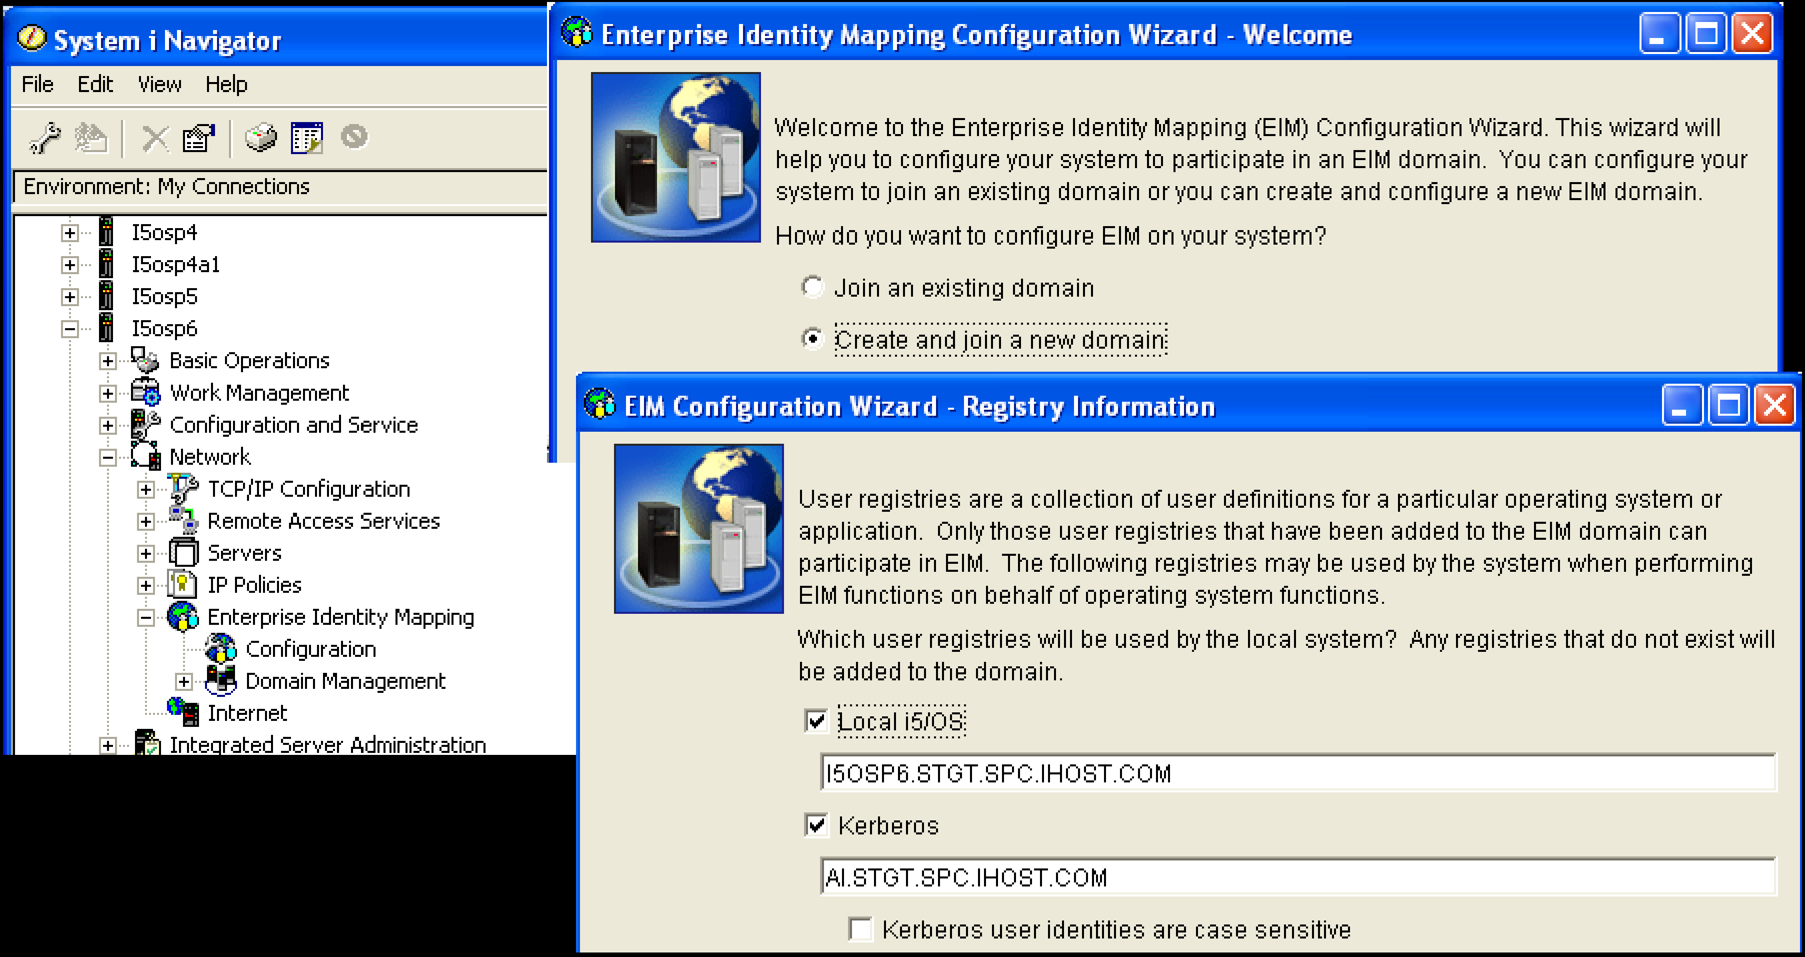
System i Navigator
File Edit View Help
Environment: My Connections
I5osp4
I5osp4a1
I5osp5
I5osp6
Basic Operations
Work Management
Configuration and Service
Network
TCP/IP Configuration
Remote Access Services
Servers
IP Policies
Enterprise Identity Mapping
Configuration
Domain Management
Internet
Integrated Server Administration
Enterprise Identity Mapping Configuration Wizard - Welcome
EIM Configuration Wizard - Registry Information

Welcome to the Enterprise Identity Mapping (EIM) Configuration Wizard. This wizard will help you to configure your system to participate in an EIM domain. You can configure your system to join an existing domain or you can create and configure a new EIM domain.

How do you want to configure EIM on your system?

○ Join an existing domain

● Create and join a new domain

User registries are a collection of user definitions for a particular operating system or application. Only those user registries that have been added to the EIM domain can participate in EIM. The following registries may be used by the system when performing EIM functions on behalf of operating system functions.

Which user registries will be used by the local system? Any registries that do not exist will be added to the domain.

☑ Local i5/OS

I5OSP6.STGT.SPC.IHOST.COM

☑ Kerberos

AI.STGT.SPC.IHOST.COM

□ Kerberos user identities are case sensitive

Notes: Setting up an EIM Domain Controller

- When creating a new EIM Domain, the wizard looks for an existing Directory server configuration on the system. If one is not configured, the wizard prompts you with the option to create a basic configured Directory Service. You will need an LDAP Directory user (distinguished name) and password with authority to create the objects for EIM.

- If you setup a new Directory server, be sure to remember the administrator password.

- Typically you dedicate one system in the network as the EIM domain controller. This is the system where you select the option Create and join a new domain. The remaining systems are configured to join the existing domain. You only need to configure the Kerberos registry on the first system that you configure for the EIM domain.

Populating the EIM Domain

- EIM identifier and user associations have to be added to the EIM domain
- Additional systems can be configured to join the existing EIM domain
  - Each system appears in the domain as a registry

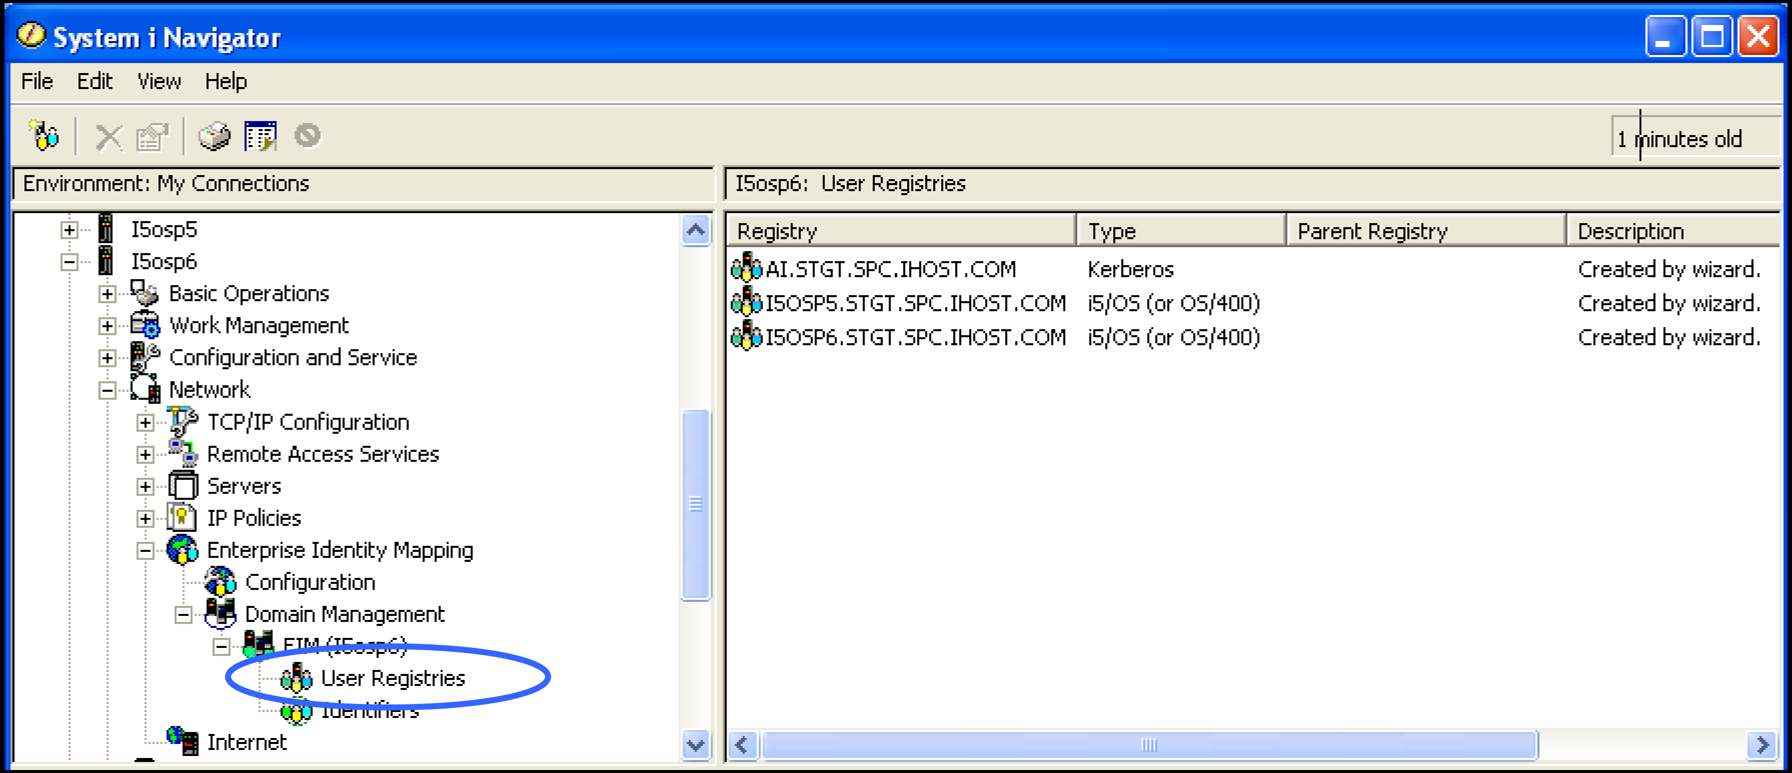

I5osp6: User Registries

Registry	Type	Parent Registry	Description
AI.STGT.SPC.IHOST.COM	Kerberos		Created by wizard.
I5OSP5.STGT.SPC.IHOST.COM	i5/OS (or OS/400)		Created by wizard.
I5OSP6.STGT.SPC.IHOST.COM	i5/OS (or OS/400)		Created by wizard.

Populating the EIM Domain (cont'd)

- User associations can be Source, Target, or Administrative

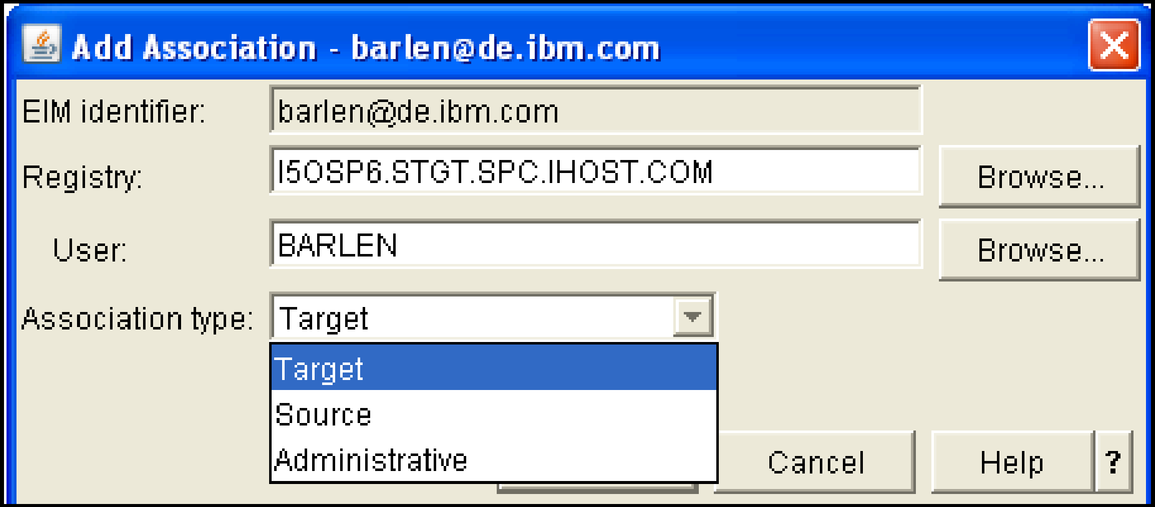

Add Association - barlen@de.ibm.com

EIM identifier: barlen@de.ibm.com

Registry: I5OSP6.STGT.SPC.IHOST.COM

User: BARLEN

Association type: Target

- Source – An association (user) that is presented as an user identifier during an incoming authentication request, such as the Kerberos user principal
- Target – An association that a source user should be mapped to (signed on with) on the target system
- Administrative – For documentation purposes only, cannot be used for EIM lookup operations

*1 = Multiple target support was added with V6R1 → special requirements

Multiple target association support for IBM i

- Prior to V6R1, you could only map one or more EIM source associations to a single target association for a particular IBM i partition
- The new features allows you to have multiple targets on a single target IBM i partition
  - Requires that the IBM i partition has multiple IP interfaces defined
  - If more than one target association exists for a given user, the IP addresses need to be defined as additional lookup information

Associations to a single target registry with IBM i services prior V6R1

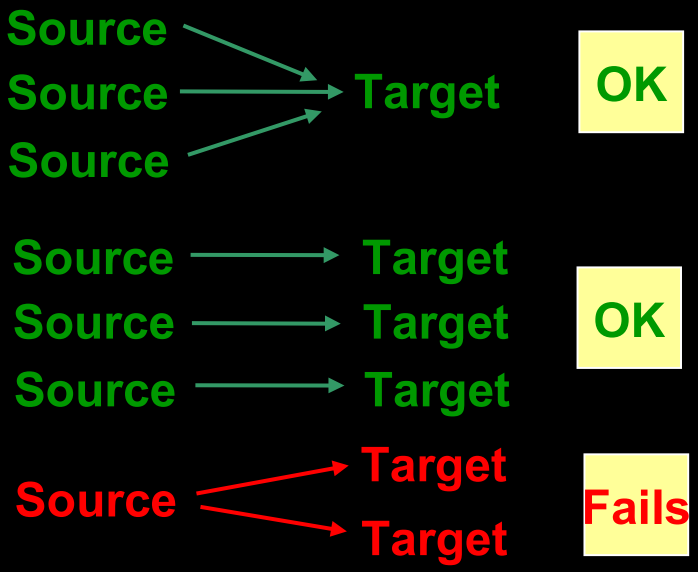

Associations to a single target registry with IBM i services with V6R1 and higher

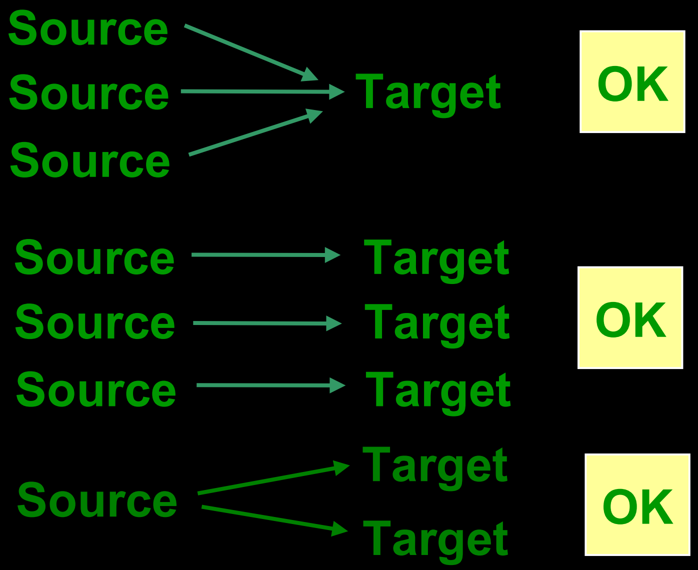

Multiple target association support for IBM i

Current environment

- Single person logs in with different user profiles for different applications

Example:

- Windows user John Dow in Windows domain WIN.DOM.LOCAL signs on via three different IBM i user profiles to system IPROD1 during a day
  - User 1: JOHND (ERP application)
  - User 2: JOHNADM (Administration profile)
  - User 3: DOWJ (Invoice processing)

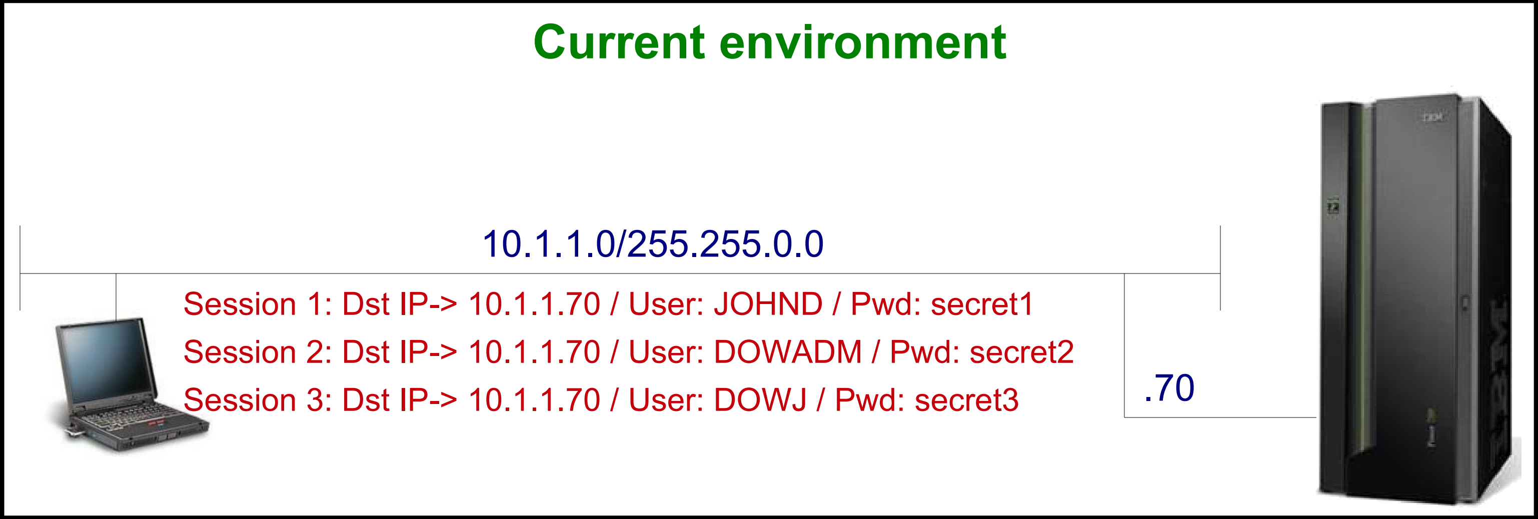

Current environment

Session 1: Dst IP-> 10.1.1.70 / User: JOHND / Pwd: secret1
Session 2: Dst IP-> 10.1.1.70 / User: DOWADM / Pwd: secret2
Session 3: Dst IP-> 10.1.1.70 / User: DOWJ / Pwd: secret3

Multiple target association support for IBM i

Requirements for multiple target user support

- You need a separate IP interface address on IBM i for multiple targets

Example:

- Windows user John Dow in Windows domain WIN.DOM.LOCAL signs on via three different IBM i user profiles to system IPROD1 during a day
  - User 1: JOHND (ERP application)
  - User 2: JOHNADM (Administration profile)
  - User 3: DOWJ (Invoice processing)

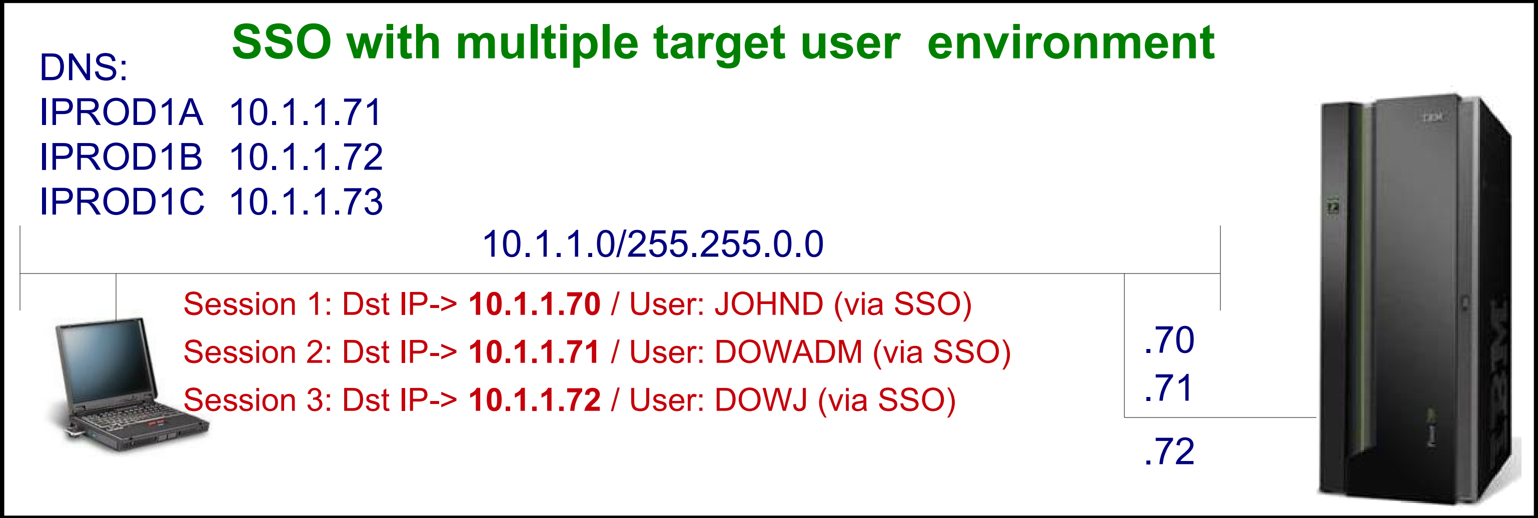

SSO with multiple target user environment

DNS:

IPROD1A	10.1.1.71
IPROD1B	10.1.1.72
IPROD1C	10.1.1.73

10.1.1.0/255.255.0.0

Session 1: Dst IP-> **10.1.1.70** / User: JOHND (via SSO)

Session 2: Dst IP-> **10.1.1.71** / User: DOWADM (via SSO)

Session 3: Dst IP-> **10.1.1.72** / User: DOWJ (via SSO)

Multiple target association support for IBM i

IBM i IP interfaces

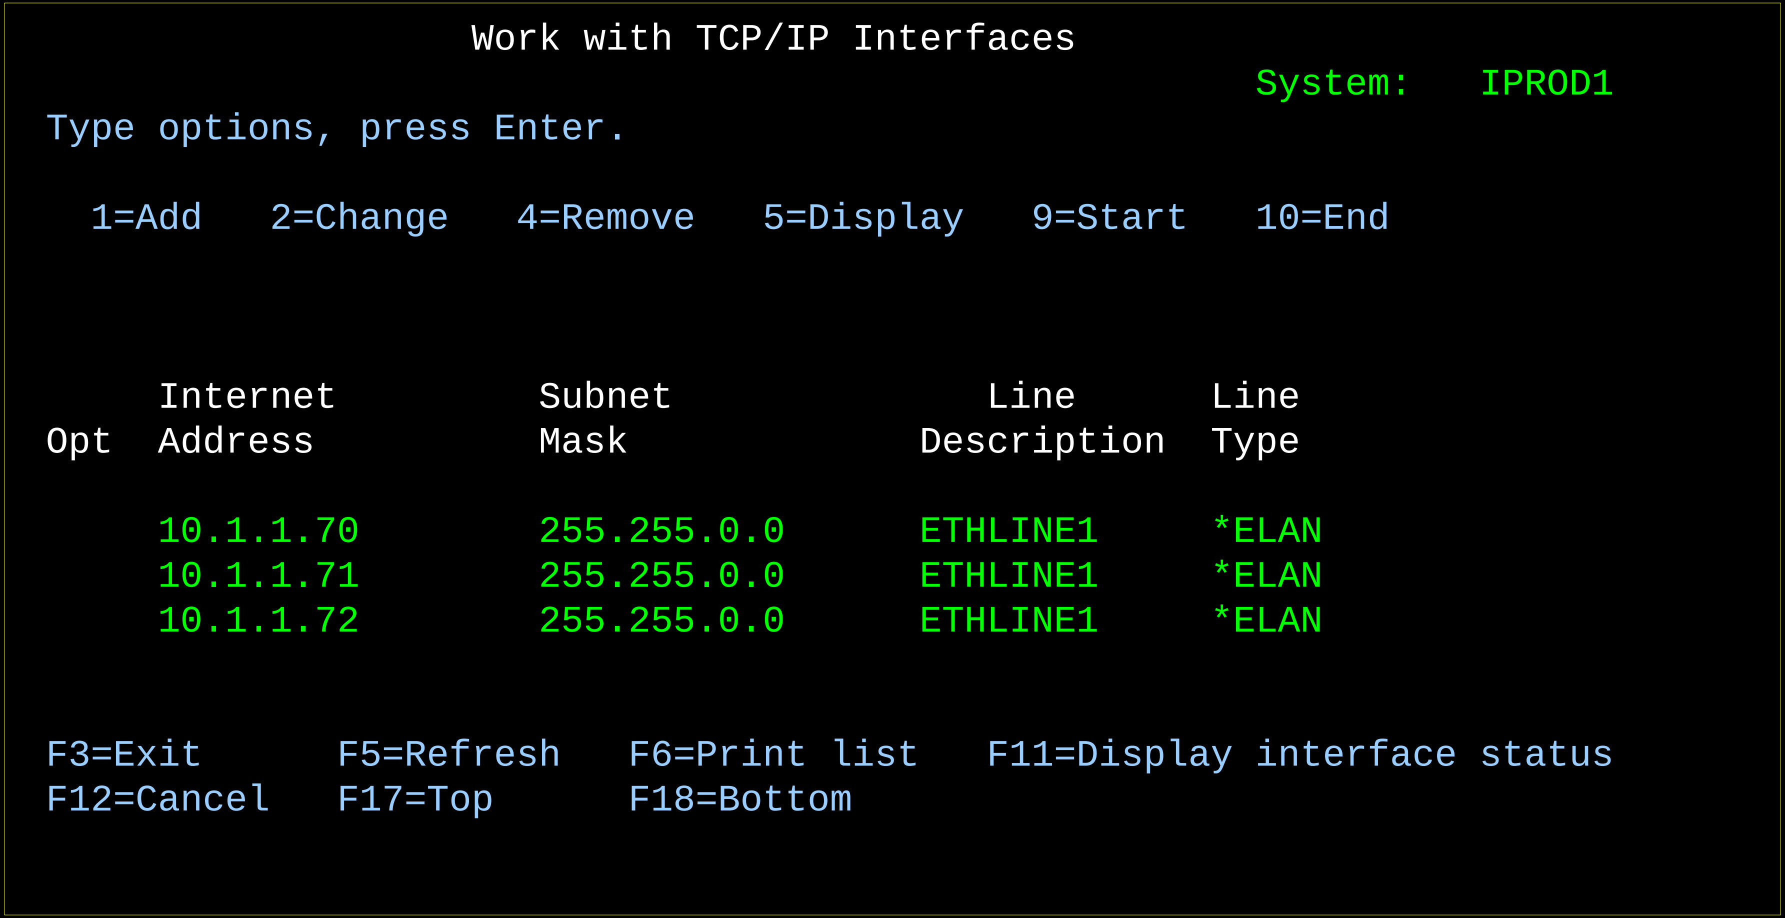

```
                    Work with TCP/IP Interfaces
                                                            System:   IPROD1
 Type options, press Enter.

   1=Add   2=Change   4=Remove   5=Display   9=Start   10=End



      Internet         Subnet              Line       Line
 Opt  Address          Mask             Description  Type

      10.1.1.70        255.255.0.0      ETHLINE1     *ELAN
      10.1.1.71        255.255.0.0      ETHLINE1     *ELAN
      10.1.1.72        255.255.0.0      ETHLINE1     *ELAN


 F3=Exit      F5=Refresh   F6=Print list   F11=Display interface status
 F12=Cancel   F17=Top      F18=Bottom
```

Multiple target association support for IBM i

EIM target assocations

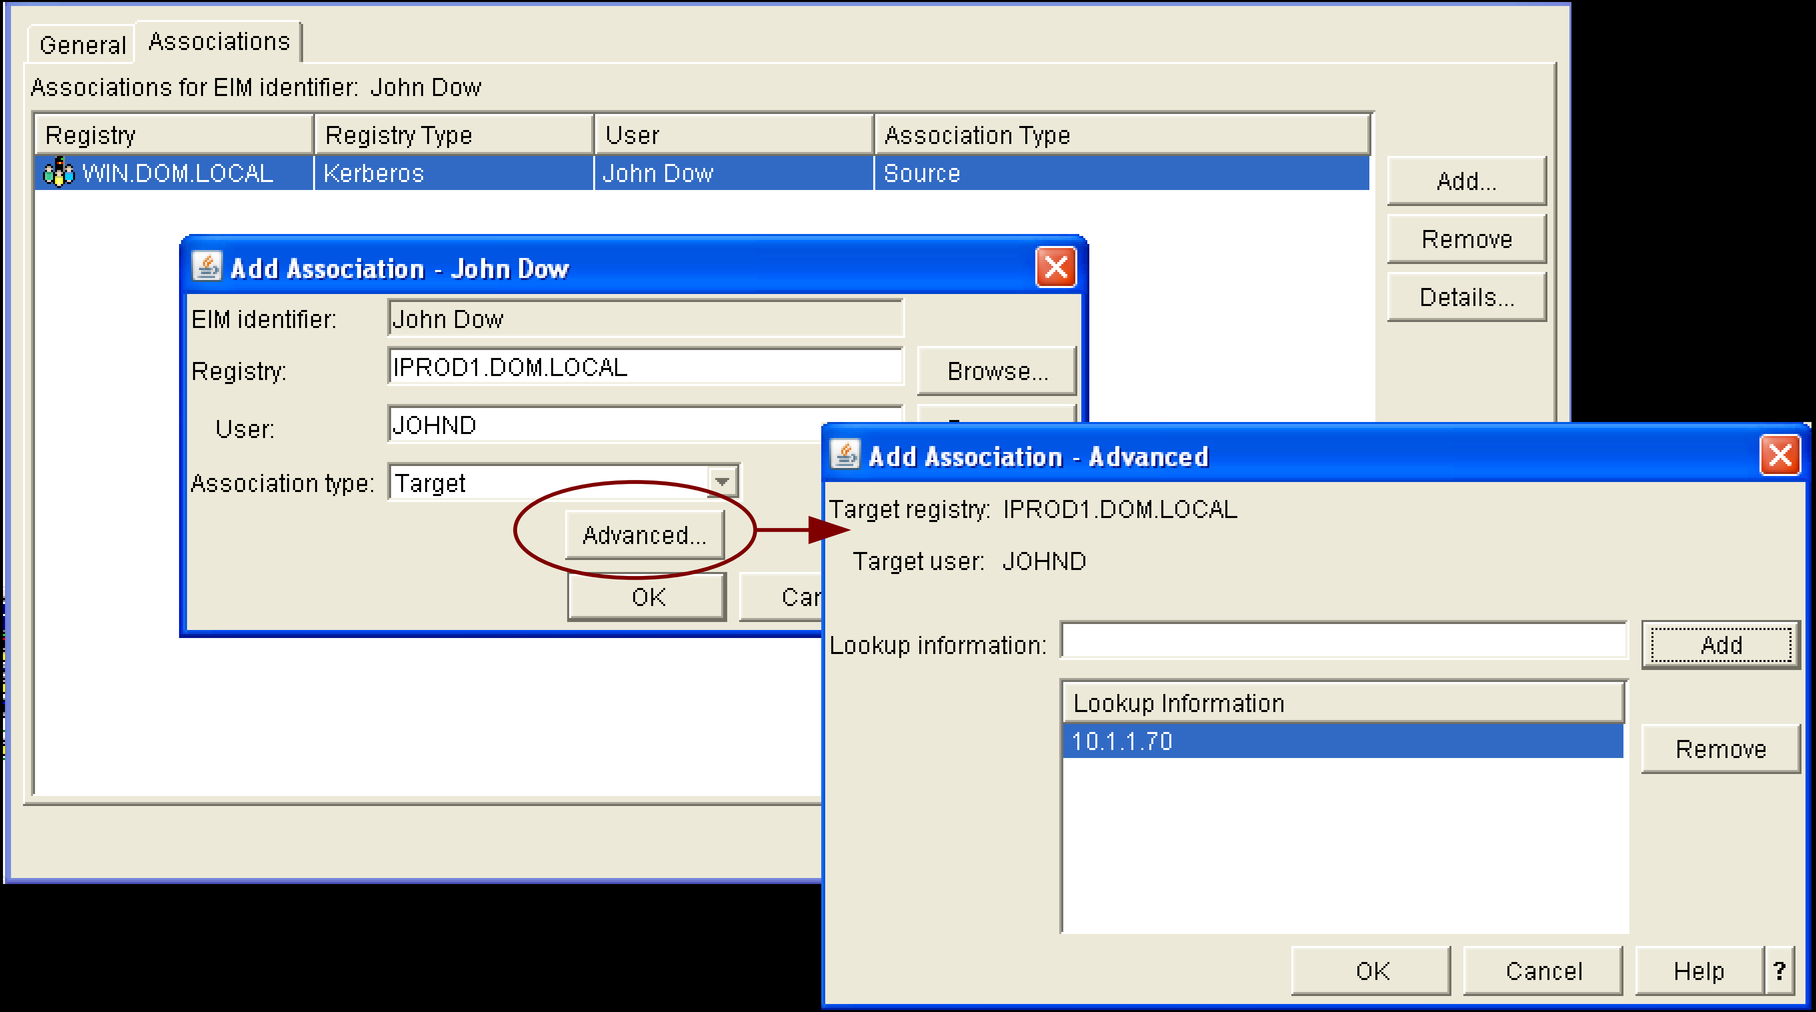

Associations for EIM identifier: John Dow

Registry	Registry Type	User	Association Type
WIN.DOM.LOCAL	Kerberos	John Dow	Source

EIM identifier: John Dow

Registry: IPROD1.DOM.LOCAL

User: JOHND

Association type: Target

Target registry: IPROD1.DOM.LOCAL

Target user: JOHND

Lookup information:

Lookup Information
10.1.1.70

Multiple target association support for IBM i

EIM target assocations

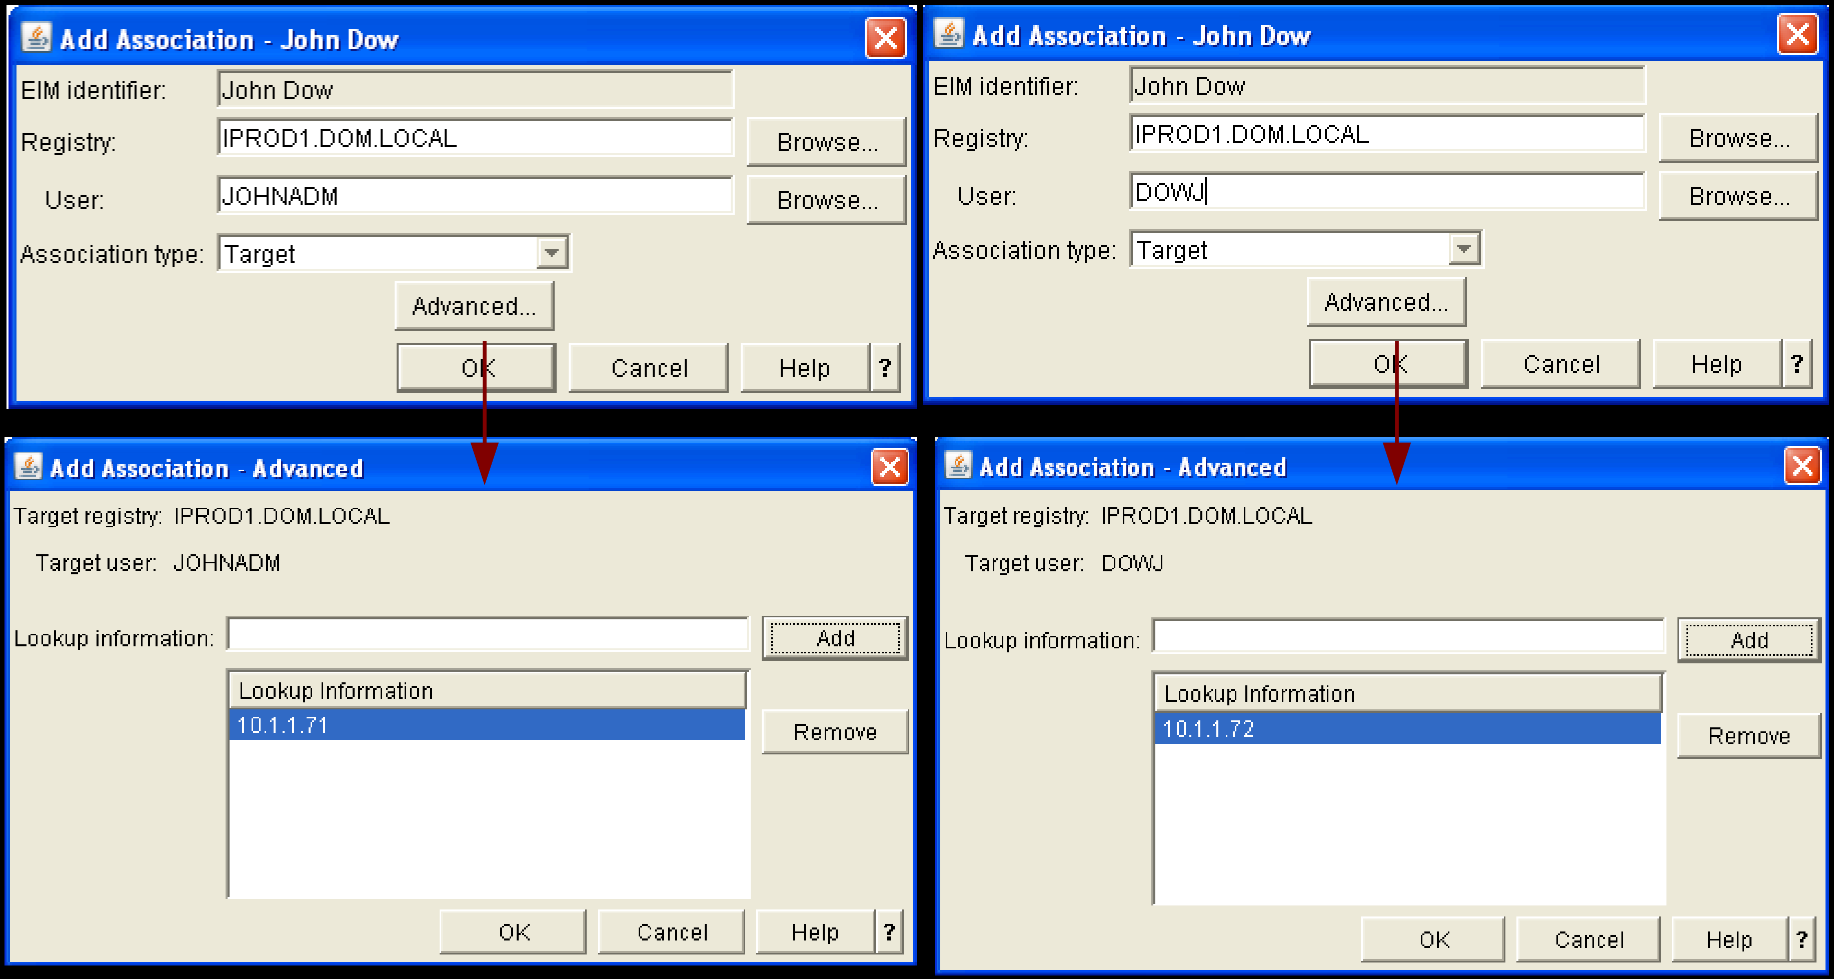

Add Association - John Dow

EIM identifier: John Dow

Registry: IPROD1.DOM.LOCAL

User: JOHNADM

Association type: Target

Add Association - John Dow

EIM identifier: John Dow

Registry: IPROD1.DOM.LOCAL

User: DOWJ

Association type: Target

Add Association - Advanced

Target registry: IPROD1.DOM.LOCAL

Target user: JOHNADM

Lookup information:

Lookup Information
10.1.1.71

Add Association - Advanced

Target registry: IPROD1.DOM.LOCAL

Target user: DOWJ

Lookup information:

Lookup Information
10.1.1.72

Multiple target association support for IBM i

EIM target assocations

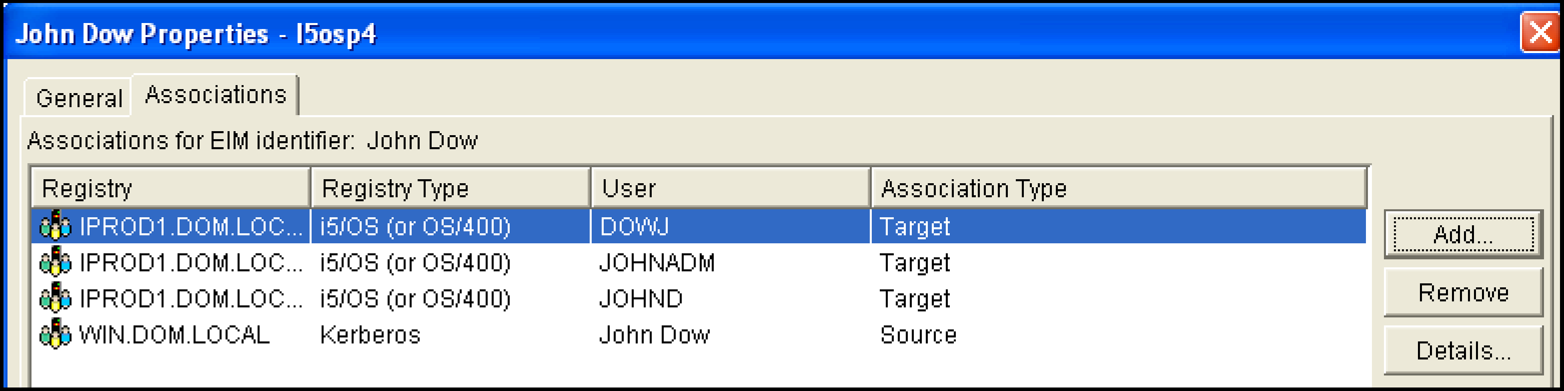

Associations for EIM identifier: John Dow

Registry	Registry Type	User	Association Type
IPROD1.DOM.LOC...	i5/OS (or OS/400)	DOWJ	Target
IPROD1.DOM.LOC...	i5/OS (or OS/400)	JOHNADM	Target
IPROD1.DOM.LOC...	i5/OS (or OS/400)	JOHND	Target
WIN.DOM.LOCAL	Kerberos	John Dow	Source

5250 session configuration

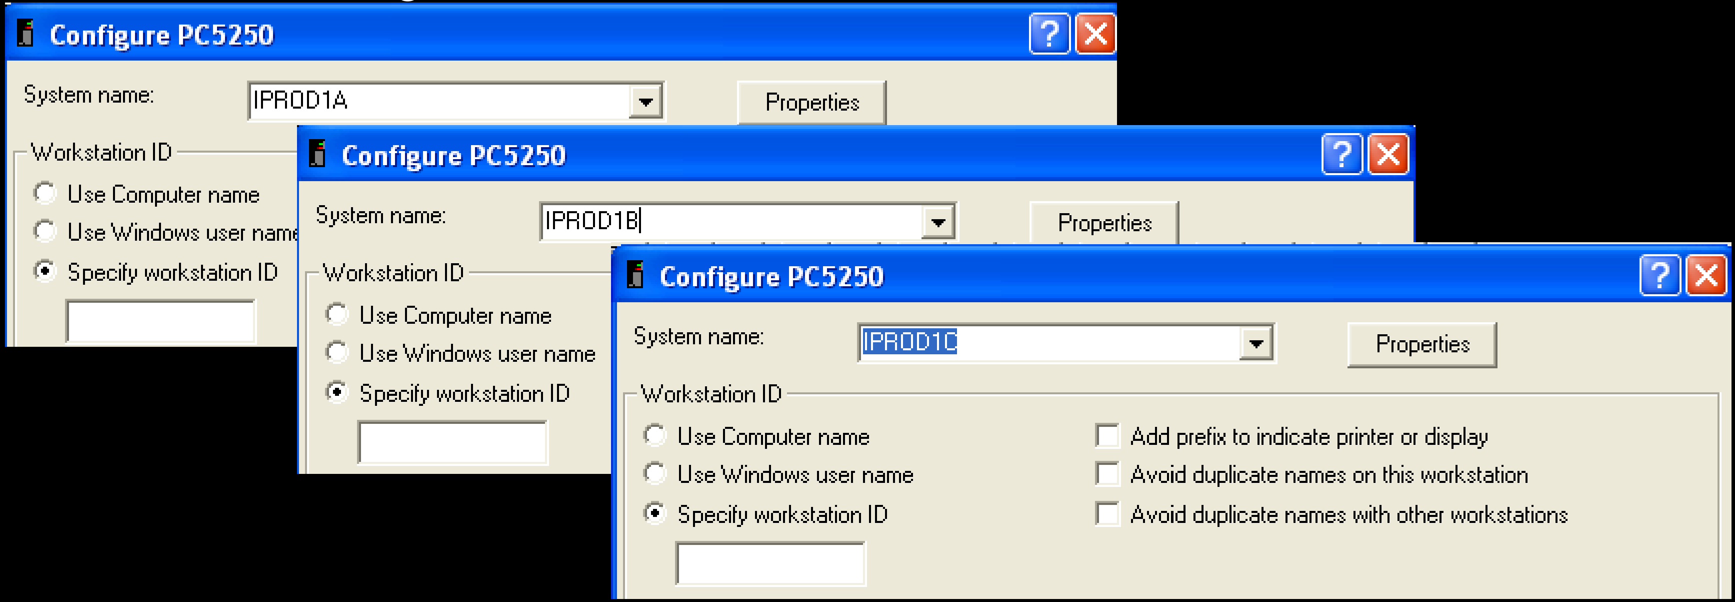

System name: IPROD1A

○ Use Computer name
○ Use Windows user name
● Specify workstation ID

System name: IPROD1B

○ Use Computer name
○ Use Windows user name
● Specify workstation ID

System name: IPROD1C

○ Use Computer name
○ Use Windows user name
● Specify workstation ID

□ Add prefix to indicate printer or display
□ Avoid duplicate names on this workstation
□ Avoid duplicate names with other workstations

Populating the EIM Domain (cont'd)

- EIM identifier and target IBM i user profile associations can either be added via the IBM System i Navigator EIM interface or through the create and change user profile command
- Source associations need to be added via the System i Navigator interface or via Windows Integration
- Third-party EIM management products or IBM Systems Lab Services & Training tools can also manager EIM data

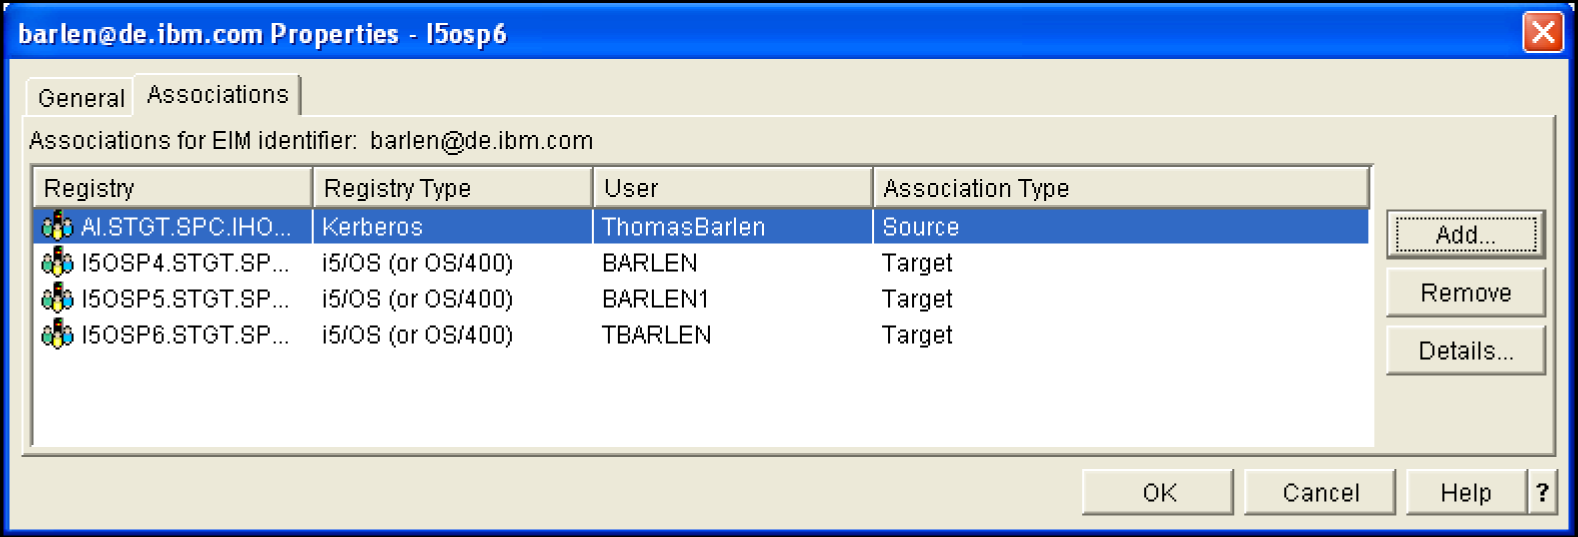

barlen@de.ibm.com Properties - I5osp6

Associations for EIM identifier: barlen@de.ibm.com

Registry	Registry Type	User	Association Type
AI.STGT.SPC.IHO...	Kerberos	ThomasBarlen	Source
I5OSP4.STGT.SP...	i5/OS (or OS/400)	BARLEN	Target
I5OSP5.STGT.SP...	i5/OS (or OS/400)	BARLEN1	Target
I5OSP6.STGT.SP...	i5/OS (or OS/400)	TBARLEN	Target

Alternative methods for EIM population

- Standard approach with IBM System i Navigator is good for testing and manual changes
  - Cumbersome to defines hundreds of identifiers and associations
- Which method is best depends on whether the Windows user name is the same as the IBM i user profile name
  - If the names are the same, a small program using Java classes or ILE APIs can be used to automatically populate EIM identifier and associations
  - If the names are different, a program can still be used, but the user mappings have to be defined prior to running the program (i.e. in a CSV file)

Alternative methods for EIM population

- Tools can be written to automatically create EIM identifier and associations
  - Example 1: An exit program for the QIBM_QSY_CRT_PROFILE and QIBM_QSY_DLT_PROFILE exit points can be written to create or delete EIM data
  - Example 2: A program can be written that listens for Windows Active Directory changes and depending on a group membership create IBM i user profiles and EIM information
    - Such as tool has been written by IBM Systems Lab Services & Training

Populating the EIM Domain (cont'd)

- Example of automatic population of EIM data via AD event notification
  - Tool registers itself with Microsoft Active Directory
  - Too is notified for every change in AD user accounts
  - Depending on group memberships IBM i user profiles get created/changed/deleted/enabled/disabled and EIM information created

```
<mapping>
    <windows>CN=Administrators,CN=Users,DC=windom,DC=ibm,DC=com</windows>
    <iseries dftcrtpassword="hfkj4m" usrcls="*SECOFR">sys11.ibm.com</iseries>
</mapping>

<mapping>
        <windows>CN=Clerks,CN=Users,DC=windom,DC=ibm,DC=com</windows>
        <iseries dftcrtpassword="*NONE" usrcls="*PGMR"
        GRPPRF="SSUGRP">sys2.ibm.com</iseries>
</mapping>
```

Notes: Populating the EIM Domain

- For EIM mapping lookup operations to work, the EIM domain controller must be filled with EIM identifier and corresponding associations. An EIM identifier uniquely identifies a person or entity within an EIM domain. The name of the identifier must be unique.

- The **Add Association** dialog allows you to define an individual association between the selected EIM identifier and a specific user identity in a specific user registry. An EIM registry definition for the user registry must exist prior to defining an association between a user identity in the user registry and the selected EIM identifier.

- With each user association, you also need to specify the type of association. The various types are as follows:
  - **Source**: A source association allows the user identity to be used as the source in an EIM lookup operation to find a different user identity that is associated with the same EIM identifier. When a user identity is used for authentication, that user identity should have a source association with an EIM identifier. For example, you might create a source association for a Kerberos principal because this form of user identity is used for authentication.
  - **Target**: A target association allows the user identity to be returned as the result of an EIM lookup operation. User identities that represent end users normally need a target association only. When a user identity is used for authorization rather than for authentication, that user identity should have a target association with an EIM identifier. For example, you might create a target association for an IBM i user profile because this form of user identity determines what resources and privileges the user has on a specific IBM i partition.

Notes: Populating the EIM Domain (cont'd)

  - **Administrative**: An administrative association for an EIM identifier is typically used to show that the person or entity represented by the EIM identifier owns a user identity that requires special considerations for a specified system. This type of association can be used, for example, with highly sensitive user registries. Due to the special nature of administrative associations, this type of association can not participate in EIM mapping lookup operations. Consequently, an EIM lookup operation that supplies a source user identity with an administrative association returns no results. Similarly, a user identity with an administrative association is never returned as the result of an EIM lookup operation.

- Note: If one identifier has multiple "targets" or one "source" points to multiple identifiers, it is up to the requesting registry to handle the returned multiple entries. IBM i system applications normally reject authentication requests where one source association is mapped to multiple target associations on the same system. With V6R1 a feature was added that allows one source to be mapped to multiple targets on a single system. Requirements for this function to work are multiple IP interfaces and the use of the additional lookup information for a target EIM association.

Enabling Kerberos on the Client Side

- The last step in enabling SSO with Kerberos and EIM is to activate Kerberos for the client applications
- EIM is typically not being used on the client side
- Activation of Kerberos is application dependent
- NetServer (SMB client) is automatically active when a service principal exists
- Browser settings have to be changed to activate Kerberos with the Apache server
- IBM System i Access for Windows including the ODBC driver is activated via System i Navigator connection settings

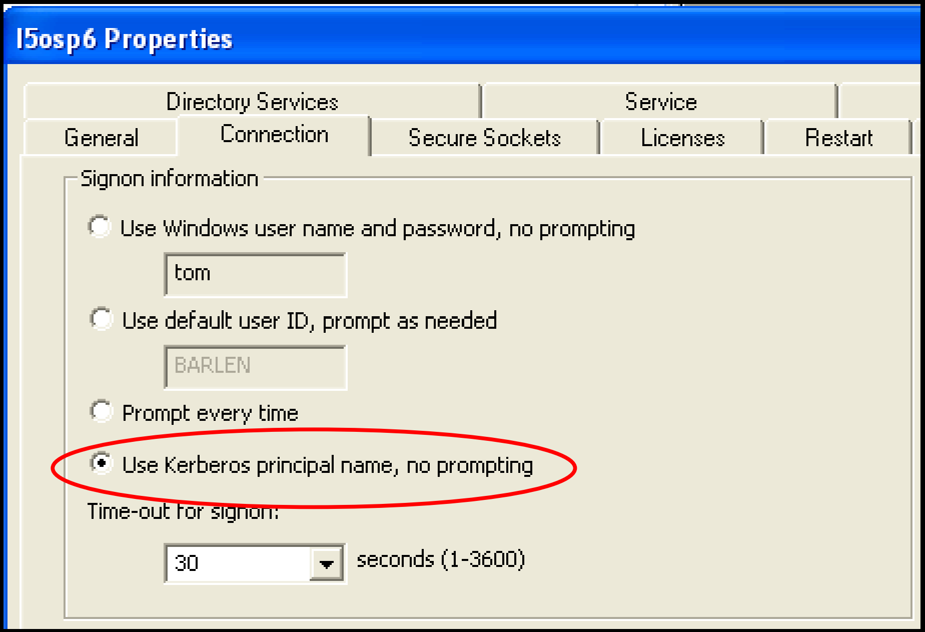

I5osp6 Properties

Signon information

○ Use Windows user name and password, no prompting

tom

○ Use default user ID, prompt as needed

BARLEN

○ Prompt every time

● Use Kerberos principal name, no prompting

Time-out for signon:

30 seconds (1-3600)

Notes: Enabling Kerberos on the Client Side

- Assuming EIM and Kerberos works on the server and in the network, you should now be able to use the Kerberos authentication method with System i Navigator. The IBM i host servers, in turn, use the EIM functionality to map the incoming (source) identifier to a target user profile.

- On the properties field for the IBM i connection, select **Use Kerberos System name, no prompting**. Restart the Navigator. The sign-on process should be quick and seamless.

- Note: The Kerberos environment allows for caching of tickets and session keys. If the client still has a valid ticket/session key in its cache, it attempts to reconnect without requesting a new service ticket from the TGS. To renew the actual Kerberos information for the Windows User, you have to log off the computer.

- One exception where no change is required on the client side to activate Kerberos, is the SMB client. Whenever a Windows 2000, XP, Vista, Windows 7, or Windows 8 client tries to map a share to a local drive letter, the Windows SMB client tries to obtain a service ticket for the NetServer. If one exists in the KDC and is returned by the KDC, the client expects the NetServer to support Kerberos as well. When Kerberos is not working on the IBM i partition or EIM associations are missing, the authentication request fails. Note that Windows does not fall back to user ID/password authentication in that case.

Thank You !

How to involve IBM Systems Lab Services & Training

- Visit out Web site at http://www.ibm.com/systems/services/labservices/

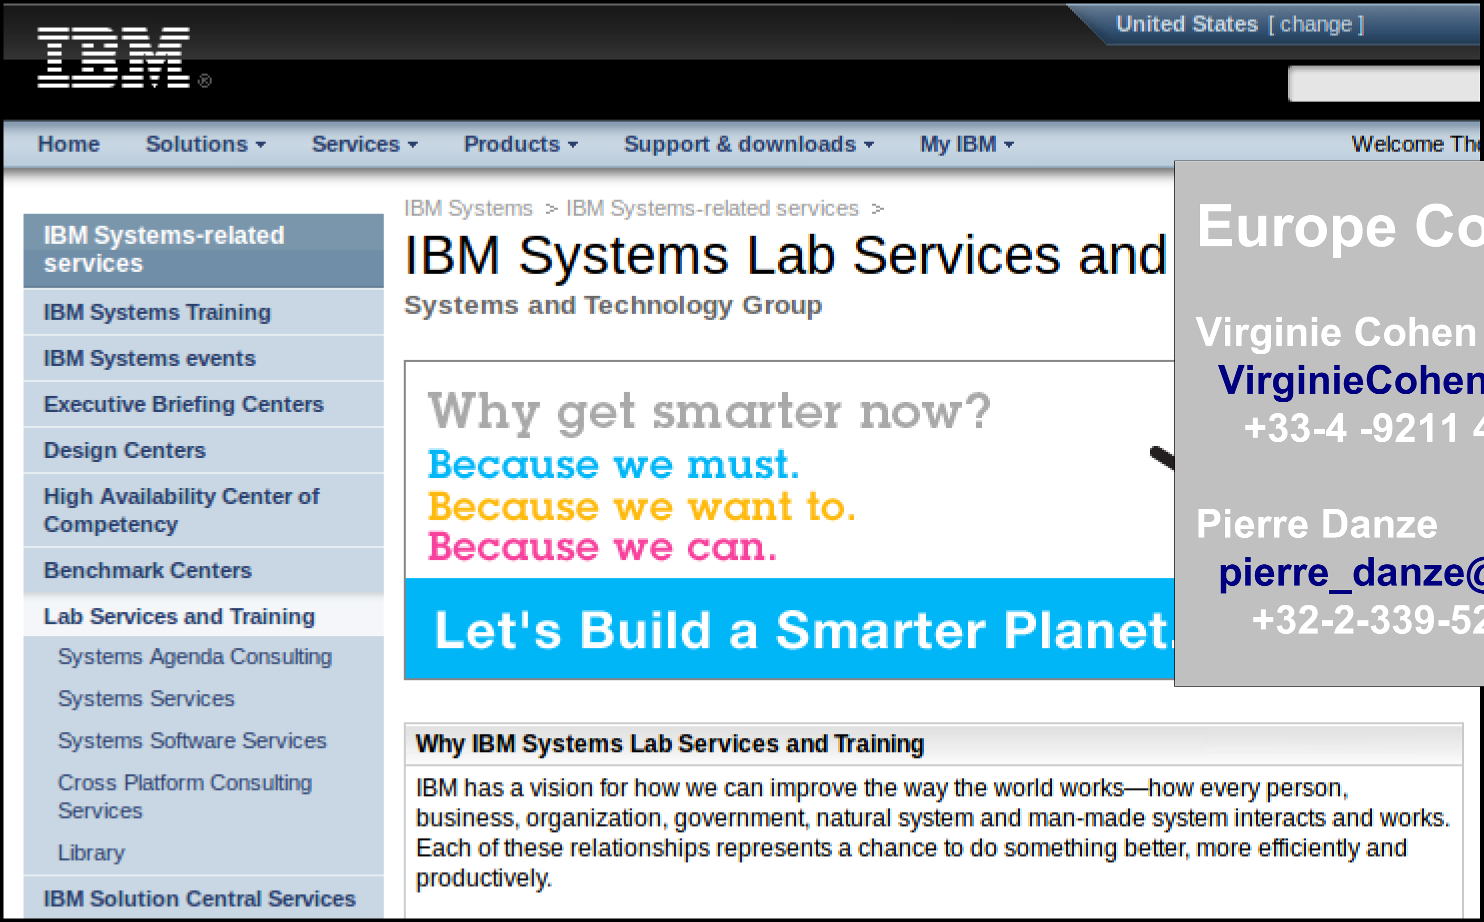

United States [change]

Home | Solutions | Services | Products | Support & downloads | My IBM | Welcome Th

IBM Systems > IBM Systems-related services >

IBM Systems-related services
- IBM Systems Training
- IBM Systems events
- Executive Briefing Centers
- Design Centers
- High Availability Center of Competency
- Benchmark Centers
- Lab Services and Training
  - Systems Agenda Consulting
  - Systems Services
  - Systems Software Services
  - Cross Platform Consulting Services
  - Library
- IBM Solution Central Services

IBM Systems Lab Services and

Systems and Technology Group

Why get smarter now?
Because we must.
Because we want to.
Because we can.

Let's Build a Smarter Planet.

Why IBM Systems Lab Services and Training

IBM has a vision for how we can improve the way the world works—how every person, business, organization, government, natural system and man-made system interacts and works. Each of these relationships represents a chance to do something better, more efficiently and productively.

Europe Contacts

Virginie Cohen
VirginieCohen@fr.ibm.com
+33-4 -9211 41 33

Pierre Danze
pierre_danze@be.ibm.com
+32-2-339-5241

Advanced Topics

Windows SMB client and IBM i NetServer

- All IBM i related client applications need to be configured to use Kerberos authentication rather than user/password
- EXCEPT ------> The Microsoft SMB client

SMB client behavior when mapping a drive from IBM i NetServer

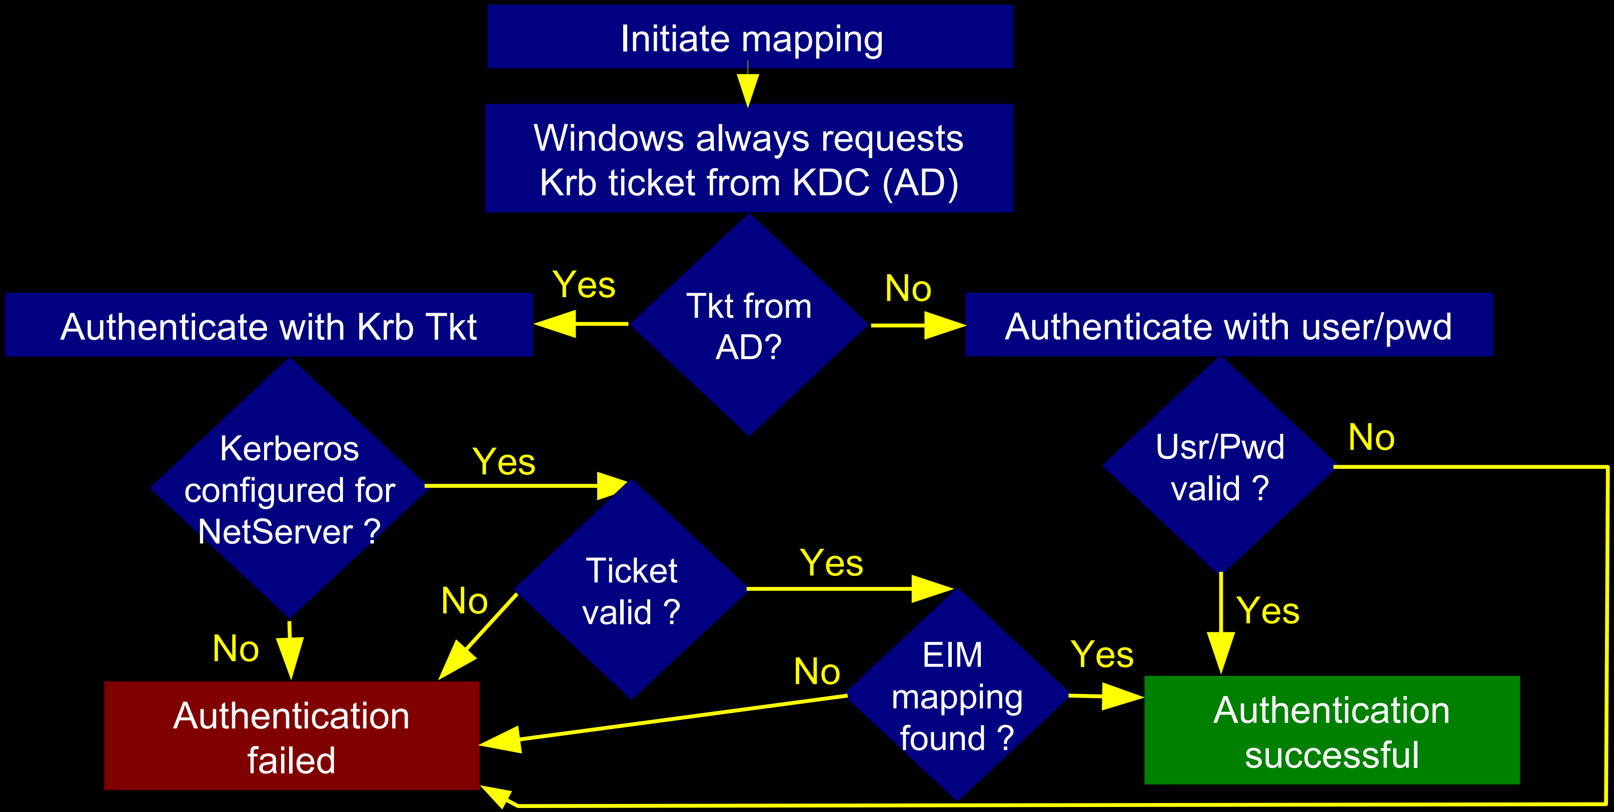
Initiate mapping
Windows always requests Krb ticket from KDC (AD)
Tkt from AD?
Yes
No
Authenticate with Krb Tkt
Authenticate with user/pwd
Kerberos configured for NetServer ?
Yes
No
Ticket valid ?
Yes
No
EIM mapping found ?
Yes
No
Usr/Pwd valid ?
Yes
No
Authentication failed
Authentication successful

Windows SMB client and IBM i NetServer

- What ticket does Windows request from AD?
  - Windows 2000 requests HOST/hostname@WINDOWS.DOMAIN
  - Windows XP and higher request cifs/hostname@WINDOWS.DOMAIN
- What does hostname refer to?

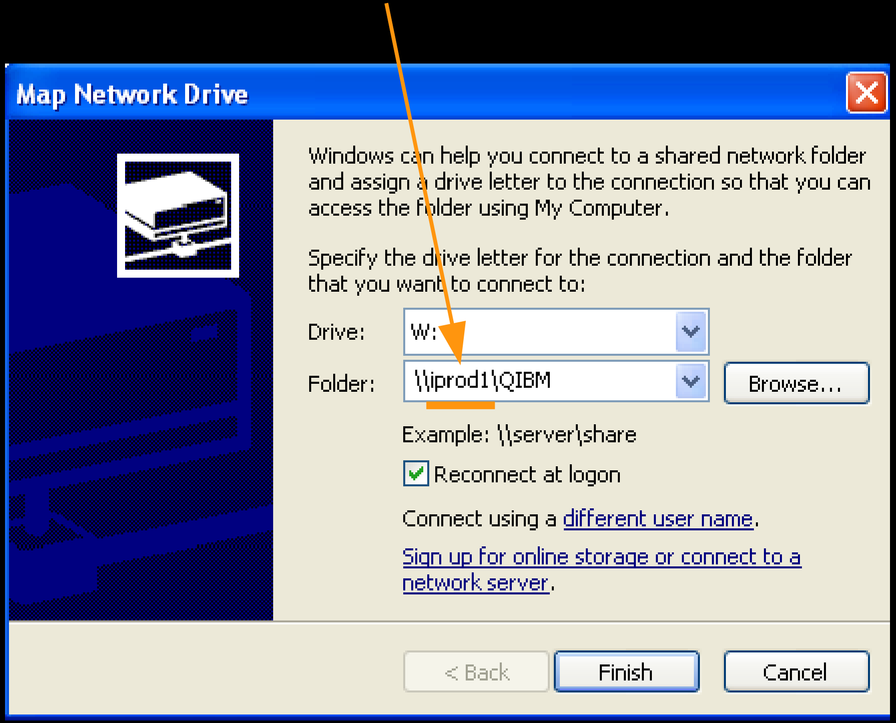

Map Network Drive

Windows can help you connect to a shared network folder and assign a drive letter to the connection so that you can access the folder using My Computer.

Specify the drive letter for the connection and the folder that you want to connect to:

Drive: W:

Folder: \\iprod1\QIBM

Example: \\server\share

☑ Reconnect at logon

Connect using a different user name.

Sign up for online storage or connect to a network server.

Typical hostnames include:

iprod1
iprod1.dom.local
qiprod1
10.1.1.70

Windows SMB client and IBM i NetServer

- What did we learn so far?
  - Windows clients always request a Kerberos ticket from AD (KDC)
  - If the clients gets a ticket, it sends it to the NetServer
  - NetServer verifies the ticket and performs an EIM lookup to obtain IBM i user profile name
  - If EIM lookup fails, the mapping process fails
- What can you do to test SSO for the NetServer or enable it only for a subset of workstations?
  - Define a service principle name (SPN) that is currently not in use and test with this name (assuming FQDN hostname not used today)

    cifs/iprod1.dom.local@WINDOWS.DOMAIN
    ~~cifs/iprod1@WINDOWS.DOMAIN~~
    ~~cifs/qiprod1@WINDOWS.DOMAIN~~
    ~~cifs/10.1.1.70@WINDOWS.DOMAIN~~
  - Once all testing has been completed, define all EIM mappings and then add the remaining SPNs to the KDC (AD)

    cifs/iprod1.dom.local@WINDOWS.DOMAIN
    cifs/iprod1@WINDOWS.DOMAIN
    cifs/qiprod1@WINDOWS.DOMAIN
    cifs/10.1.1.70@WINDOWS.DOMAIN

Mass roll-out of SSO

- Enabling Kerberos on the Client Side
  - The last step in enabling SSO with Kerberos and EIM is to activate Kerberos for the client applications
  - EIM is typically not being used on the client side
  - Activation of Kerberos is application dependent
  - NetServer (SMB client) is automatically active when a service principal exists
  - Browser settings have to be changed to activate Kerberos with the Apache server
  - System i Access for Windows including the System i ODBC driver is activated via System i Navigator connection settings, but can be overridden in the ODBC datasource or PC5250 workstation profile

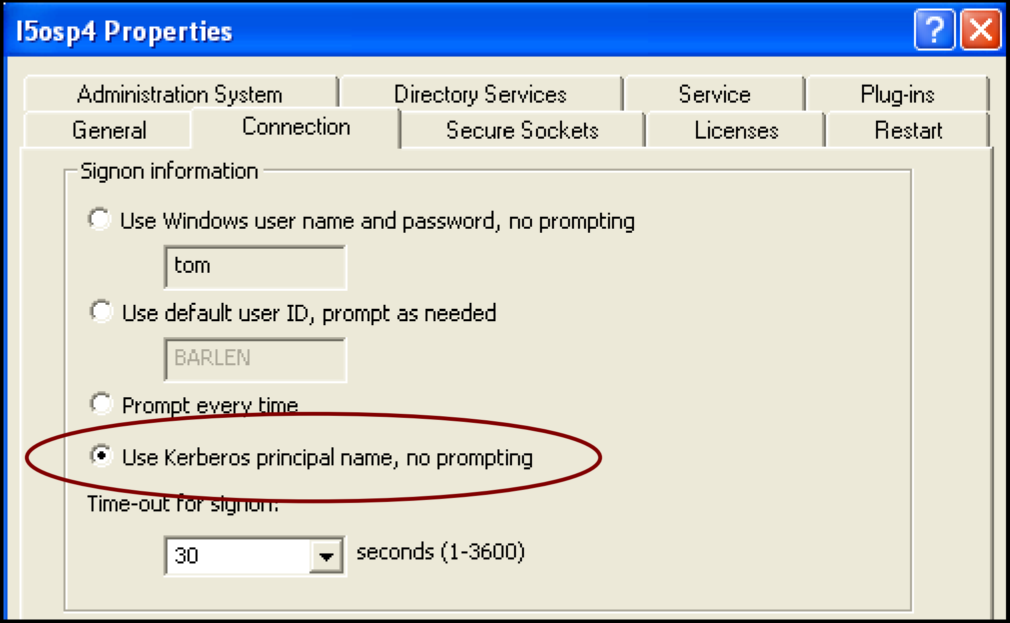
I5osp4 Properties

Signon information

○ Use Windows user name and password, no prompting

tom

○ Use default user ID, prompt as needed

BARLEN

○ Prompt every time

● Use Kerberos principal name, no prompting

Time-out for signon:

30 seconds (1-3600)

Mass roll-out of SSO

- SSO configuration settings are stored in various places
  - System i Navigator provides a central switch to turn SSO on or off for System I Navigator, PC5250, ODBC
    - Each application can override the Navigator settings
  - System i Navigator stores the configuration setting in the Windows registry

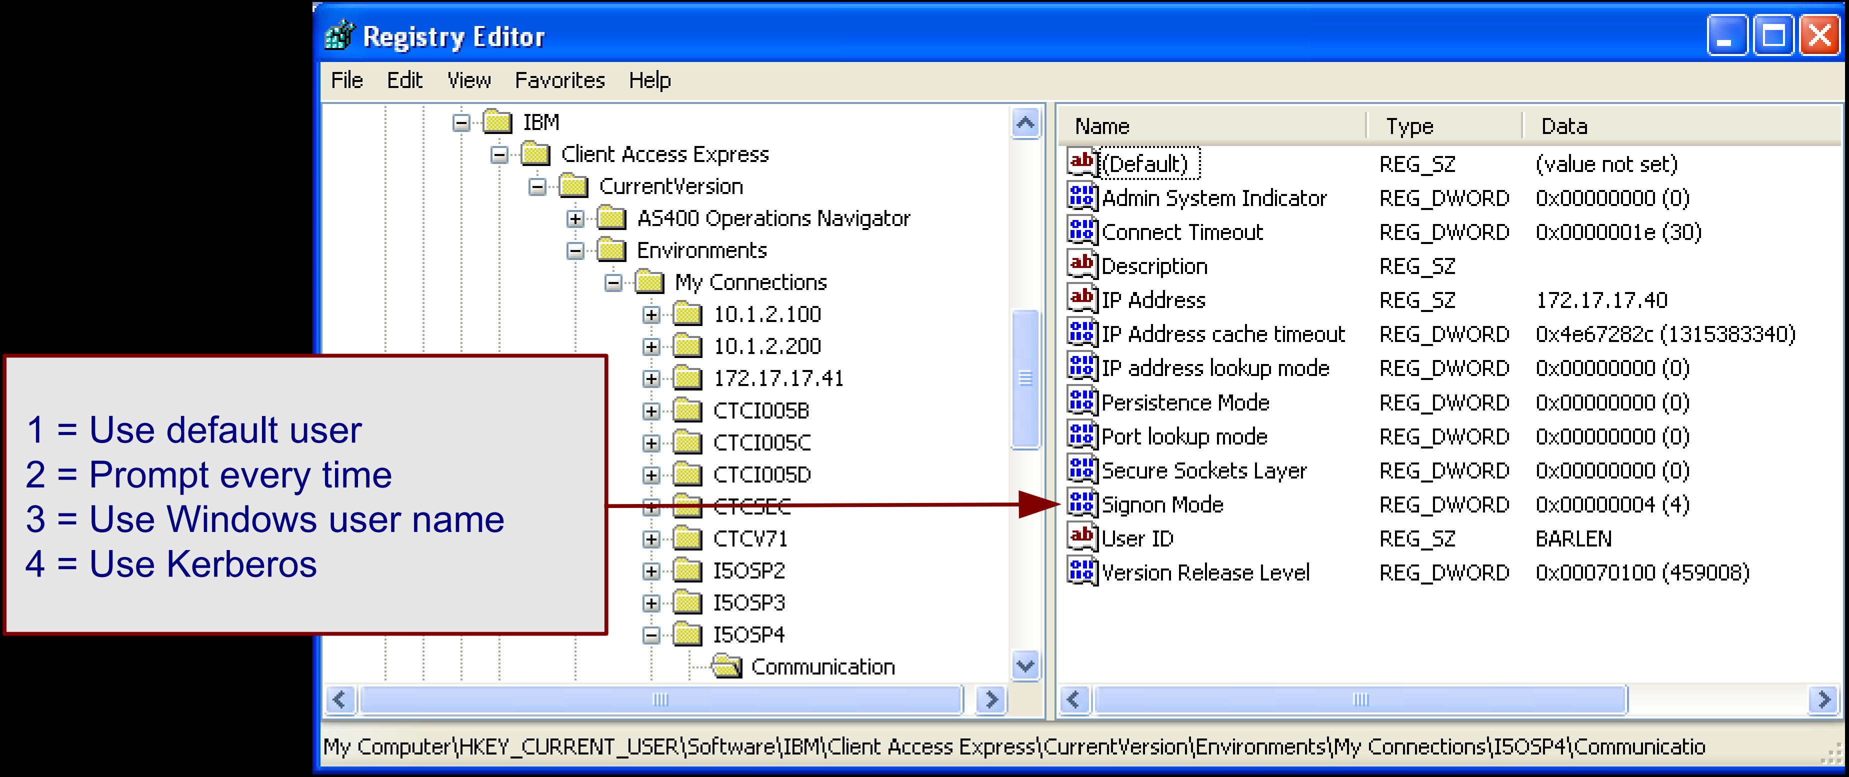

Name	Type	Data
(Default)	REG_SZ	(value not set)
Admin System Indicator	REG_DWORD	0x00000000 (0)
Connect Timeout	REG_DWORD	0x0000001e (30)
Description	REG_SZ	
IP Address	REG_SZ	172.17.17.40
IP Address cache timeout	REG_DWORD	0x4e67282c (1315383340)
IP address lookup mode	REG_DWORD	0x00000000 (0)
Persistence Mode	REG_DWORD	0x00000000 (0)
Port lookup mode	REG_DWORD	0x00000000 (0)
Secure Sockets Layer	REG_DWORD	0x00000000 (0)
Signon Mode	REG_DWORD	0x00000004 (4)
User ID	REG_SZ	BARLEN
Version Release Level	REG_DWORD	0x00070100 (459008)

My Computer\HKEY_CURRENT_USER\Software\IBM\Client Access Express\CurrentVersion\Environments\My Connections\I5OSP4\Communicatio

1 = Use default user
2 = Prompt every time
3 = Use Windows user name
4 = Use Kerberos

Mass roll-out of SSO

- System i Navigator registry setting can be exported to .reg file and used for automatic import via Login scripts

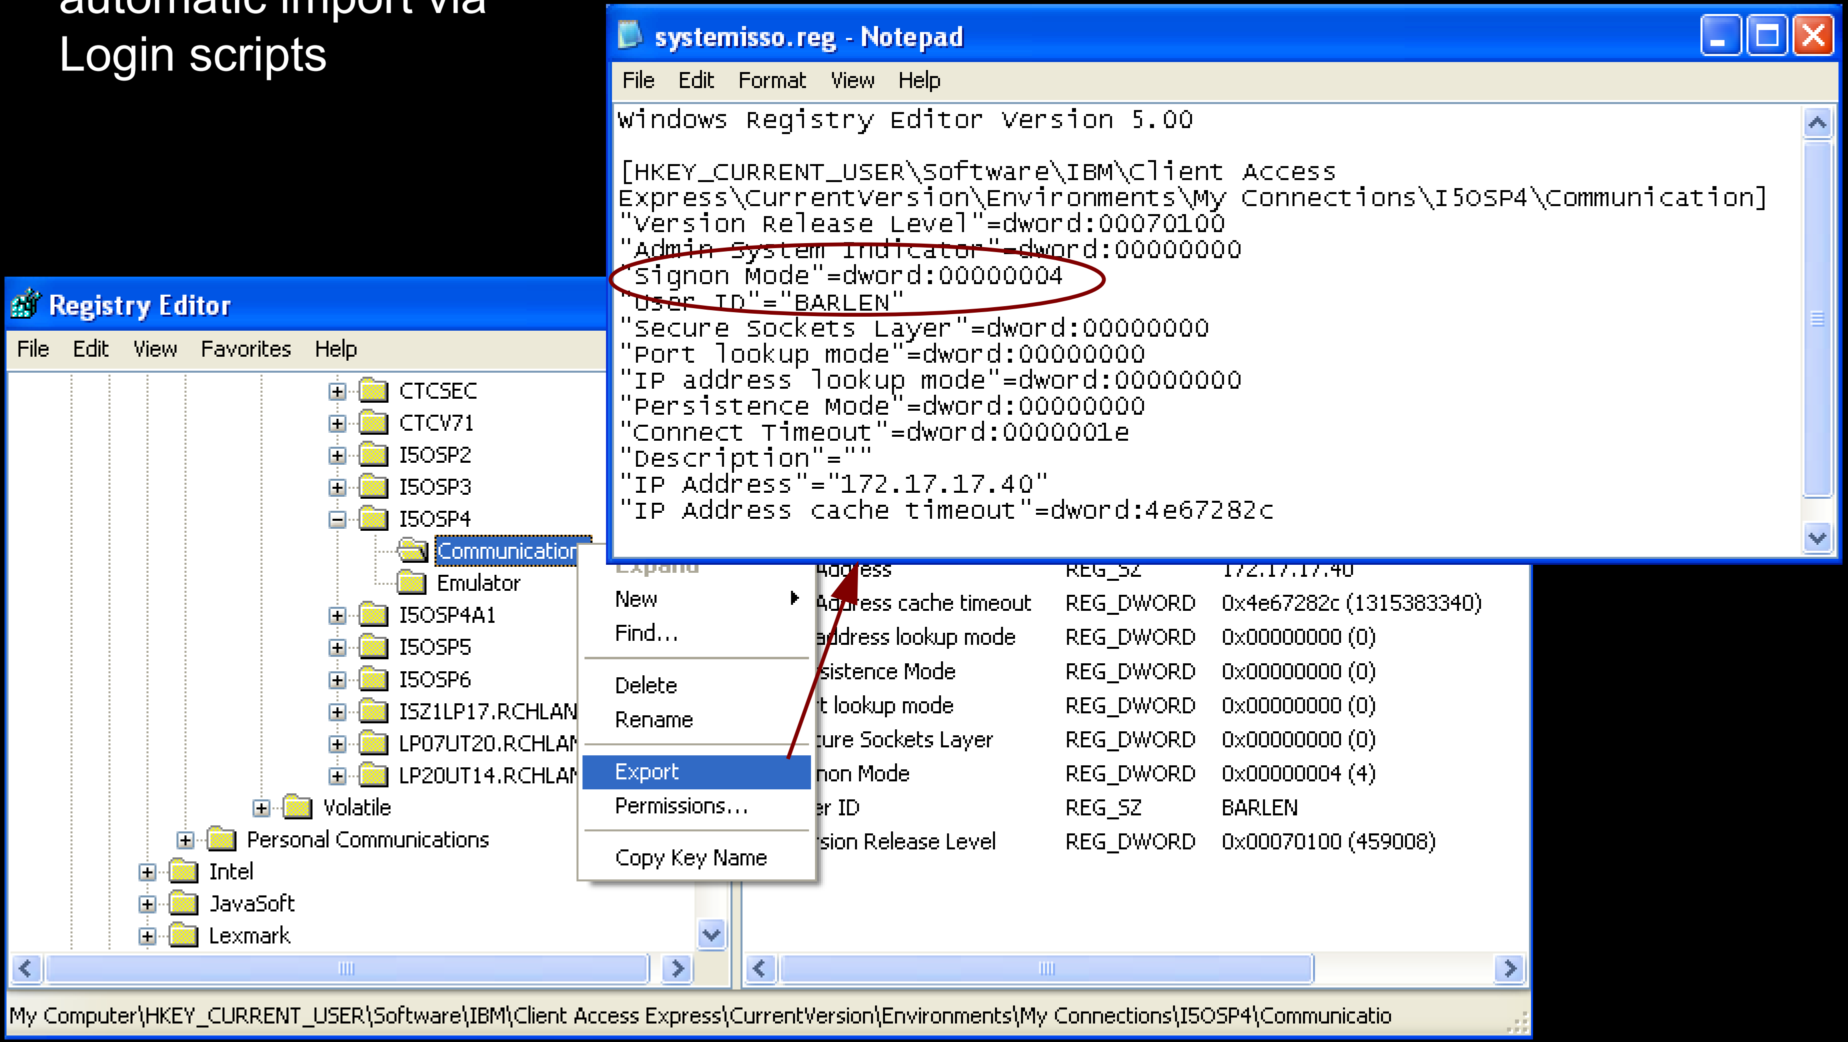

```
Windows Registry Editor Version 5.00

[HKEY_CURRENT_USER\Software\IBM\Client Access
Express\CurrentVersion\Environments\My Connections\I5OSP4\Communication]
"Version Release Level"=dword:00070100
"Admin System Indicator"=dword:00000000
"Signon Mode"=dword:00000004
"User ID"="BARLEN"
"Secure Sockets Layer"=dword:00000000
"Port lookup mode"=dword:00000000
"IP address lookup mode"=dword:00000000
"Persistence Mode"=dword:00000000
"Connect Timeout"=dword:0000001e
"Description"=""
"IP Address"="172.17.17.40"
"IP Address cache timeout"=dword:4e67282c
```

Address	REG_SZ	172.17.17.40
Address cache timeout	REG_DWORD	0x4e67282c (1315383340)
address lookup mode	REG_DWORD	0x00000000 (0)
sistence Mode	REG_DWORD	0x00000000 (0)
t lookup mode	REG_DWORD	0x00000000 (0)
ure Sockets Layer	REG_DWORD	0x00000000 (0)
non Mode	REG_DWORD	0x00000004 (4)
er ID	REG_SZ	BARLEN
sion Release Level	REG_DWORD	0x00070100 (459008)

My Computer\HKEY_CURRENT_USER\Software\IBM\Client Access Express\CurrentVersion\Environments\My Connections\I5OSP4\Communicatio

Mass roll-out of SSO

- As an alternative to the registry approach for System i Navigator you can also use the IBM System i Access for Windows cwbenv command

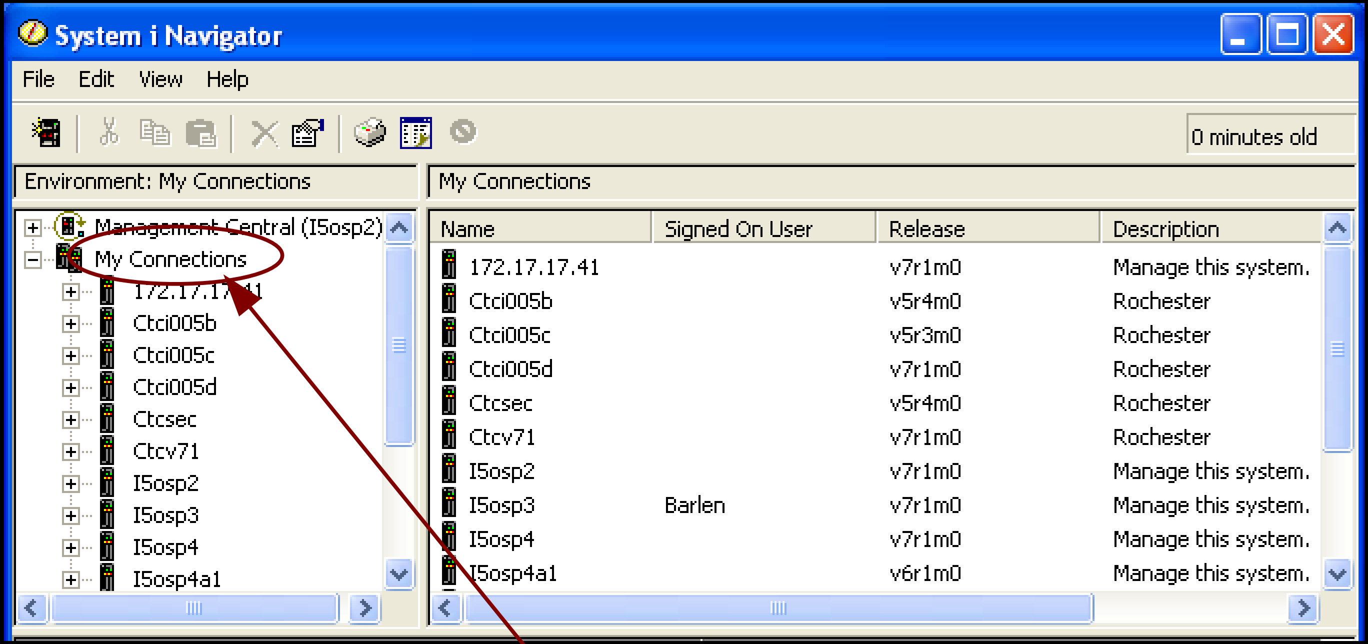

Environment: My Connections

- Management Central (I5osp2)
- My Connections
  - 172.17.17.41
  - Ctci005b
  - Ctci005c
  - Ctci005d
  - Ctcsec
  - Ctcv71
  - I5osp2
  - I5osp3
  - I5osp4
  - I5osp4a1

My Connections

Name	Signed On User	Release	Description
172.17.17.41		v7r1m0	Manage this system.
Ctci005b		v5r4m0	Rochester
Ctci005c		v5r3m0	Rochester
Ctci005d		v7r1m0	Rochester
Ctcsec		v5r4m0	Rochester
Ctcv71		v7r1m0	Rochester
I5osp2		v7r1m0	Manage this system.
I5osp3	Barlen	v7r1m0	Manage this system.
I5osp4		v7r1m0	Manage this system.
I5osp4a1		v6r1m0	Manage this system.

- Export a connection environment including its settings (includes all connections)
  - `cwbenv /E "My connections" ibmienv.fil`
- Import a connection environment
  - `cwbenv /I /O ibmienv.fil`
  - /O overrides existing connections with new settings

Mass roll-out of SSO - PC5250

- Manually specifying SSO (Kerberos) for a PC5250 session
  - PC5250 sessions are stored in workstation profiles (*.WS files)
  - Certain settings are stored in *.CAE files, i.e. whether SSO is used
- Example: Workstation profile I5OSP4.WS exists on desktop

- Start the PC5250 session and Communication->Configure

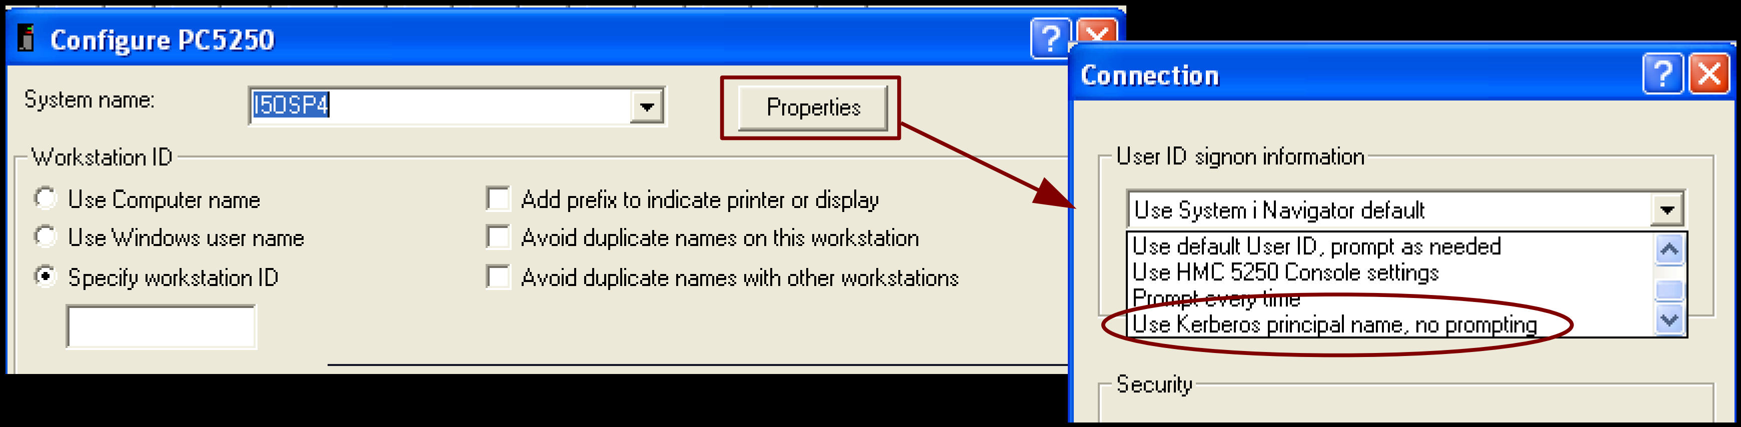

Configure PC5250

System name: I5OSP4

Workstation ID

○ Use Computer name

○ Use Windows user name

● Specify workstation ID

□ Add prefix to indicate printer or display

□ Avoid duplicate names on this workstation

□ Avoid duplicate names with other workstations

Connection

User ID signon information

Use System i Navigator default

Use default User ID, prompt as needed

Use HMC 5250 Console settings

Prompt every time

Use Kerberos principal name, no prompting

Security

- Save the changed configuration → A new file is created on the desktop
  - Same name as workstation profile, but different extension .CAE

Mass roll-out of SSO - PC5250

- .CAE file contains additional IBM i specific configuration settings
  - Important configuration directive is: UserIDSource
    - Same parameter values than in registry 4 → Kerberos

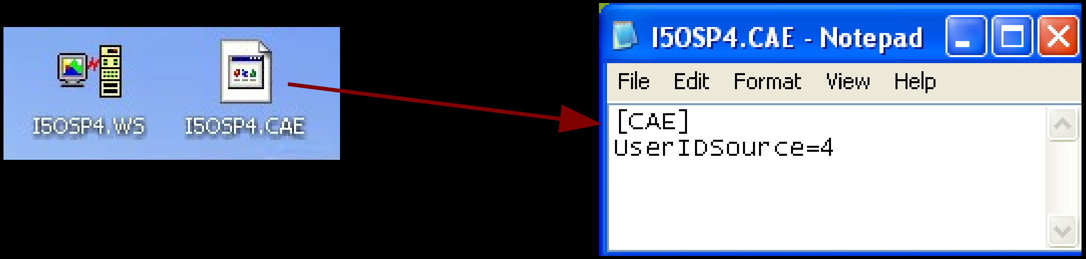

```
[CAE]
UserIDSource=4
```

- The file can be centrally distributed to client workstations, i.e. via login scripts

Mass roll-out of SSO - ODBC

- By default, IBM i ODBC data sources inherit connection settings from IBM i Navigator
  - ODBC settings are also stored in Windows registry
  - Can be centrally deployed the same way as IBM i Navigator registry settings

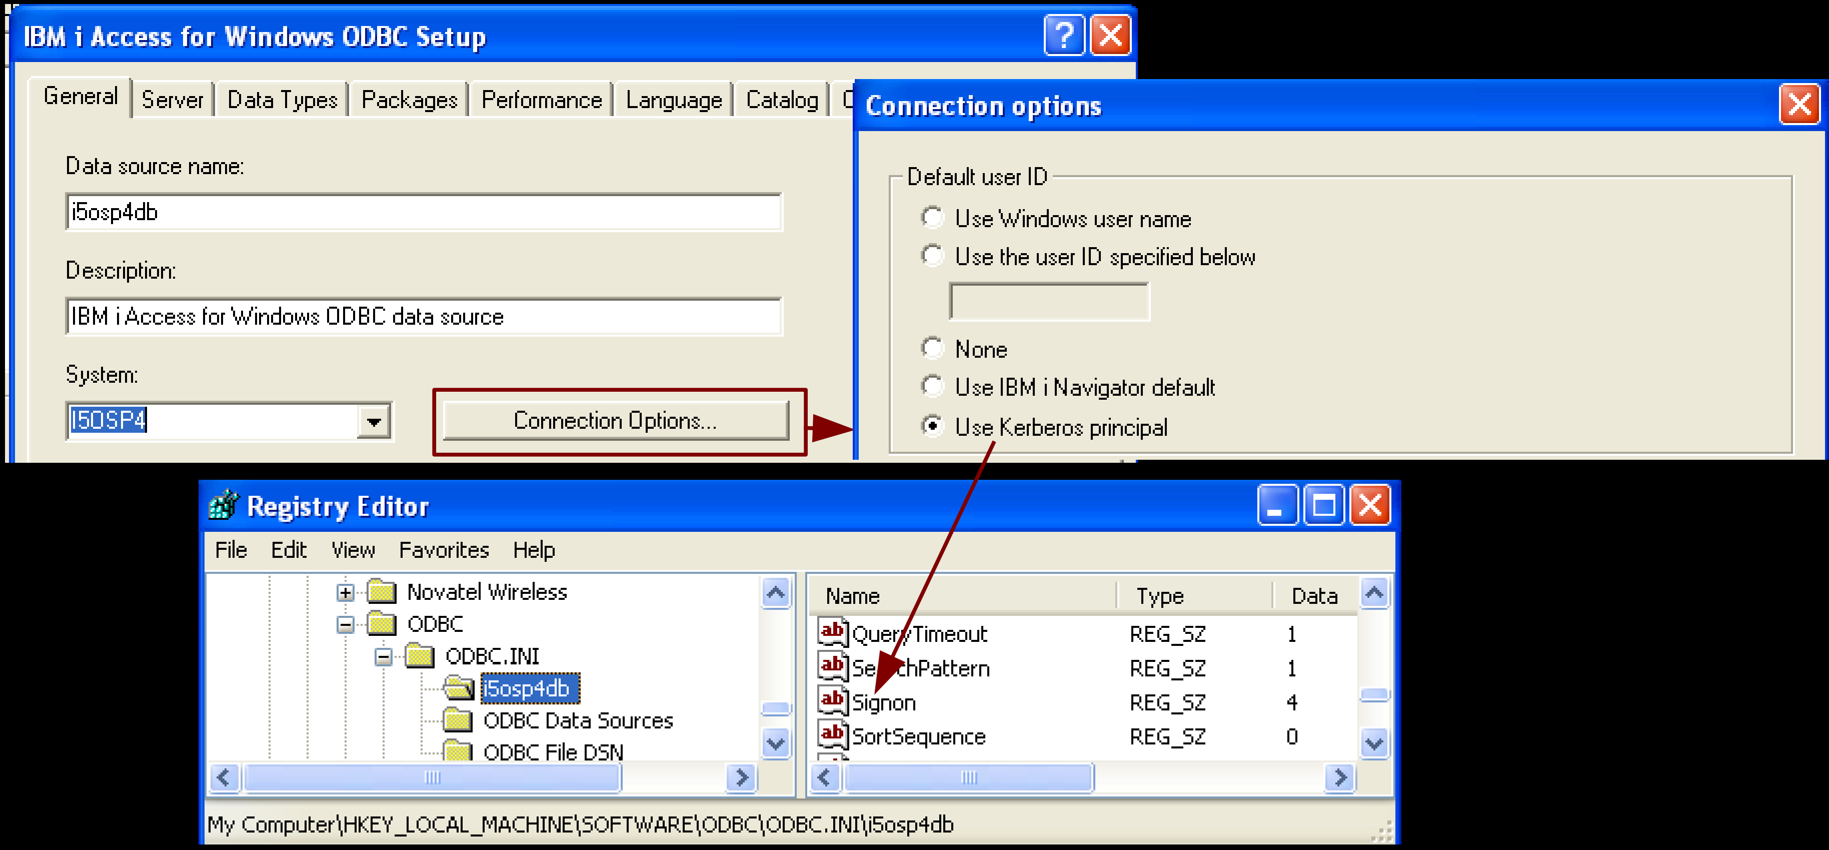

IBM i Access for Windows ODBC Setup

Data source name: i5osp4db

Description: IBM i Access for Windows ODBC data source

System: I5OSP4

Connection options

Default user ID

○ Use Windows user name

○ Use the user ID specified below

○ None

○ Use IBM i Navigator default

● Use Kerberos principal

Registry Editor

Name	Type	Data
QueryTimeout	REG_SZ	1
SearchPattern	REG_SZ	1
Signon	REG_SZ	4
SortSequence	REG_SZ	0

My Computer\HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBC.INI\i5osp4db