

---

# z/OS Communications Server

## Usability enhancements



© 2011 IBM Corporation

This presentation describes the usability enhancements in z/OS® V1R13 Communications Server. Three enhancements fall within the simplification theme. The enhancements involve Configuration Assistant support for multiple releases, multiple stacks, and for the discovery of stack IP addresses.

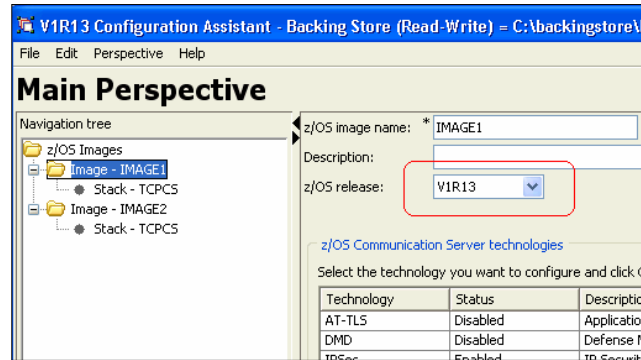



## Configuration assistant: Background

- GUI to simplify configuration for Communications Server technologies
- Unsupported (Windows® only) version initially shipped with V1R7
- Beginning in V1R11, a supported version ships with z/OSMF

The IBM Configuration Assistant for z/OS Communications Server is a graphical user interface (GUI) that simplifies configuration of many Communications Server technologies. It initially shipped for V1R7 and was included in z/OS Management Facility V1R11, which was the first release of z/OSMF.

## Multiple z/OS Communications Server releases: Introduction



3

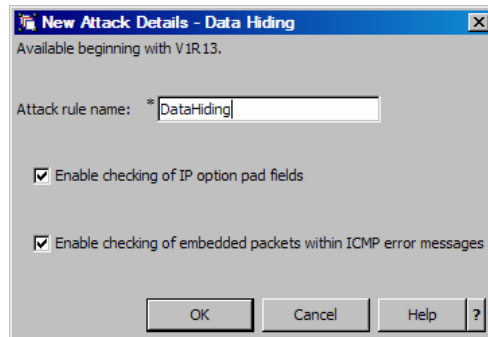
Usability enhancements

© 2011 IBM Corporation

Beginning in z/OS V1R13, the Configuration Assistant makes it easier to manage a diverse environment by supporting the configuration of multiple z/OS Communications Server releases. You no longer have to maintain multiple installations of z/OSMF in order to manage multiple releases.

In V1R13, the Configuration Assistant supports both V1R13 and V1R12 configuration. The pictures on this slide show that you can select the release level for a new image and change the release level for an existing image.

## Multiple z/OS Communications Server releases: Reusable objects



- Reusable objects can apply to all releases
- Settings not applicable to a release are ignored

Reusable objects have settings that are applied to multiple TCP/IP stacks. The stacks can be either V1R12 or V1R13. If reusable objects have any to configuration that is not applicable the release of the stack using the object, it is ignored. It is possible the configuration can be rejected and not allowed.

## Multiple z/OS Communications Server releases: Warnings

### Image Release Change Report

You changed the z/OS release of the image from V1R13 to V1R12. Some settings that are available for V1R13 are not available for V1R12. These settings can remain, but they will not be used in any z/OS Communications Server configuration file that is generated. The following report shows the details of the unavailable settings for each affected stack.

#### TCP/IP Stack: TCPCS

##### Technology: IDS

Selected requirement map: IDS\_Default

**Attacks:** The following attack settings are configured. The TCP/IP stack is at level V1R12. The attack is not available for V1R12 stacks. The requirement map is configured with this attack, but the attack will be ignored because the requirement map is selected for a V1R12-level stack.

Attack Type	Rule Name	Ignored Settings
Data Hiding Attack	DataHiding	Attack type not available; the rule is not used
IPv6 Outbound Raw Attack	IPv6OutboundRaw	Attack type not available; the rule is not used
IPv6 Destination Options Attack	IPv6DestinationOptions	Attack type not available; the rule is not used
IPv6 Hop-by-Hop Options Attack	IPv6HopByHop	Attack type not available; the rule is not used
IPv6 Next Header Attack	IPv6NextHeader	Attack type not available; the rule is not used
TCP Queue Size Attack	TcpQueueSize	Attack type not available; the rule is not used
Global TCP Stall Attack	GlobalTCPStall	Attack type not available; the rule is not used
EE Malformed Packet Attack	EEMalformedPacket	Attack type not available; the rule is not used
EE LDLC Check Attack	EELDLCCheck	Attack type not available; the rule is not used
EE Port Check Attack	EEPortCheck	Attack type not available; the rule is not used
EE XID Flood Attack	EEXIDFlood	Attack type not available; the rule is not used

If settings are ignored, the Configuration Assistant issues warnings to make you aware of this.

## Multiple z/OS Communications Server releases: Things to think about

- For migration
  - V1R12 stays V1R12
  - Others become V1R13
- All available configuration is shown
- Might be ignored or not allowed for down-level releases

When migrating prior release backing-store files into V1R13 Configuration Assistant, V1R12 level backing store files keep the releases for images at V1R12. For other release levels, the images are migrated to V1R13.

Regardless of the release level of an image or stack, all configuration settings are shown. Some settings available for only certain releases are noted. Settings not applicable to a release are either ignored or not allowed.

## Configuring multiple stacks: Introduction

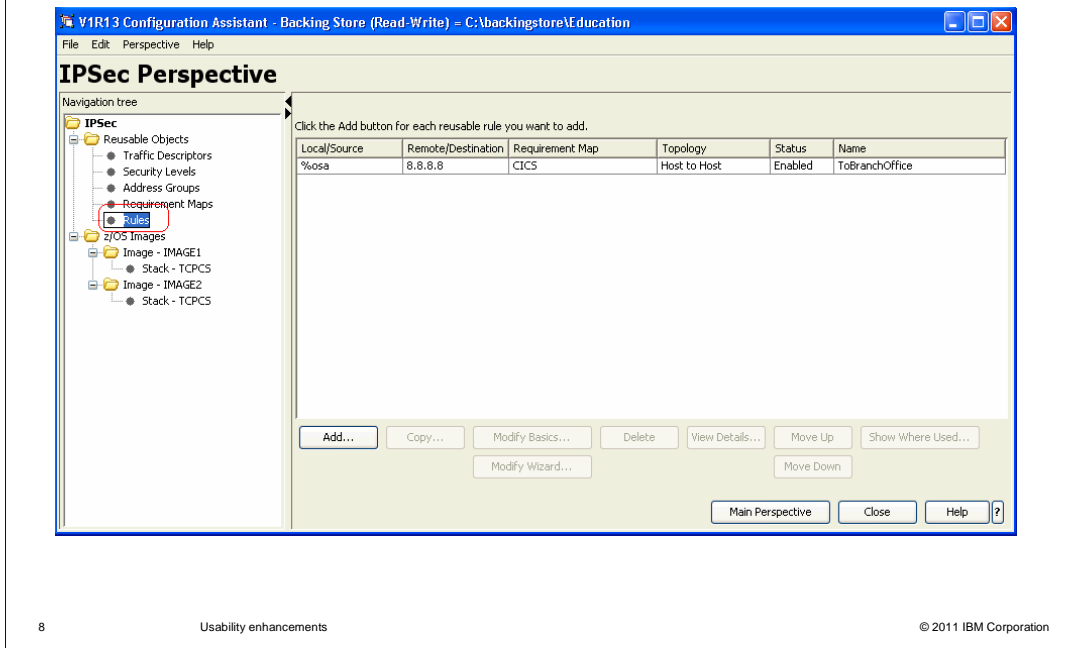
- Reusable rules
- Named variables
  - Local IP address
  - Local IKE identity

IPSec is an example of a technology which is supported by the Configuration Assistant. IPSec is configured by creating a set of rules for each TCP/IP stack. Each stack might have a large number of rules, and many of these rules are exactly the same on each stack except for the local IP addresses. You can have a lot of rules to manage, and if an update is required, it might need to be done for all stacks.

Beginning in V1R13, the Configuration Assistant supports common configuration of multiple stacks. This introduces a new reusable object called rules. Reusable rules are created a single time and assigned to one or more TCP/IP stacks. If a reusable rule needs to be updated, only a single rule needs to be modified and the changes are propagated to all stacks.

In some cases, local IP addresses and internet key exchange (IKE) identities can be shared among multiple stacks. For example, dynamic virtual IP addresses (DVIPAs) share IP addresses. In other cases, local IP addresses and IKE identities can differ from stack to stack. Reusable rules can reference variable names for both local IP addresses and IKE identities, and these names can be assigned different values for each stack.

## Configuring multiple stacks: Reusable rules



The image on the slide illustrates the new IPsec reusable rule object. You can create reusable rules by clicking on the Rules node in the navigation tree and clicking the Add button. A wizard will walk you through the creation of the rule the same as if creating connectivity rules in a TCP/IP stack. You can create reusable rules for both filtering and for dynamic tunnels.

Note that the reusable rule shown on this slide uses the name %osa for the local IP address. Later on you will see how this local address name is mapped to specific IP addresses on each stack.



## Configuring multiple stacks: Adding reusable rules

The screenshot displays the V1R13 Configuration Assistant interface. The main window is titled "IPSec Perspective" and shows a navigation tree on the left with "Stack - TCPCS" selected. The main pane shows the "Connectivity Rules" tab for the selected stack, with fields for "TCP/IP stack name" (TCPCS), "Description", and "z/OS release" (V1R13). Below these fields is a table with columns for "Local/Source", "Remote/Destination", "Requirement Map", and "Topology". An "Add..." button is highlighted with a red box. A "New Connectivity Rule - Welcome" dialog box is open in the foreground, showing the "Welcome" step. The "Indicate connectivity rule type:" section has the "Reusable rule:" radio button selected, and a drop-down menu is open showing "ToBranchOffice" selected. The "Special case:" section has the "Mobile User" option selected. The dialog box also has a "Welcome" step selected in the left pane.

You can add a reusable rule to an individual stack by selecting the stack tree node and clicking the Add button below the list of connectivity rules. Then select the Reusable rule radio button, and choose the reusable rule name from the drop-down list.

## Configuring multiple stacks: Converting stack rules

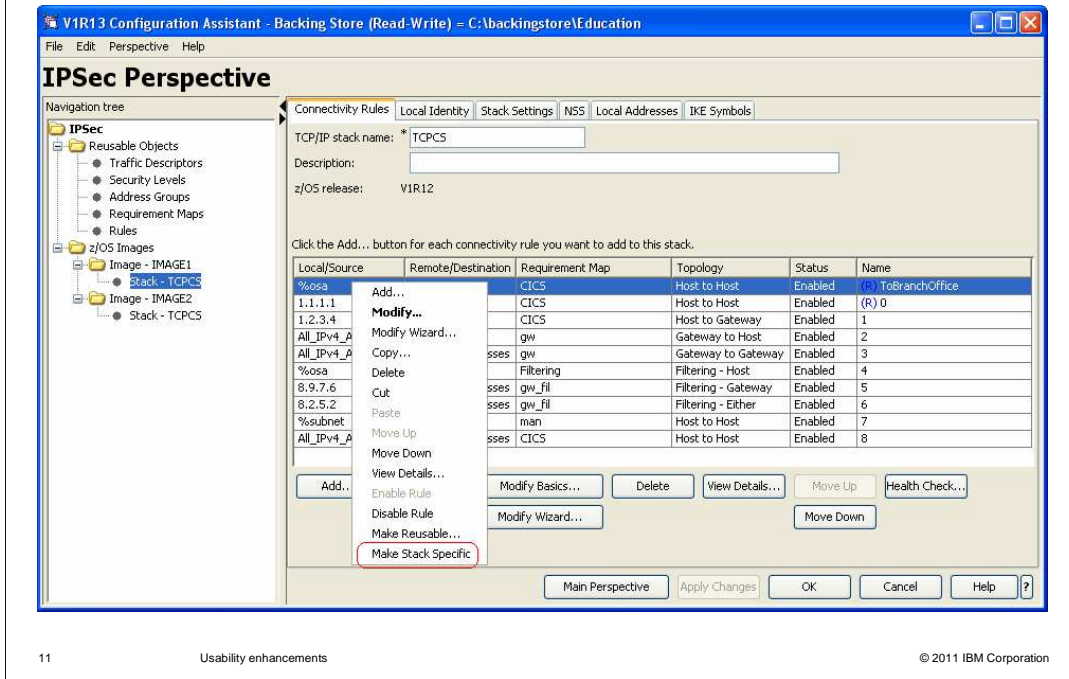
The screenshot shows the 'IPSec Perspective' window. The navigation tree on the left includes 'Reusable Objects' and 'z/OS Images'. The main area shows the 'Connectivity Rules' tab for stack 'TCPCS'. A table lists connectivity rules with columns: Local/Source, Remote/Destination, Requirement Map, Topology, Status, and Name. A context menu is open over the table, with 'Make Reusable...' highlighted. The table contains the following data:

Local/Source	Remote/Destination	Requirement Map	Topology	Status	Name
%osa	8.8.8.8	CICS	Host to Host	Enabled	(R) ToBranchOffice
1.1.1.1	2.2.2.2	CICS	Host to Host	Enabled	(R) 0
1.2.3.4	9.9.9.9	CICS	Host to Gateway	Enabled	1
All_IPv		gw	Gateway to Host	Enabled	2
All_IPv		dresses gw	Gateway to Gateway	Enabled	3
%osa		Filtering	Filtering - Host	Enabled	4
8.9.7.6		dresses gw_fil	Filtering - Gateway	Enabled	5
8.2.5.2		dresses gw_fil	Filtering - Either	Enabled	6
%subn		man	Host to Host	Enabled	7
All_IPv		dresses CICS	Host to Host	Enabled	8

Reusable rules appear in a stack's list of connectivity rules alongside stack-specific connectivity rules. The reusable rules are identified with the blue (R) symbol as shown on this slide. You can change the order of reusable rules just like stack-specific connectivity rules.

To help you migrate from stack-specific rules to reusable rules, you can change a stack rule into a reusable rule. To do this, select a stack-specific rule and click Make Reusable. A copy of the stack rule is added to the reusable rule set. The stack rule will remain in place, but is marked as derived from a reusable rule. The (R) symbol will now appear to indicate this rule is a reusable rule.

## Configuring multiple stacks: Converting reusable rules

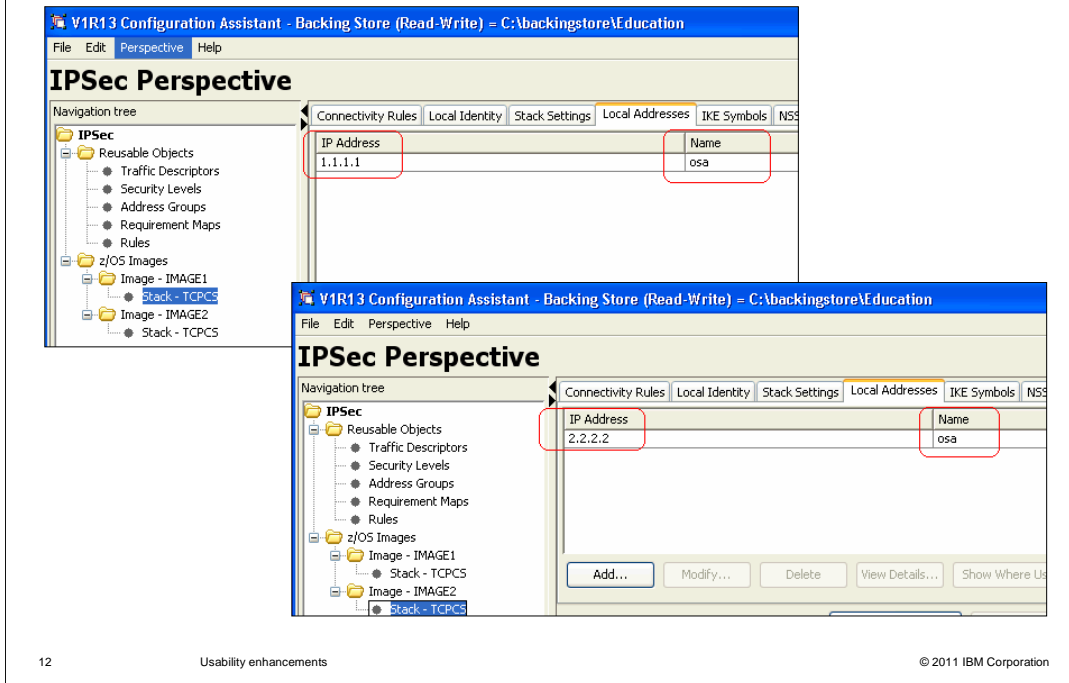


You might need to change a rule setting, but only on a specific TCP/IP stack. For example, you might want to change the filter logging settings on one stack without impacting other stacks.

If the rule you want to change is a reusable rule, you have the option to change it to a stack-specific rule. After the change, the rule will remain in the stack, but is marked as no longer derived from a reusable rule. The reusable rule will remain in the set of reusable rules, but it will no longer be associated with the stack-specific rule.

If this change is temporary, you can later delete the stack-specific rule and add the reusable rule back to the stack.

## Configuring multiple stacks: Local IP address names



Some reusable rules might have the same local IP address for each stack; this is true for a DVIPA or a distributed DVIPA. But other rules might have different local IP addresses for each stack. In order for reusable rules to be used for this case, you can assign variable names to local IP addresses, and assign different values for these addresses on each stack.

After selecting a stack in the navigation tree, the Local Addresses tab shows all address names known to the Configuration Assistant. You can add new names and assign values to existing names in this tab. Once a name is configured for one stack, that name is available to all stacks.

The images on this slide show that the name “osa” is assigned different values for two different stacks.

## Configuring multiple stacks: Address names in rules

**New Connectivity Rule - Data Endpoints**

Welcome  
 Typical Rule  
 Network Topology  
 Data Endpoints  
 Requirement Map  
 Local Security Endpoint  
 Remote Security Endpoint  
 Manual Tunnel Keys  
 Special Case: Mobile User  
 Special Case: IP V6 OSPF IP Security  
 Finish

**Data Endpoints**  
Use this panel to identify the data endpoints.  
These are the IP addresses of the host endpoints of the traffic you want to protect.

Host To Host - Data Endpoints

Connectivity rule name: \* 8

**Local data endpoint**

Address group:  
All\_IPv4\_Addresses  
New... Copy... Modify... View Details... Show Where Used...

IPv4 or IPv6 address, subnet or range:  
\*  
Examples: x.x.x.x, x.x.x.x/yy, x.x.x.x-y.y.y

Local IP address name:  
osa

**Remote data endpoint**

Address group:  
All\_IPv4\_Addresses  
New... Copy...

IPv4 or IPv6 address:  
\*  
Examples: x.x.x.x, x.x.x.x/yy, x.x.x.x-y.y.y

When a local address name is configured, you can choose it as the local data endpoint for a connectivity rule. Local address names can be configured in both reusable rules and stack-specific connectivity rules.

If you use a local address name in a reusable rule which is assigned to several stacks, the name resolves to the unique IP address assigned to that name on each stack.

## Configuring multiple stacks: Address names in groups

**New IP Address Group**

Use this panel to configure a group of IP addresses.

9.37.236.137	4.98.124.187
5.96.158.184	
2.45.197.242	3.15.141.211
7.37.253.241	
8.63.138.187	8.63.138.187

Name: \* localAddresses

Description:

Type an IP address, range, or subnet directly into the table below or click the Select Local Address Name button.

IP address
%osa

Add More Rows    Select Local Address Name...

**Select Local Addresses by Name**

Select the names. Multiple selections are allowed.

- osa
- single
- subnet
- v4
- v6

You can also add address names to address groups. Address groups containing local address names can be chosen as the local data endpoint for both reusable rules and stack-specific rules. You can type the local address names directly into the table, or select the name by clicking Select Local Address Name. On this slide, the local IP address name “osa” is included in the address group “localAddresses.”

## Configuring multiple stacks: Local IKE identities

The image displays two screenshots of the V1R13 Configuration Assistant interface, specifically the 'IPSec Perspective' for the 'Local Identity' tab. Both screenshots show a navigation tree on the left and a table of IKE identities on the right. Red boxes highlight the 'Value' and 'Name' columns in the table.

**Top Screenshot:** Shows the configuration for a stack where the 'Value' is 'IP Address: 1.1.1.1' and the 'Name' is 'LocalIkeIdentity'.

Type	Value	Name
IKE Identity	IP Address: 1.1.1.1	LocalIkeIdentity

**Bottom Screenshot:** Shows the configuration for a stack where the 'Value' is 'User ID @ FQDN: user1@f.q.d.n' and the 'Name' is 'LocalIkeIdentity'.

Type	Value	Name
IKE Identity	User ID @ FQDN: user1@f.q.d.n	LocalIkeIdentity

Both screenshots show a navigation tree with 'IPSec' expanded, containing 'Reusable Objects' (Traffic Descriptors, Security Levels, Address Groups, Requirement Maps, Rules) and 'z/OS Images' (Image - IMAGE1, Image - IMAGE2, Stack - TCPCS).

15

Usability enhancements

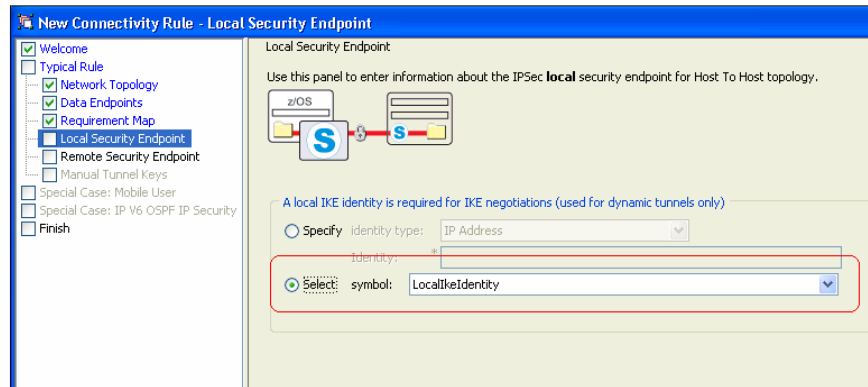
© 2011 IBM Corporation

Just as with local IP addresses, it is likely that you will configure different local IKE identities for each stack that uses a reusable rule. You can therefore assign variable names to local IKE identities, and assign different values for these identities on each stack.

After selecting a stack in the navigation tree, the IKE Symbols tab shows all IKE identities known to the Configuration Assistant. You can add new names and assign values to existing names in this tab. Once a name is configured for one stack, that name is available to all stacks.

The images on this slide show that the name "LocalIkeIdentity" is assigned different values for two different stacks. Note that the identities do not have to share the same identity type on each stack.

## Configuring multiple stacks: Identity names in rules

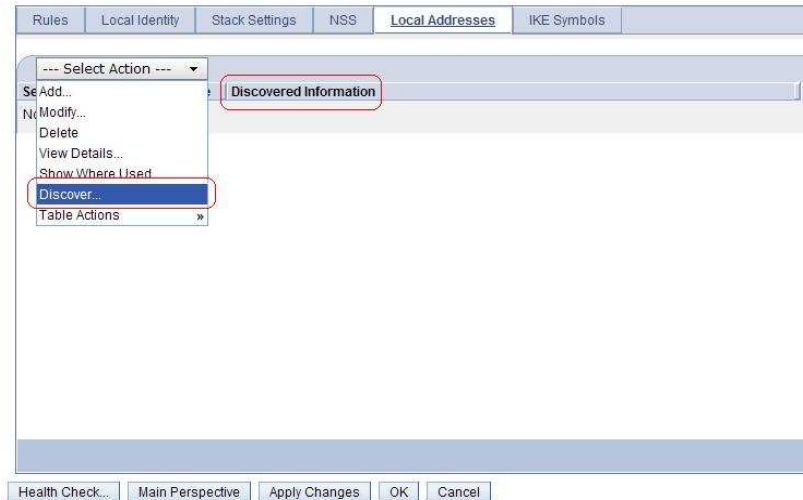


When a local IKE identity is configured, you can choose it as the local security endpoint identity for a connectivity rule. Local IKE identities can be configured in both reusable rules and stack-specific connectivity rules.

If you use a local identity name in a reusable rule which is assigned to several stacks, the name resolves to the unique IKE identity assigned to that name on each stack.



## Discovery of stack IP addresses: Introduction



17

Usability enhancements

© 2011 IBM Corporation

With the introduction of reusable rules and local IP address names, the Configuration Assistant also supports the discovery of IP addresses for each stack. Thus, you can avoid manually entering IP addresses for your reusable rules.

The discover function is only available on z/OSMF. It is not available in the Configuration Assistant workstation GUI. You can initiate the discovery process from the Local Addresses tab for a stack by choosing the Discover action, as shown on this slide.

## Discovery of stack IP addresses: Discovery

Discover Stack Local Addresses

Image: VIC136  
Stack: TCPCS1

Host connection

\*Host name:   
\*Port:   
\*User name:  \*Password:   Save password  
 Use SSL

Click Go to begin discovery

When you select the discover action, the Discover Stack Local Addresses panel is displayed. Because the discover function is invoked while you are editing a stack, the image and stack names are already known, as shown on the panel. You fill in the connection information to connect to the Policy Agent and select the Go button. The Policy Agent must be active and configured with the ServicesConnection statement in order to successfully connect. This is the same Policy Agent configuration statement you use for the Policy Data Import function.

You will see a message when the discovery function completes, indicating success or failure. In some cases, IP addresses cannot be added to the Local Addresses tab, for example when a name is already used for another IP address. The message will indicate any such failures and the discovery report will contain more detail. You can access the discovery report by clicking the View Report button shown here.

## Discovery of stack IP addresses: Discovery report

### Discovery Local Addresses Summary

The following data set names were used to discover local address information:

CS390.BASE.TCPPARMS(TCP.CS1)  
 CS390.BASE.TCPPARMS(VPADEF1)  
 CS390.BASE.TCPPARMS(TCP.CS1B)  
 CS390.BASE.TCPPARMS(TN3270S1)  
 USER.TCPPARMS(TCP.CS1)  
 CS390.BASE.TCPPARMS(STRTDEV1)

The following IP addresses did not already exist, and were added to the local addresses tab. The associated interface names were used as the local address names.

IP Address	Local Address Name	Discovered Information
10.61.0.1	DYNXCF4	Type=Dynamic XCF, Name=DYNXCF4
2001:db8:10::61:0:1	DYNXCF6	Type=Dynamic XCF, Name=DYNXCF6
127.0.0.2	LOOPBACK	Type=Loopback, Name=LOOPBACK
10.81.1.1	VIPA4811L	Type=Static VIPA, Name=VIPA4811L

...

The following IP addresses did not already exist, and were added to the local addresses tab. The associated interface names were not used, because of the reasons shown in the table.

IP Address	Discovered Information	Reason
127.0.0.3	Type=Loopback, Name=LOOPBACK	The interface name is already used as a local address name.
127.0.0.4	Type=Loopback, Name=LOOPBACK	The interface name is already used as a local address name.
10.91.1.1/24	Type=Dynamic VIPA Define	This IP address does not have an interface name.

This slide shows parts of a discovery report. This report is available at any time by selecting the History action from the main Action menu. The report shows details of the connection, the complete list of TCP/IP configuration data sets, and details about all of the interfaces. It also includes details about cases where the interface name cannot be used.

## Discovery of stack IP addresses: After discovery

Rules	Local Identity	Stack Settings	NSS	Local Addresses	IKE Symbols
--- Select Action ---					
Select	IP Address	Name	Discovered Information		
<input type="radio"/>	2001:db8:10::92:1:1	VIPA6921	Type=Dynamic VIPA Define, Name=VIPA6921		
<input type="radio"/>	2001:db8:10::91:1:1	VIPA6911	Type=Dynamic VIPA Define, Name=VIPA6911		
<input type="radio"/>	10.71.0.0/16		Type=Dynamic VIPA Range		
<input type="radio"/>	10.91.3.3		Type=Dynamic VIPA Backup, backup rank=100		
<input type="radio"/>	10.93.1.1/24		Type=Dynamic VIPA Define		
<input type="radio"/>	10.92.1.1/24		Type=Dynamic VIPA Define		
<input type="radio"/>	10.91.1.1/24		Type=Dynamic VIPA Define		
<input type="radio"/>	::14:0	LOOPBACK6	Type=Loopback, Name=LOOPBACK6		
<input type="radio"/>	2001:db8::64	MPC1IPV6	Type=MPC, Name=MPC1IPV6		
<input type="radio"/>	2001:db8:172::16:2:1	QDIO6201	Type=OSAD, Name=QDIO6201		
<input type="radio"/>	2001:db8:172::16:1:1	QDIO6101	Type=OSAD, Name=QDIO6101		

Health Check... Main Perspective Apply Changes OK Cancel

This slide has an example of the Local Addresses tab after performing the discovery process. Addresses associated with an interface or VIPA are assigned its name as their local address name. Other addresses are imported without a name, but you can assign names to them. Additional information discovered about these addresses is shown in the Discovered Information column.

## Discovery of stack IP addresses: things to think about

- Only strategic interfaces are discovered
  - OSA-Express
  - Hipersockets
  - MPC Point-to-Point
  - Dynamic XCF
  - Static VIPA
  - Distributed VIPA (DVIPA)
  - Loopback
- Address names assigned by discovery are available in all other stacks, but without IP address values

Only the interface types listed on this slide are recognized by the discovery process. Non-strategic interfaces such as CLAW and LCS are not subject to discovery.

When you use the discovery function on a given stack, you will have IP addresses with names assigned to most or all of the addresses. Address names are common to all stacks, so you will see those names appear on other stacks. When you use the discovery function on any stack that already contains names without assigned values, the Configuration Assistant will update the IP address values for all address names that match discovered interface names.



## Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

[mailto:iea@us.ibm.com?subject=Feedback\\_about\\_usability.ppt](mailto:iea@us.ibm.com?subject=Feedback_about_usability.ppt)

This module is also available in PDF format at: [../usability.pdf](..../usability.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.



## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, and z/OS are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2011. All rights reserved.