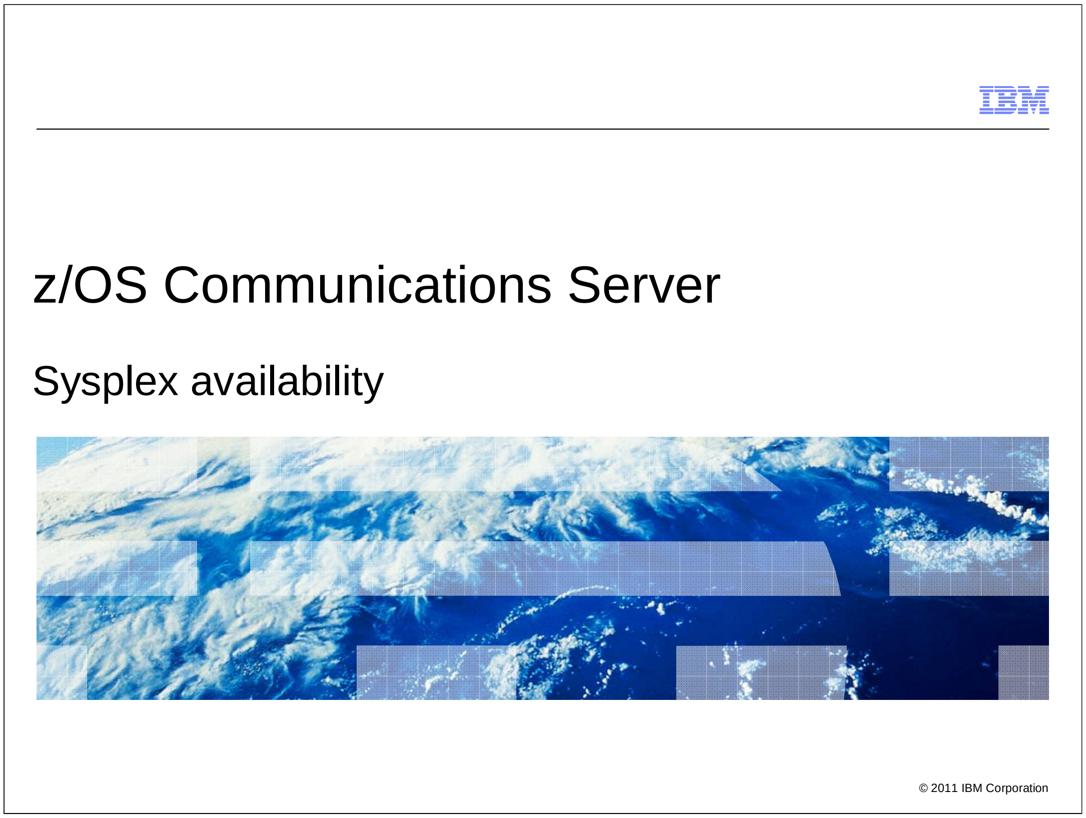

z/OS Communications Server

Sysplex availability

This presentation describes the updates to sysplex availability in z/OS® V1R13 Communications Server.

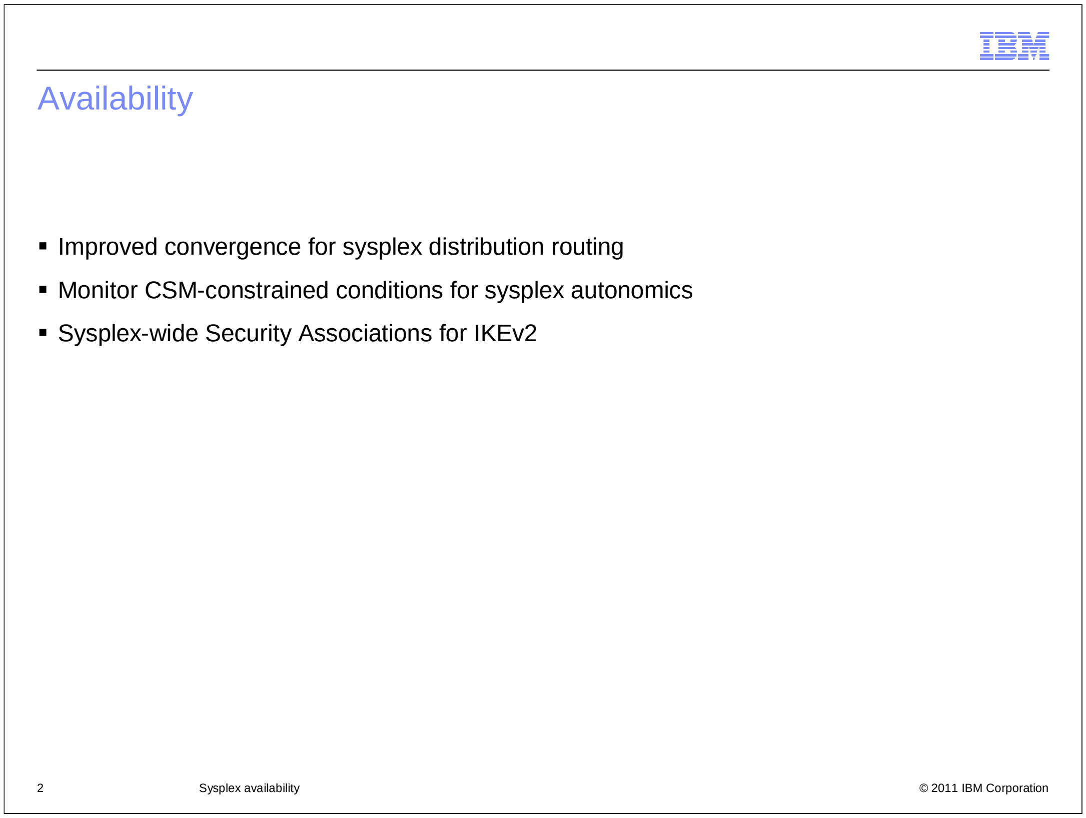

Availability

- Improved convergence for sysplex distribution routing
- Monitor CSM-constrained conditions for sysplex autonomics
- Sysplex-wide Security Associations for IKEv2

2 Sysplex availability © 2011 IBM Corporation

The availability theme includes sysplex distribution and monitoring, and sysplex-wide security associations (SWSA).

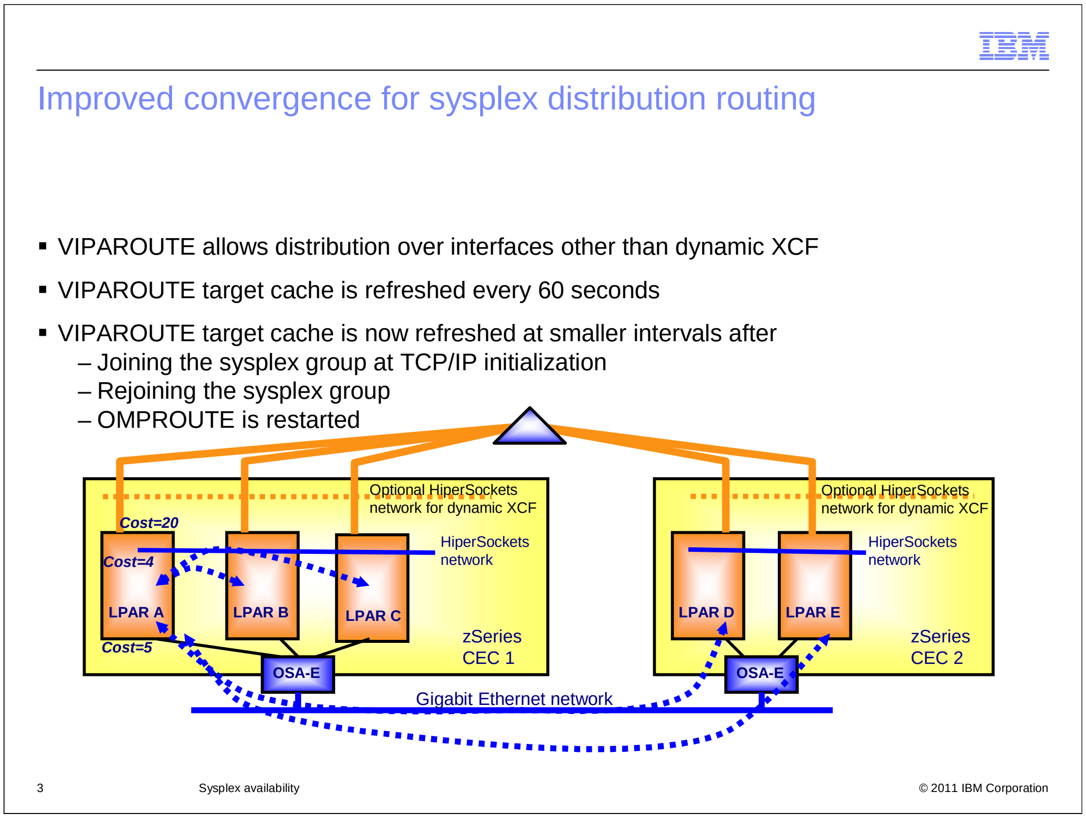

Improved convergence for sysplex distribution routing

- VIPAROUTE allows distribution over interfaces other than dynamic XCF
- VIPAROUTE target cache is refreshed every 60 seconds
- VIPAROUTE target cache is now refreshed at smaller intervals after
  - Joining the sysplex group at TCP/IP initialization
  - Rejoining the sysplex group
  - OMPROUTE is restarted

3 Sysplex availability © 2011 IBM Corporation

The VIPAROUTE optimized sysplex distributor routing function allows the distribution of sysplex distributor traffic over interfaces other than dynamic XCF.

VIPAROUTE maintains a target cache for target routes. This target cache is refreshed every 60 seconds. However, during a takeover this can result in a period of non-optimal routing or even traffic disruption until the target cache is refreshed.

To avoid this, z/OS V1R13 Communications Server changes the cache refresh interval when joining or rejoining the sysplex, and when OMPROUTE is restarted. During a takeover that occurs because the primary routing stack is restarted, a route is unavailable until its interface finishes activating. It uses an interval pattern of five, five, 15, 35 and 60 seconds to allow for faster convergence of target routes.

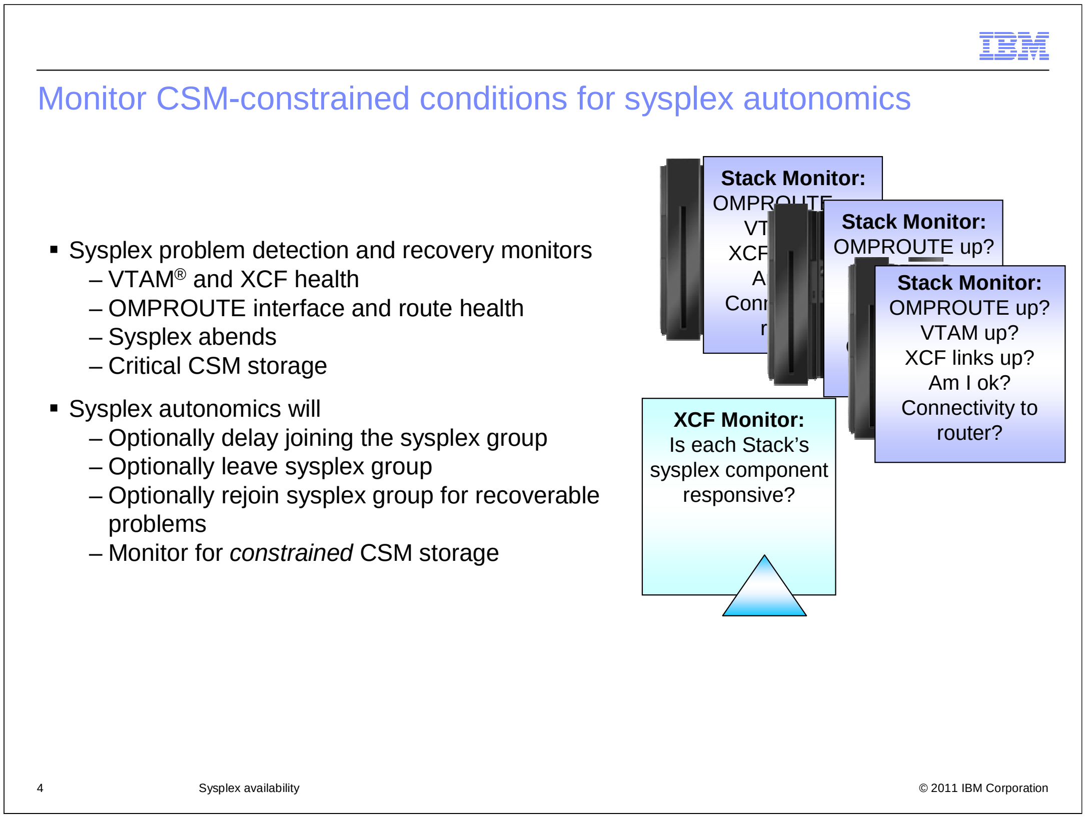

Monitor CSM-constrained conditions for sysplex autonomics

- Sysplex problem detection and recovery monitors
  - VTAM® and XCF health
  - OMPROUTE interface and route health
  - Sysplex abends
  - Critical CSM storage
- Sysplex autonomics will
  - Optionally delay joining the sysplex group
  - Optionally leave sysplex group
  - Optionally rejoin sysplex group for recoverable problems
  - Monitor for *constrained* CSM storage

Stack Monitor:
OMPROUTE
VT
XCF
A
Conn
r

Stack Monitor:
OMPROUTE up?

Stack Monitor:
OMPROUTE up?
VTAM up?
XCF links up?
Am I ok?
Connectivity to router?

XCF Monitor:
Is each Stack's sysplex component responsive?

4 Sysplex availability © 2011 IBM Corporation

Sysplex autonomics improves availability by automatically detecting and recovering from problems that can affect the health of the sysplex. It monitors various functions that are related to sysplex distribution, and can take action to delay joining the sysplex group, leave the sysplex group, and rejoin the sysplex group when problems have cleared.

Beginning in V1R13, sysplex autonomics now monitors for CSM FIXED and ECSA storage reaching constrained state. Previously it monitored only for CSM-critical storage, but CSM-constrained conditions can also impact TCP/IP's ability to process network traffic.

CSM will now monitor for CSM-constrained storage. When CSM storage is constrained for more than three times the configured TIMERSECS value, CSM will issue operator message EZD1974E. If RECOVERY was specified on the SYSPLEXMONITOR parameter of the GLOBALCONFIG statement, and this TCP/IP stack is not the only member of the TCP/IP sysplex group, the stack will take several actions. It will leave the sysplex group, inactivate all DVIPAs, and save all VIPADYNAMIC block definitions. The new operator message (EZD1974E) is deleted when CSM's constrained state is exited. If AUTOREJOIN was specified, the stack will rejoin the sysplex group and reprocess the saved VIPADYNAMIC configuration.

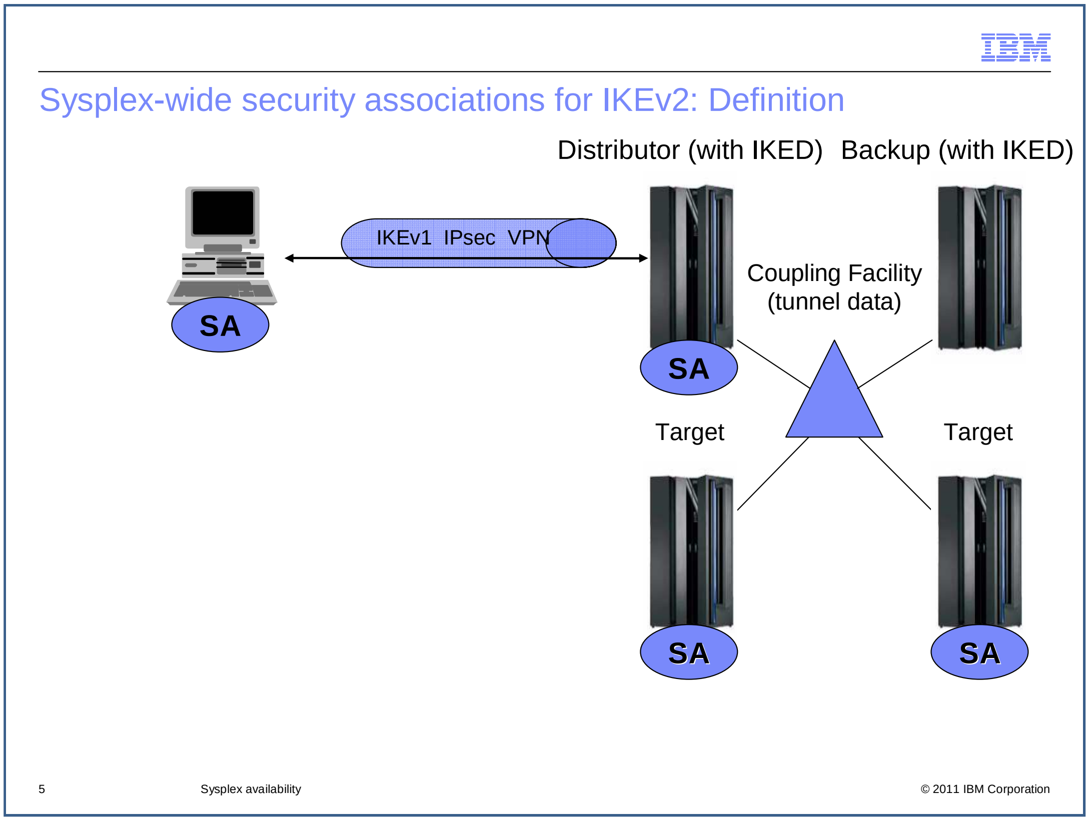

Sysplex-wide security associations for IKEv2: Definition

5 Sysplex availability © 2011 IBM Corporation

Sysplex-wide Security Associations (SWSA) represents the intersection of two Communications Server functions. The first, sysplex distributor, allows for the distribution of application workload along with backup and recovery mechanisms, using Dynamic Virtual IP Addresses (DVIPAs). The second, IPsec, protects network data using Security Associations (SAs).

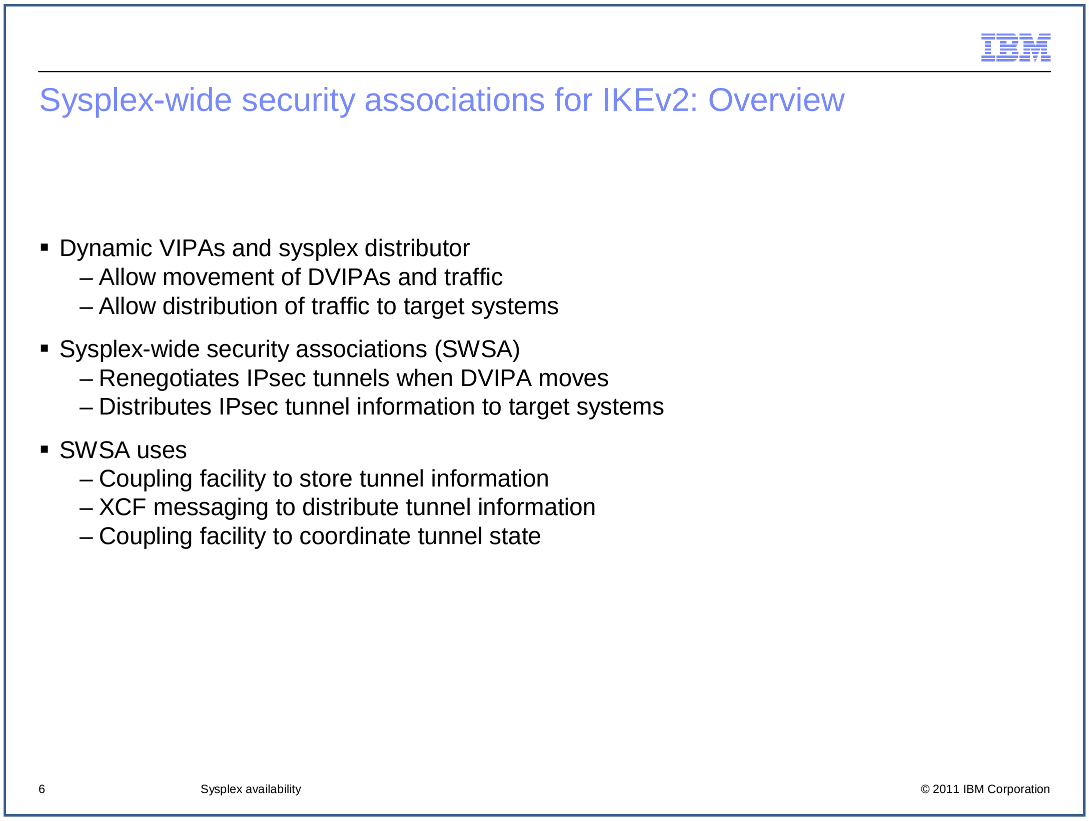

Sysplex-wide security associations for IKEv2: Overview

- Dynamic VIPAs and sysplex distributor
  - Allow movement of DVIPAs and traffic
  - Allow distribution of traffic to target systems
- Sysplex-wide security associations (SWSA)
  - Renegotiates IPsec tunnels when DVIPA moves
  - Distributes IPsec tunnel information to target systems
- SWSA uses
  - Coupling facility to store tunnel information
  - XCF messaging to distribute tunnel information
  - Coupling facility to coordinate tunnel state

6 Sysplex availability © 2011 IBM Corporation

SWSA allows you to exploit both functions together. With SWSA you can encrypt sysplex distributor workload. The distributor is responsible for negotiating an SA with a remote host. Copies of the SA, known as shadow SAs, are sent to any target stacks that can potentially receive workload for the DVIPA. In addition to distribution, you can recover SAs associated with DVIPAs that are migrated to an alternate TCP/IP stack (DIPVA takeover). Any stacks that backup the DVIPA do not receive SA data from the distributor directly, but have access to SA data in the coupling facility (CF) needed for SA recovery should a DVIPA move.

IKED must run on any stack that will potentially negotiate SAs, including the distributor and backup stacks. Target stacks need not run IKED.

Critical SA data that is shared among sysplex members is stored in the coupling facility for SAs that are distributed, or are candidates for takeover.

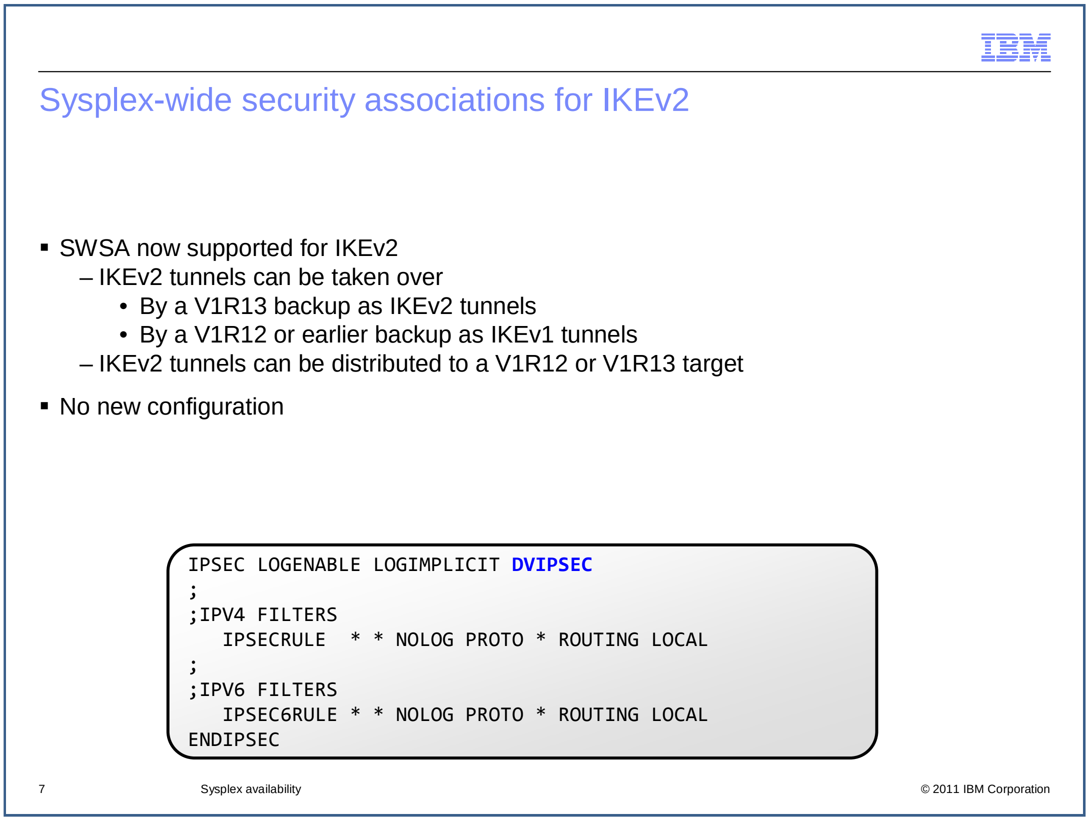

Sysplex-wide security associations for IKEv2

- SWSA now supported for IKEv2
  - IKEv2 tunnels can be taken over
    - By a V1R13 backup as IKEv2 tunnels
    - By a V1R12 or earlier backup as IKEv1 tunnels
  - IKEv2 tunnels can be distributed to a V1R12 or V1R13 target
- No new configuration

```
IPSEC LOGENABLE LOGIMPLICIT DVIPSEC
;
;IPV4 FILTERS
   IPSECRULE  * * NOLOG PROTO * ROUTING LOCAL
;
;IPV6 FILTERS
   IPSEC6RULE * * NOLOG PROTO * ROUTING LOCAL
ENDIPSEC
```

7 Sysplex availability

SWSA has been available since z/OS V1R2 Communications Server was released in 2001. IKEv2 support was introduced in z/OS V1R12 Communications Server, but SWSA distribution and takeover were not allowed for IKEv2 tunnels. In V1R12, all traffic flowing through IKEv2 tunnels for DVIPA addresses must be serviced at the distributor stack, and these tunnels cannot be recovered on DVIPA takeover.

Beginning in z/OS V1R13 Communications Server, SWSA takeover and distribution are supported for IKEv2 tunnels. No new configuration is needed; SWSA support for IKEv2 is enabled just as for IKEv1, using the DVIPSEC parameter on the IPSEC TCP/IP profile statement. IKEv1 and IKEv2 SWSA tunnels can coexist in the same sysplex.

Backup stacks running at a V1R12 or earlier release level will attempt to recover IKEv2 SWSA tunnels from a V1R13 distributor. The tunnels are negotiated using IKEv1, and the negotiations will not use any attributes unique to IKEv2 (such as port ranges, opaque traffic selectors, or ICMP types and codes).

Target stacks running at a V1R12 release level will receive IKEv2 SWSA “shadow” tunnels and can serve as targets for connections using these tunnels.

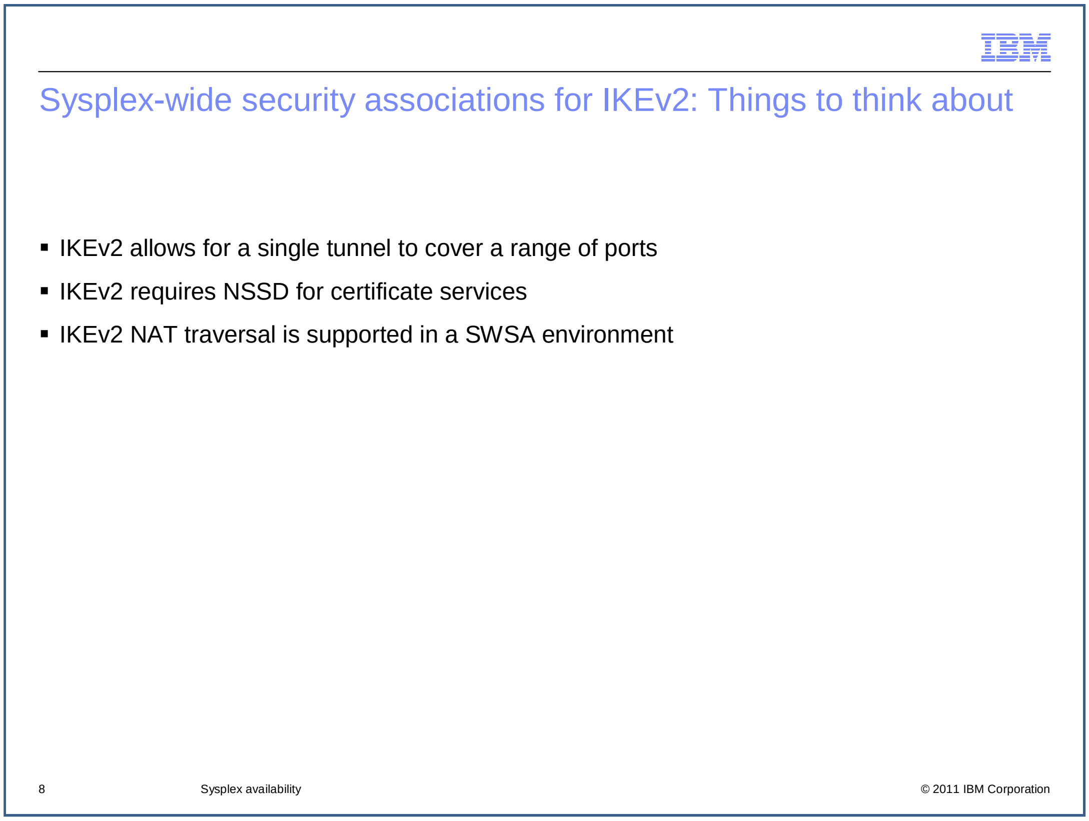

Sysplex-wide security associations for IKEv2: Things to think about

- IKEv2 allows for a single tunnel to cover a range of ports
- IKEv2 requires NSSD for certificate services
- IKEv2 NAT traversal is supported in a SWSA environment

8 Sysplex availability © 2011 IBM Corporation

IKEv1 does not allow for tunnels that cover port ranges; it only allows for tunnels protecting specific ports or all ports. However, IKEv2 does allow for tunnels protecting a port range. If you have several IKEv1 tunnels protecting ports in a port range, these tunnels can be replaced by a single IKEv2 tunnel.

IKED’s IKEv1 support allows for certificate operations to take place either using IKED’s local key ring or using NSSD’s certificate services. (NSSD is the Network Security Server Daemon.) The IKEv2 standards require the use of advanced certificate services, such as certificate revocation lists, that are available only using NSSD. If you choose to use IKEv2 and want to use certificates for identity protection, you need to configure IKED as a certificate services client to NSSD.

IKEv2 support for NAT traversal is described in a separate presentation in the security section. IKEv2 NAT traversal is also supported in a SWSA environment.

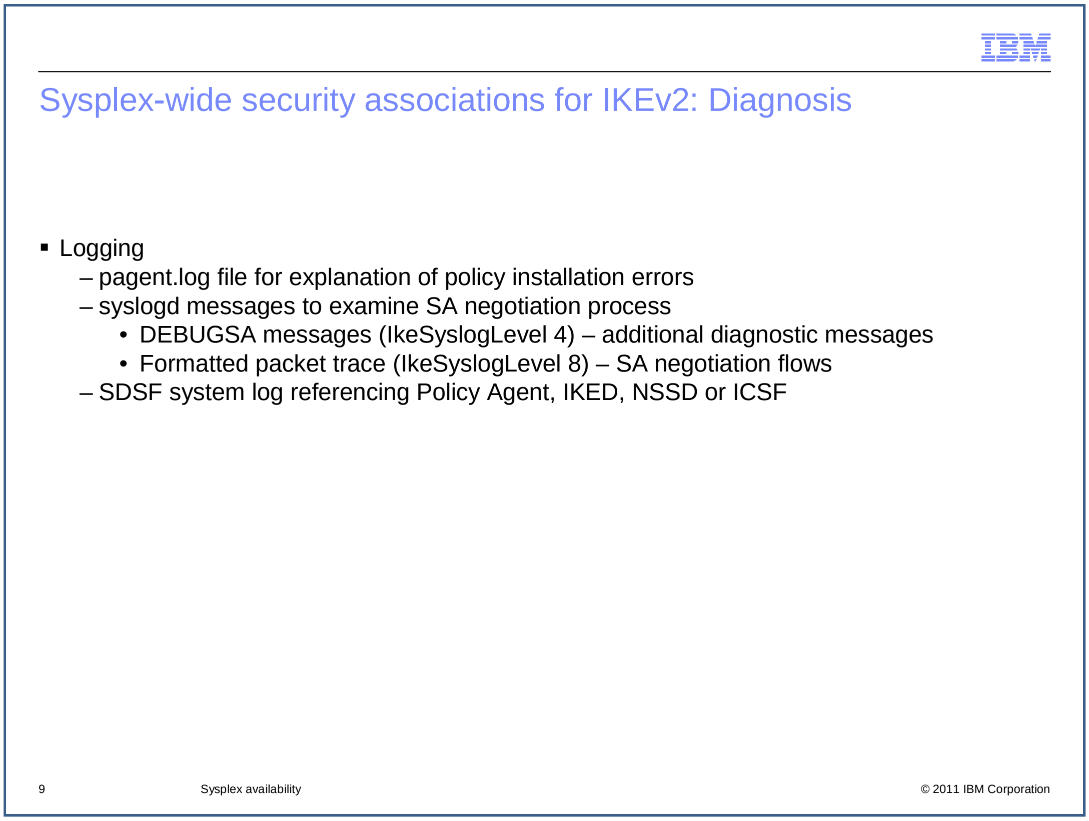

Sysplex-wide security associations for IKEv2: Diagnosis

- Logging
  - pagent.log file for explanation of policy installation errors
  - syslogd messages to examine SA negotiation process
    - DEBUGSA messages (IkeSyslogLevel 4) – additional diagnostic messages
    - Formatted packet trace (IkeSyslogLevel 8) – SA negotiation flows
  - SDSF system log referencing Policy Agent, IKED, NSSD or ICSF

Start diagnosis for IKEv2 SWSA by checking several logs. For more information about these logs, see the *z/OS Communications Server: IP Diagnosis Guide*.

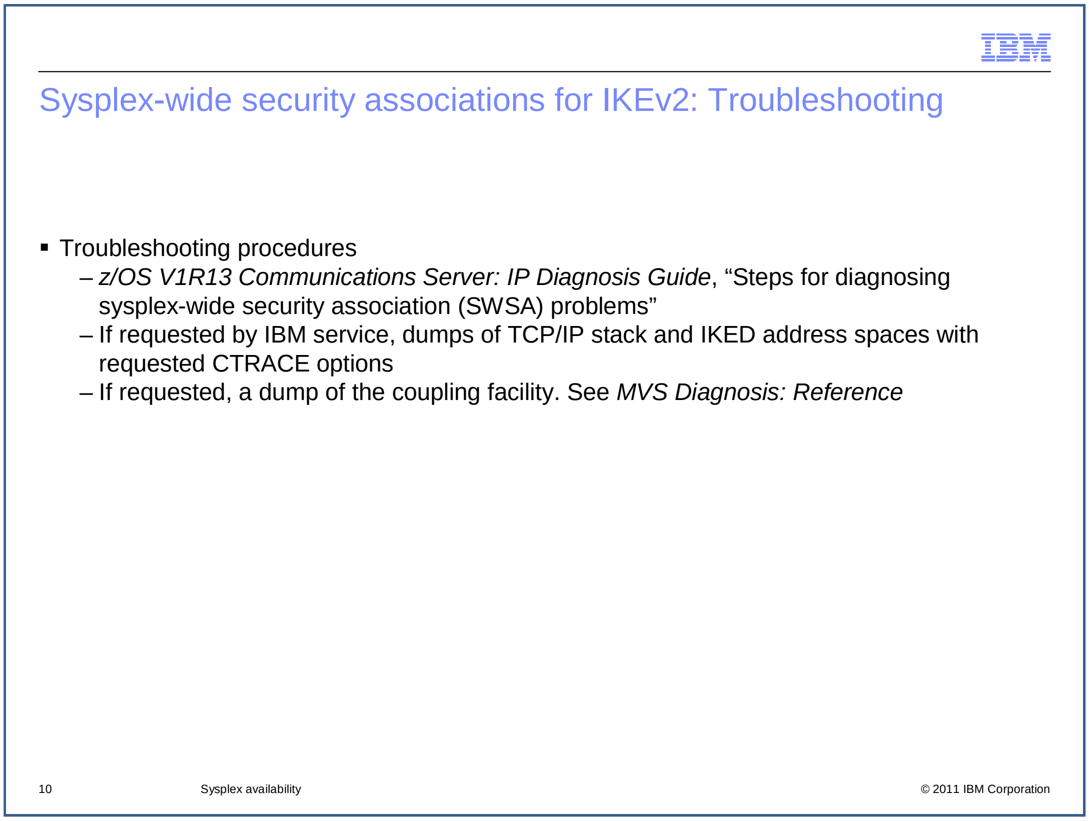

Sysplex-wide security associations for IKEv2: Troubleshooting

- Troubleshooting procedures
  - *z/OS V1R13 Communications Server: IP Diagnosis Guide*, “Steps for diagnosing sysplex-wide security association (SWSA) problems”
  - If requested by IBM service, dumps of TCP/IP stack and IKED address spaces with requested CTRACE options
  - If requested, a dump of the coupling facility. See *MVS Diagnosis: Reference*

Troubleshooting for IKEv2 SWSA is similar to troubleshooting for IKEv1 SWSA problems. See “Steps for diagnosing sysplex-wide security association (SWSA) problems” in the *z/OS Communications Server: IP Diagnosis Guide*.

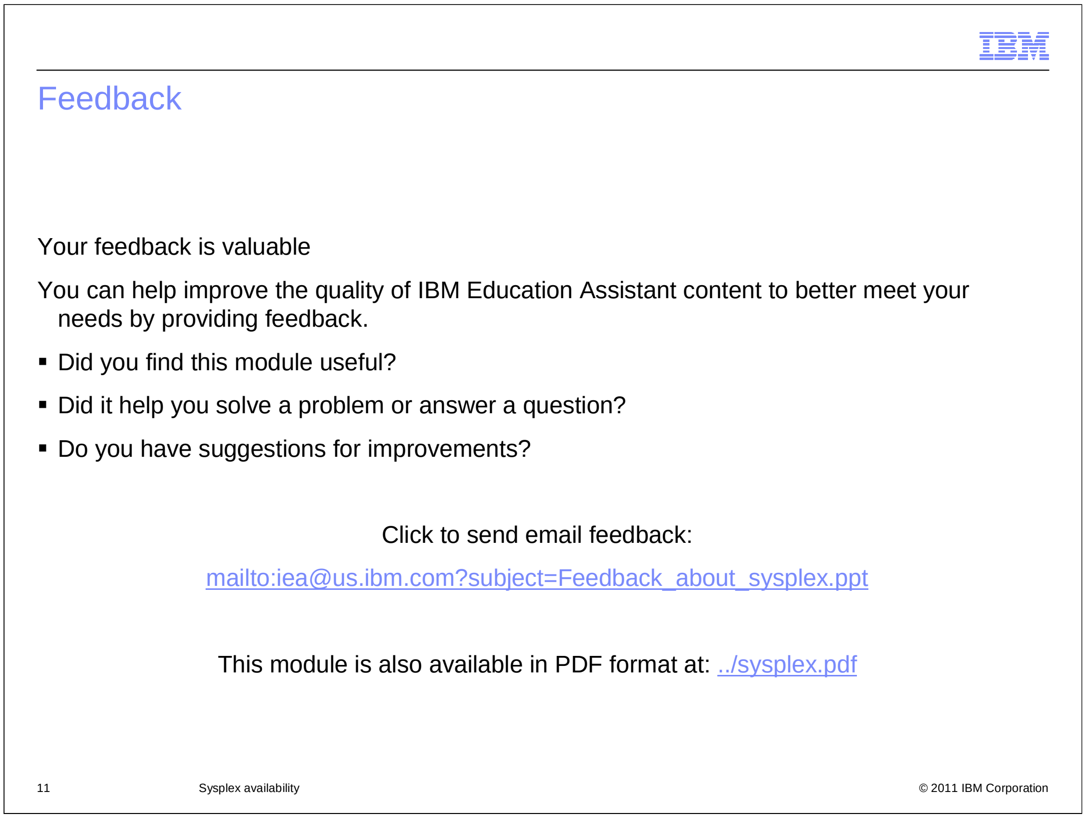

Feedback

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_sysplex.ppt

This module is also available in PDF format at: ../sysplex.pdf

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

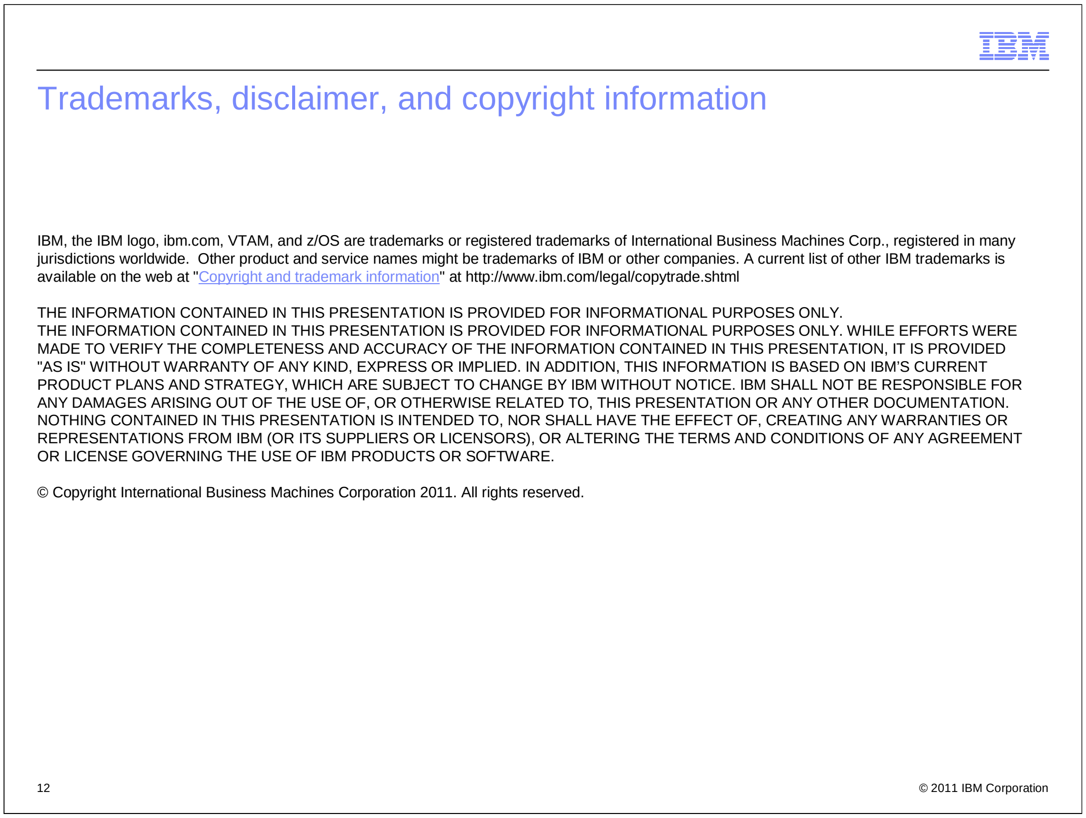

Trademarks, disclaimer, and copyright information

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.
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.