This presentation discusses the Policy enhancements in z/OS V1R9 Communications Server.

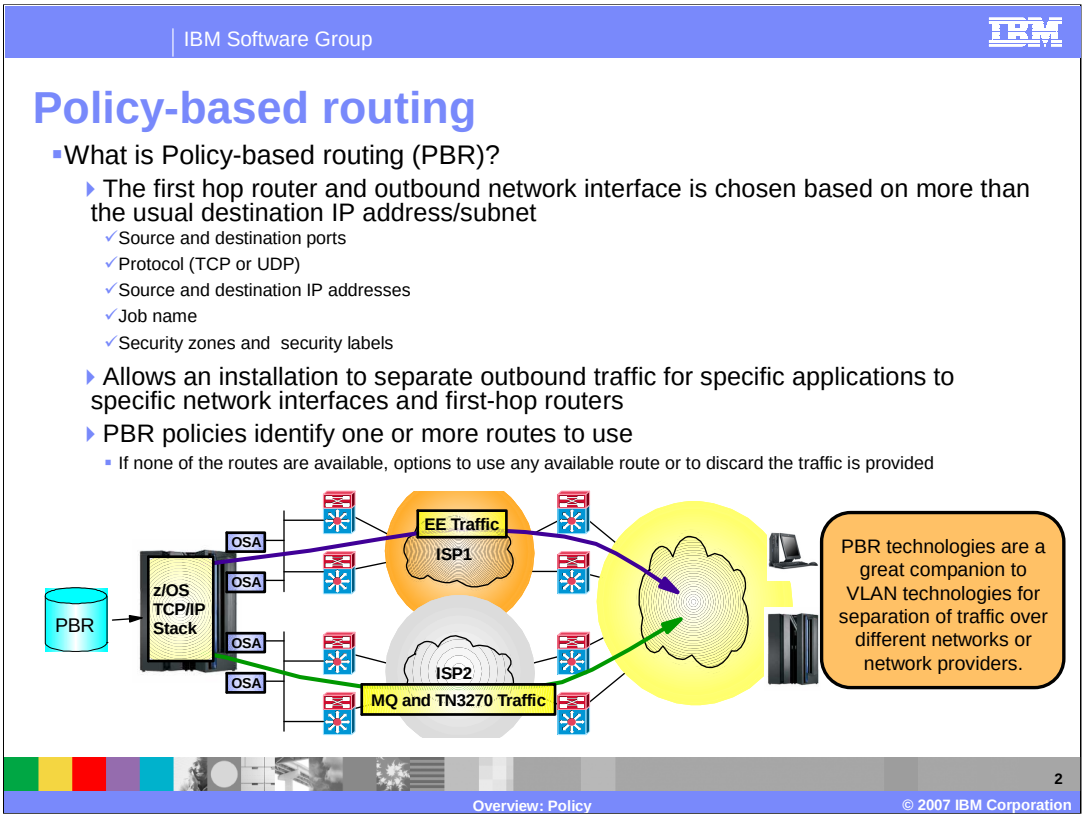

In previous releases, the TCP/IP stack did not provide a capability in routing decision that takes into consideration the type of applications, port numbers, protocol, source IP addresses and other criteria. In other words, the capability to select a suitable route (a network/interface) based on the type of applications, server/client IP addresses did not exist. Instead, it used a single metric in deciding which route to use when forwarding packets, namely the shortest path (or a static route) for a destination address. When multiple equal cost routes existed, the TCP/IP stack may use a round-robin mechanism to select among the equal cost routes, either on a per-packet or per-connection basis, but only if the appropriate multipath support is enabled on the stack.

Policy-based routing allows IP Routing to use additional route selectors. It is made possible through the use of multiple route tables. In addition to the main route table, the TCP/IP stack can now have multiple policy-based route tables. Policy-based route tables have many of the same characteristics as the main route table. They can contain both static and dynamic routes and the static routes can be configured as both replaceable and non-replaceable.

PBR allows an installation to separate outbound traffic for specific applications to specific network interfaces and first-hop routers. The separation can be based on security, choice of network provider, or isolation of certain applications.

Most often there will be one policy-based route table defined to be used for the traffic, but there may be as many as eight. Each of the policy-based route tables is searched, in the order defined, for a route to the destination. If any active route to the destination is found in a route table, the search is stopped and that route is used for the traffic. This route may be a host route, a subnet, network, or supernet route, or a default route. If no active route to the destination is found in a route table, the search continues with the next route table. If all policy-based route tables are searched without success, the main route table may also be searched if the policy indicates that the main route table can be used as a backup.

Policy-based routing is not supported for all types of IP traffic. The support is limited to locally originated IPv4 TCP and UDP traffic. All IPv6 traffic, all forwarded traffic, and all traffic using protocols other than TCP and UDP are not processed by policy-based routing and continues to be routed using only the main routing table.

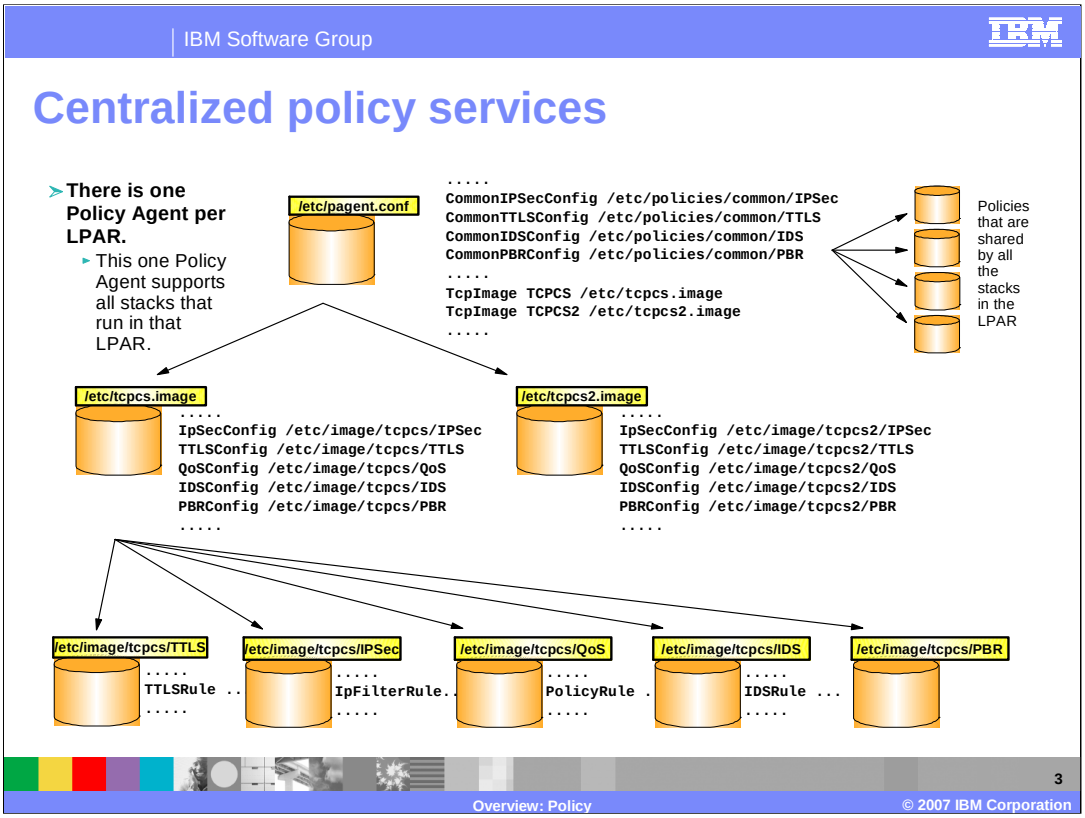

This slide graphically depicts the entire set of configuration files that can be used to define the different policy types, and the general Policy Agent configuration. It's important to understand that various subsets of the configuration files shown might be used, depending on the different policy types in use and the number of TCP/IP stacks supported by an instance of the Policy Agent.

When the Policy Agent is started the main configuration file is identified using a standard search order. This file in turn can point to one or more image configuration files using the TcpImage statement. Each image configuration file is used to configure policies for one TCP/IP stack. The image files can in turn point to image-specific files for the different policy types. The main configuration file can point to common files for all policy types except QoS. A given common configuration file applies to all TCP/IP stacks. This allows policy definitions that are not unique for each TCP/IP stack to be placed in the common file, and those that are unique to be placed in each image-specific file.

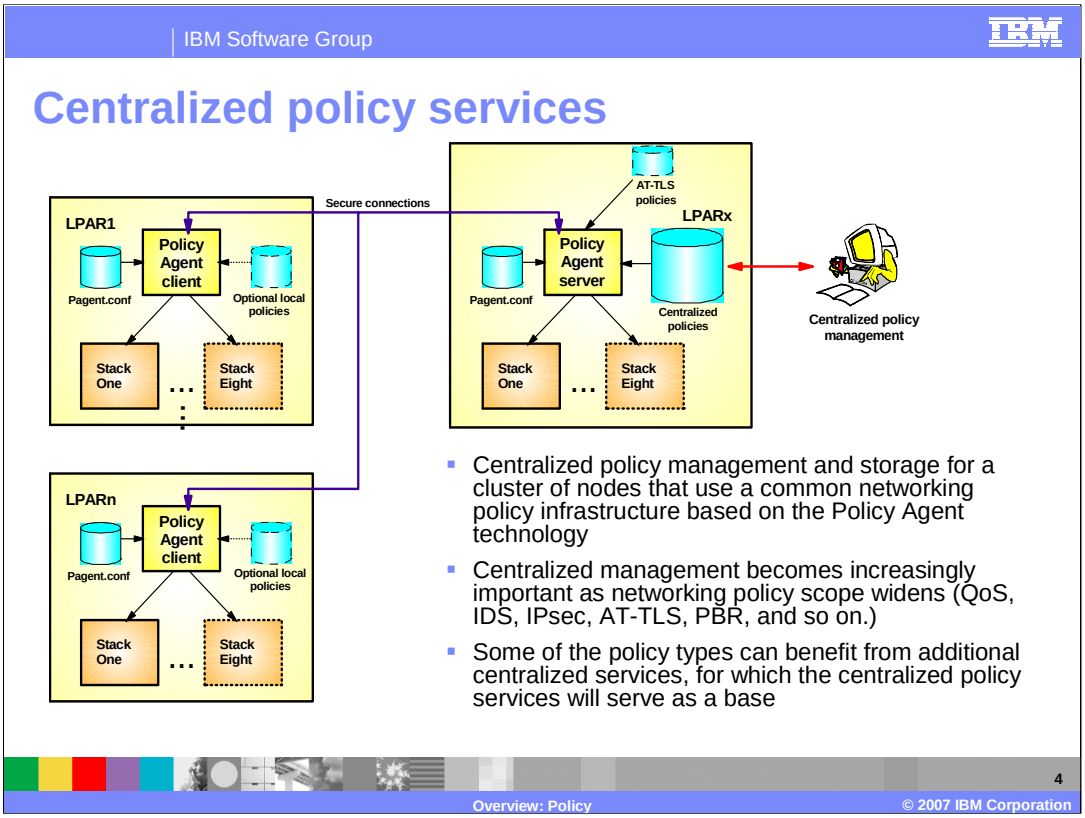

This picture shows an overview of the centralized policy services solution. On the left side are a number of policy clients. Each policy client can use local configuration file as usual, if needed. On the right side is the policy server. Centralized policies are defined, but are not installed in any TCP/IP stacks, on the policy server. These centralized policies are retrieved by the policy clients using the existing Policy Agent API (PAPI).

The IBM Configuration Assistant can be used to define the centralized policies, and local policies for the policy server and policy client (this is not shown).

To take full advantage of this solution, local policies should not be defined on the policy clients. The policy server is not itself considered a policy client, so local policies on the policy server are normal and expected.

The problem being solved by centralized policy services is primarily one of policy management. Each of the last several releases has introduced a new policy type, and the Policy Agent configuration needs to be replicated on each system. If the IBM Configuration Assistant for z/OS Communications Server is used to configure policy definitions, it also must be replicated on (or have connectivity to) each system.

Centralized policy services provides a centralized policy management and storage for a cluster of nodes that use a common networking policy infrastructure based on the Policy Agent technology. Initially a cluster of z/OS nodes is supported. However it can be extended to act as centralized networking policy server for heterogeneous nodes. Centralized management becomes increasingly important as networking policy scope widens (QoS, IDS, IPSec, AT-TLS, PBR, and so on.).

The Policy Agent is changed to take on the new roles of policy server and policy client. The policy server provides centralized policy administration and management for a set of policy clients. The policy client retrieves policies from the policy server. A single Policy Agent can be a policy client or policy server but not both.

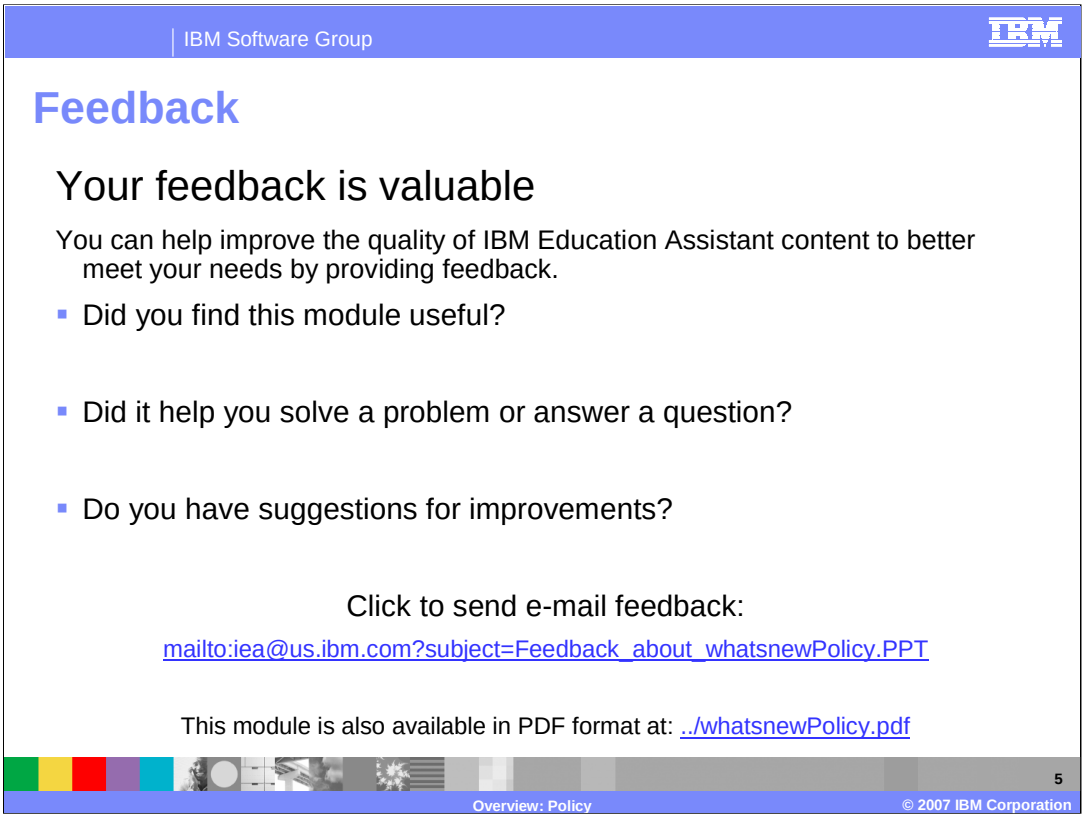

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

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

IBM z/OS

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. 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 in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.