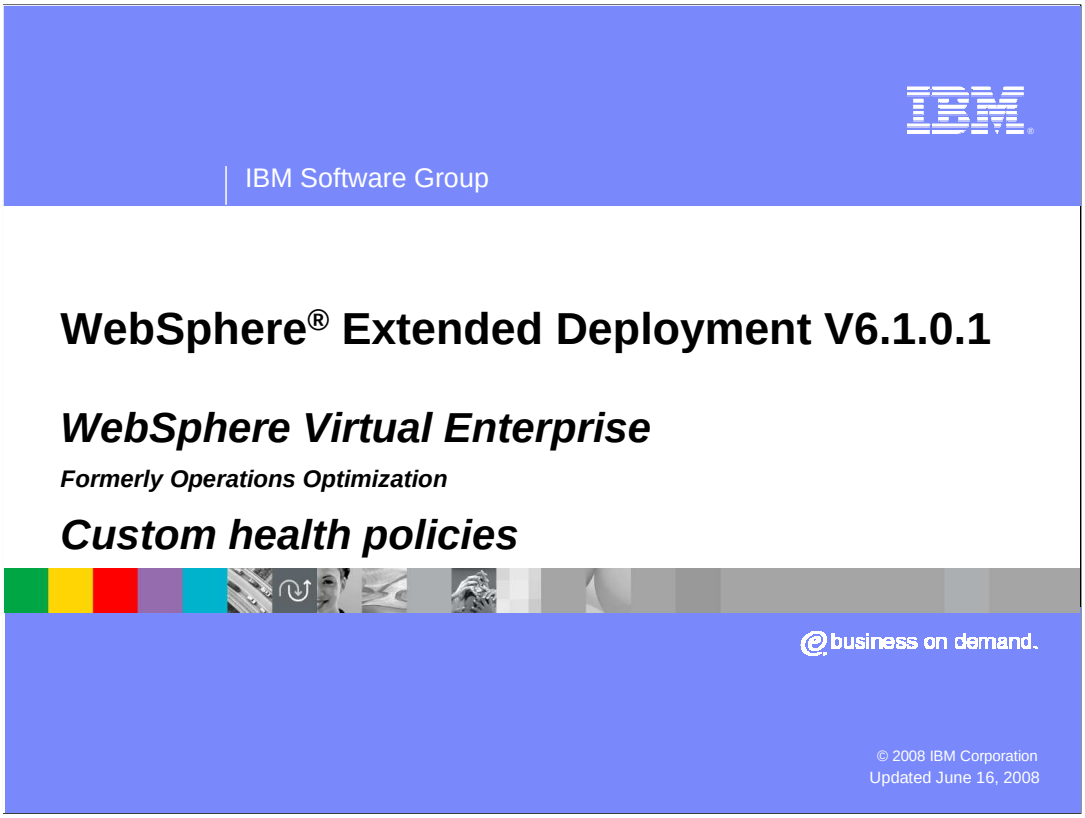

This presentation will cover WebSphere Extended Deployment's ability to monitor server health based on custom health policies.

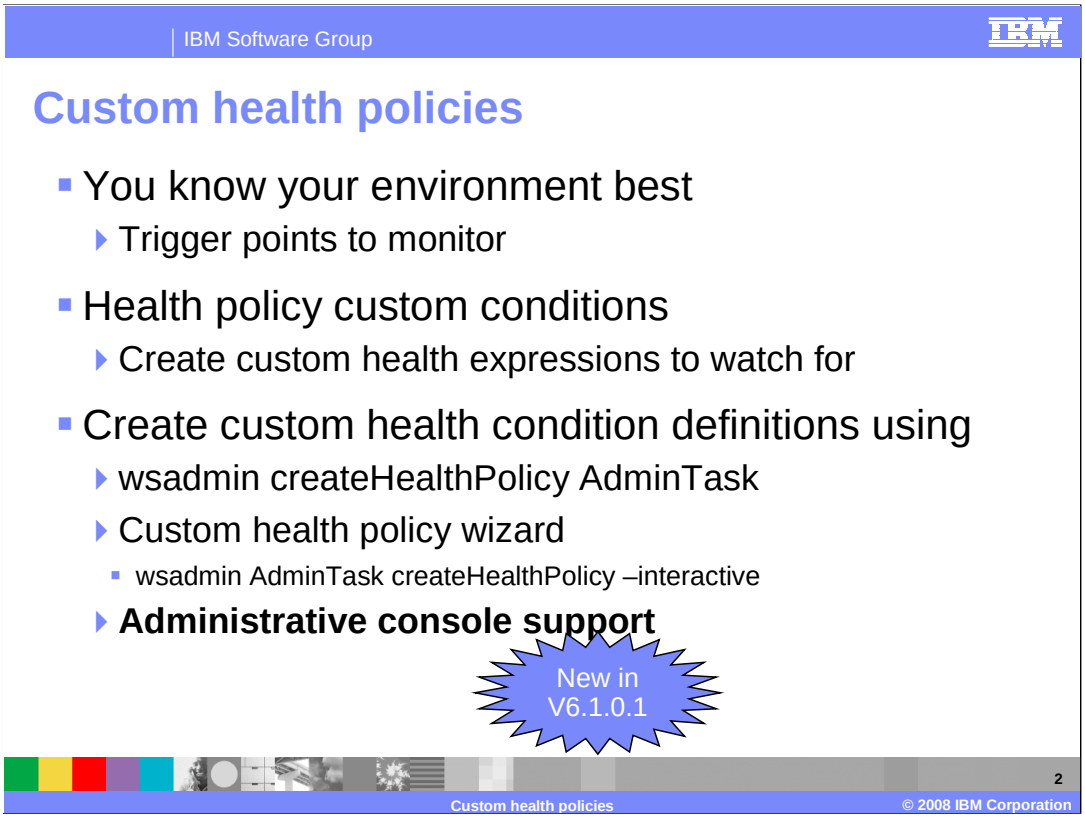

WebSphere Extended Deployment's health monitoring feature can monitor your environment for several common software health conditions, with the most commonly seen conditions predefined. Some examples are memory leaks, slow response times, or server age. You can use these health policy types by specifying a few simple parameters.

However, you might need to monitor your environment for conditions not supported by the default health policies, or require different levels of granularity.

WebSphere Extended Deployment version 6.1 introduces the concept of a custom health policy that allows you to define a custom condition when the predefined health conditions do not fit your needs. Using a custom health policy lets you create expressions to define what "unhealthy" means in your environment, rather than having WebSphere Extended Deployment define it. In version 6.1 you can create custom health condition definitions using the createHealthPolicy administrative task. Version 6.1.0.1 adds the ability to interactively define custom health policies from the administrative console.

Custom condition operands

- WebSphere Extended Deployment supported PMI modules
- On demand router metrics
  - server level
  - cell level
- MBean invocations
  - operation metric
  - attribute metric
- URL return code metric

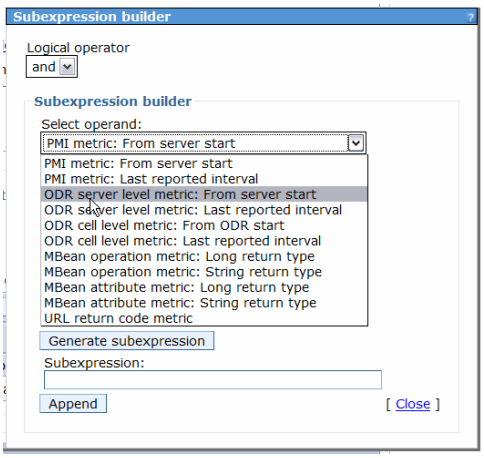

Custom health conditions can monitor metrics from a subset of PMI modules and server-level and cell-level metrics published by the on demand router. Custom health conditions can also invoke an MBean or query its attributes. In addition, you can ping any relative path, or Uniform Resource Identifier (URI), on the server that is the target of this policy.

Custom conditions are supported on all platforms, but with varying levels of support. For example, metrics from an on demand router can be used for a health policy for all server types. However, a custom health condition for a non-WebSphere server can not make use of the WebSphere PMI server metrics or MBeans.

PMI metric

- **PMI metric: From server start**
  - Uses an average number of the reported values from the time that the server started
- **PMI metric: From last reported interval**
  - Average of the reported values in the last interval
    - Interval is the length of the health controller cycle

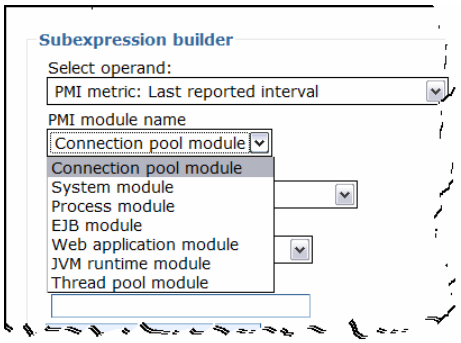

When monitoring PMI metrics from server start, the health controller maintains an average of the reported values from the time that the server being monitored was started.

When monitoring from the last interval, the health controller queries the average of the reported values since the last time the health controller ran.

You can further configure custom health conditions that use PMI modules at finer granularities than the server. For example, you can use the sub expression builder to create a webAppModule policy as a starting point, then edit the expression to monitor response time for an individual servlet.

Both PMI operand types support the same set of PMI modules.

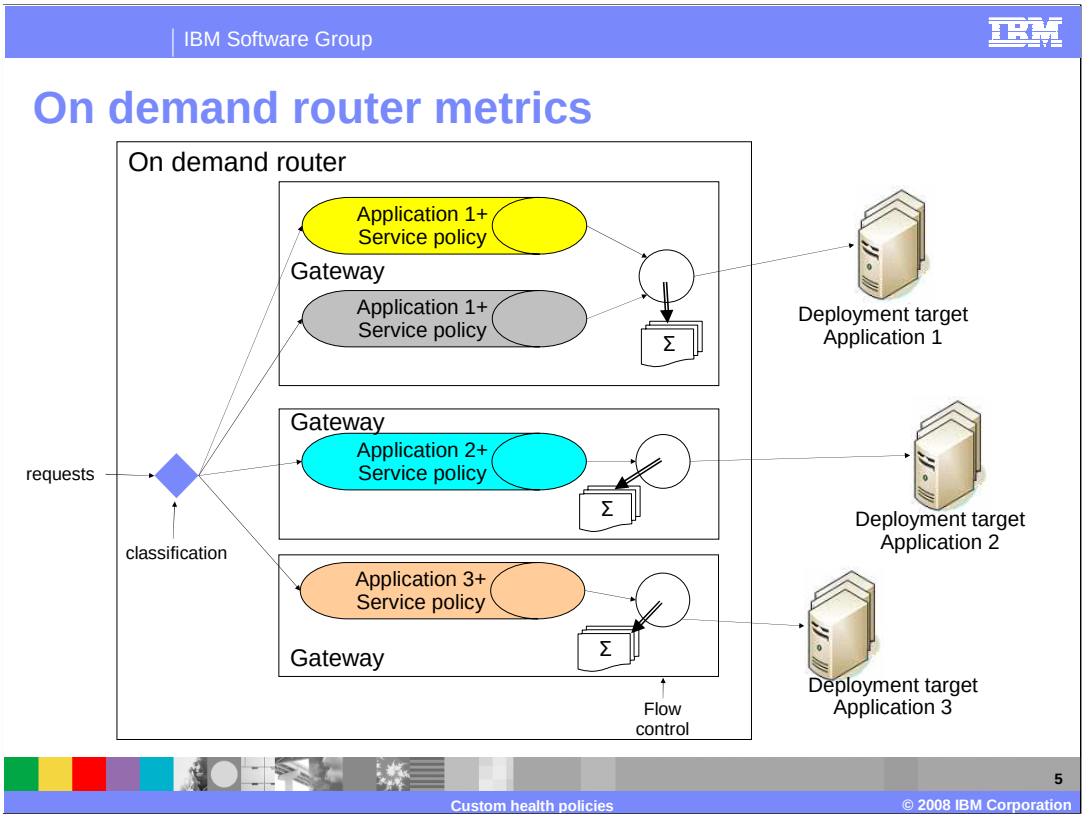

When a request enters an on demand router, it is matched to a quality of service goal represented by a service policy. The request is then placed in a queue associated with both the service policy and the deployment target of the application that will handle the request. This queue is managed by an Autonomic Request Flow Manager (ARFM) Gateway component within the on demand router. Each Gateway serves just one deployment target, which can be a cluster or an individual application server. A deployment target can be served by multiple Gateways. The Gateways maintain statistics on the queues they manage and requests flowing through them. The health controller can access many of these statistics and use them in custom health conditions.

ODR server metrics

- **ODR cell level metric:**
  - Subset of *cell level* metrics published by the On demand router
  - **From ODR start**
    - Metrics are cumulative and reported since the server start
  - **Last reported interval**
    - Average of the reported values in the last interval
    - Interval is the length of the health controller cycle

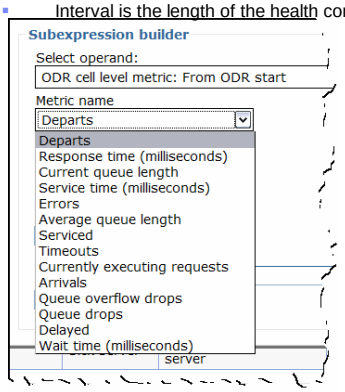

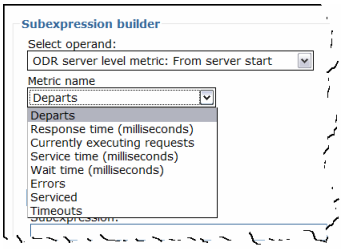

- **ODR server level metric:**
  - Subset of *server level* metrics published by the On demand router
  - **From server start**
    - Metrics are cumulative and reported since the server start
  - **Last reported interval**
    - Average of the reported values in the last interval
    - Interval is the length of the health controller cycle

The health controller can monitor these ODR-related metrics at several different levels of granularity. At the server level, the health controller aggregates low-level data from all ARFM gateways for requests as they are dispatched to individual back-end servers. For cell level metrics, the health controller aggregates the low-level data across all ARFM gateways in the cell. Cell level metrics are most useful for identifying cell-wide issues, for example a database bottleneck that causes increased queuing in the on-demand routers for all servers running an application.

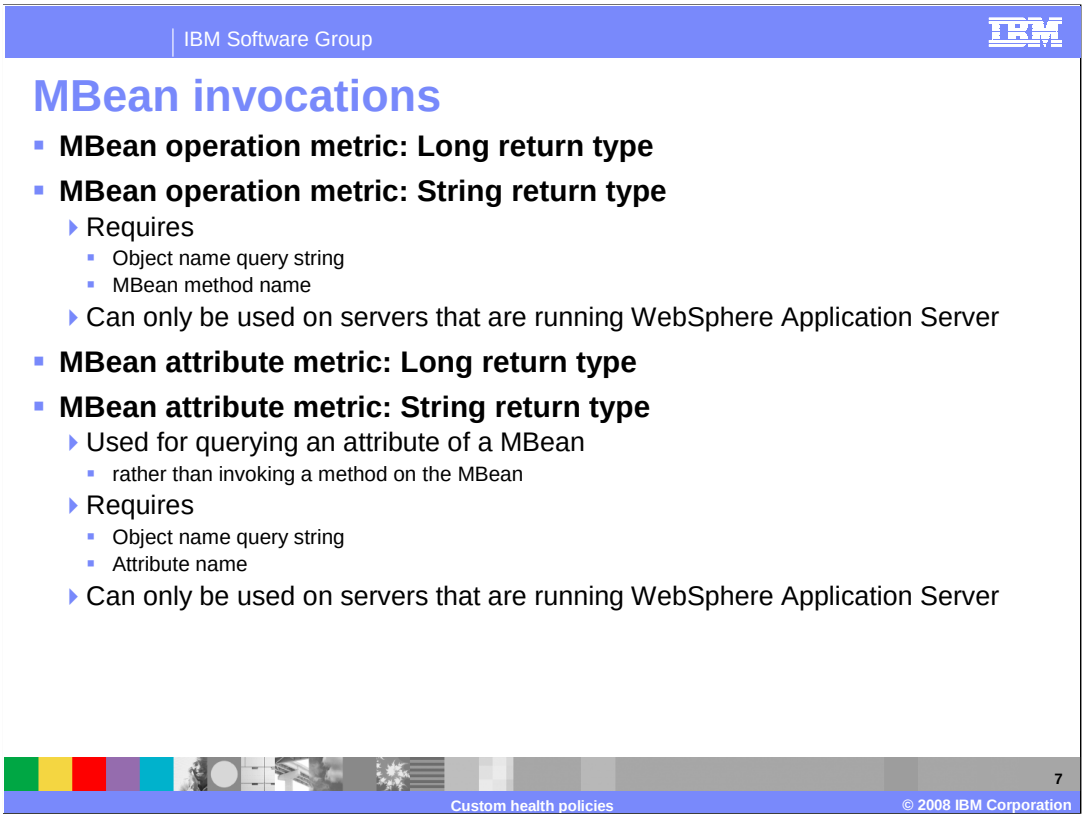

You can also create health conditions based on managed bean, or MBean, attributes, or the value returned from invoking a method on an MBean. To invoke an MBean method you must specify a query string to identify the MBean, and the method to invoke on that MBean. To access an MBean attribute you must specify the MBean query string and the attribute name.

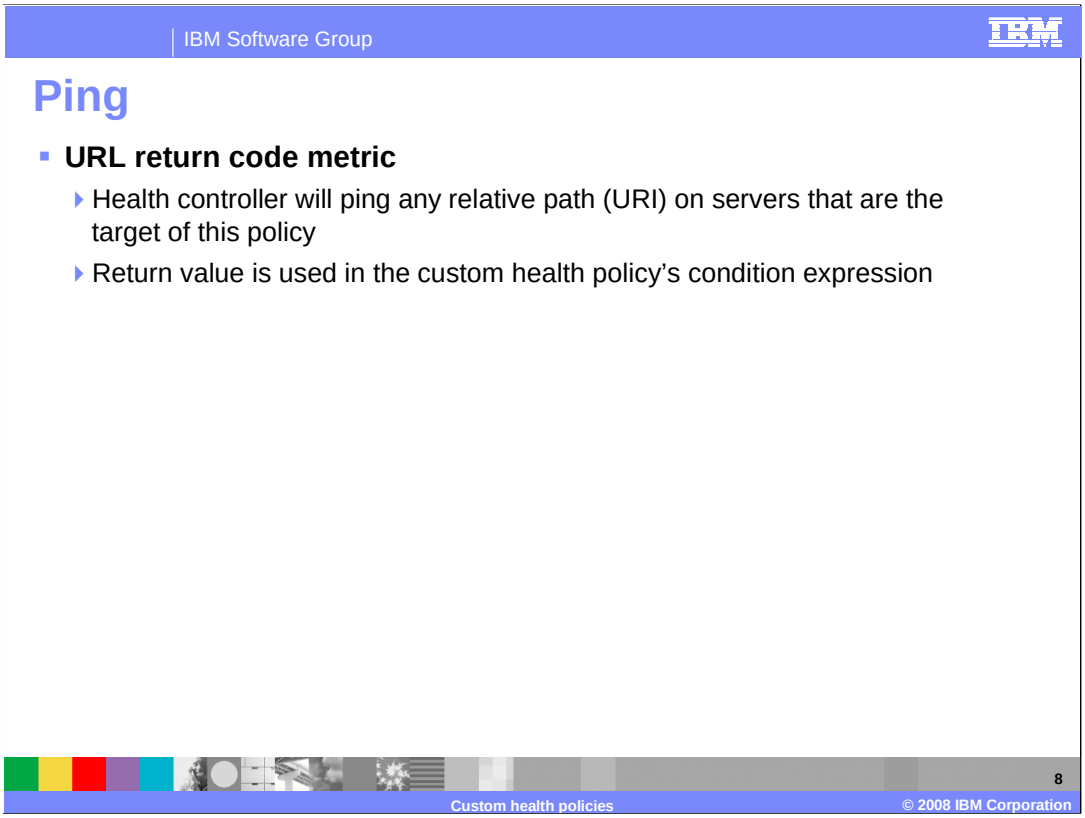

If you use the URL return code metric, the health controller will attempt to access the Universal Resource Identifier you specify on any server that is a target of this policy. The return value from the HTTP request is used in the condition expression for the custom health policy.

Custom health policy wizard

- The custom health policy wizard, or interactive mode, is invoked using wsadmin
  - JACL: $AdminTask createHealthPolicy {-interactive}
  - Jython: AdminTask.createHealthPolicy ('[-interactive]')
- Respond to the prompts to build the health policy
- Save your changes to the configuration
- Use the wsadmin AdminTask listHealthPolicies to verify that your new policy has been created

You can also create custom health policies through the administrative task "createHealthPolicy". If you invoke the task with the "-interactive" parameter it will guide you through the creation of the policy with a series of prompts.

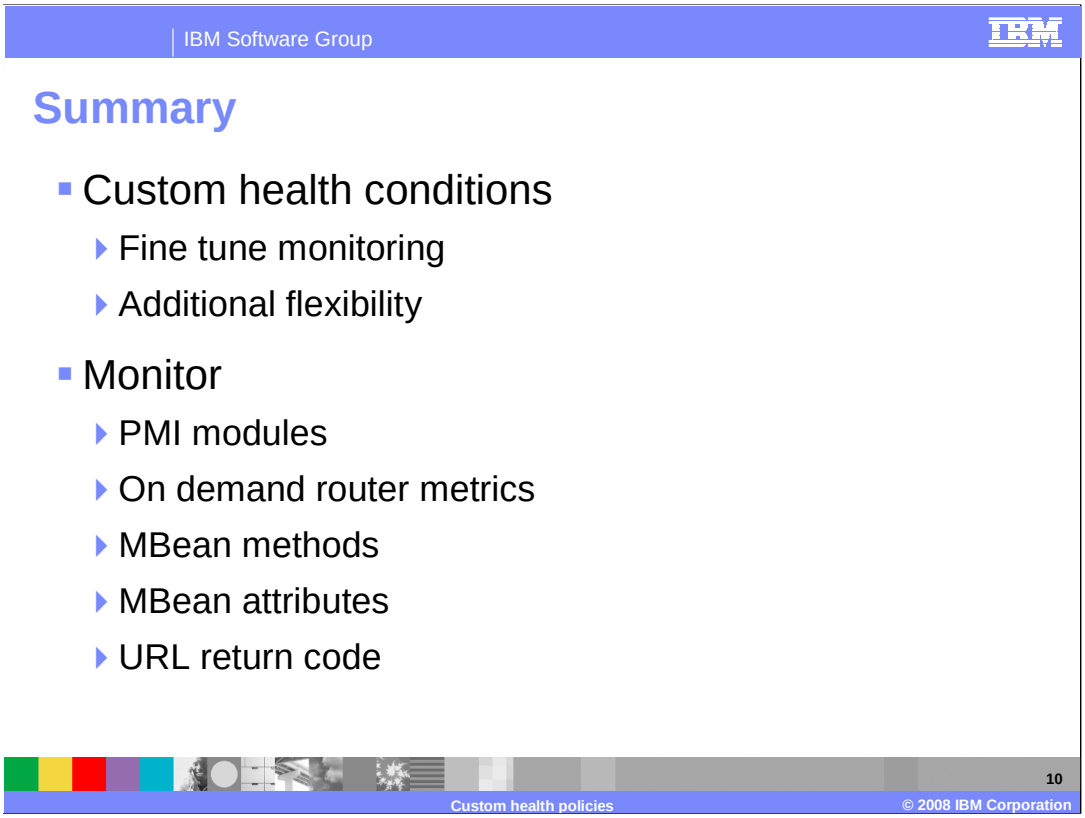

WebSphere Extended Deployment's health monitoring feature includes optimized algorithms for the most common health conditions, such as memory leak or slow response time. Custom health conditions allow you monitor PMI metrics and metrics maintained by the on demand routers. In addition, you can query managed beans and ping any Web address on a monitored server. This allows you to fine tune what you consider unhealthy in your environment. You can create Custom health policies with the createHealthPolicy administrative task or, starting in version 6.1.0.1, from the administrative console.

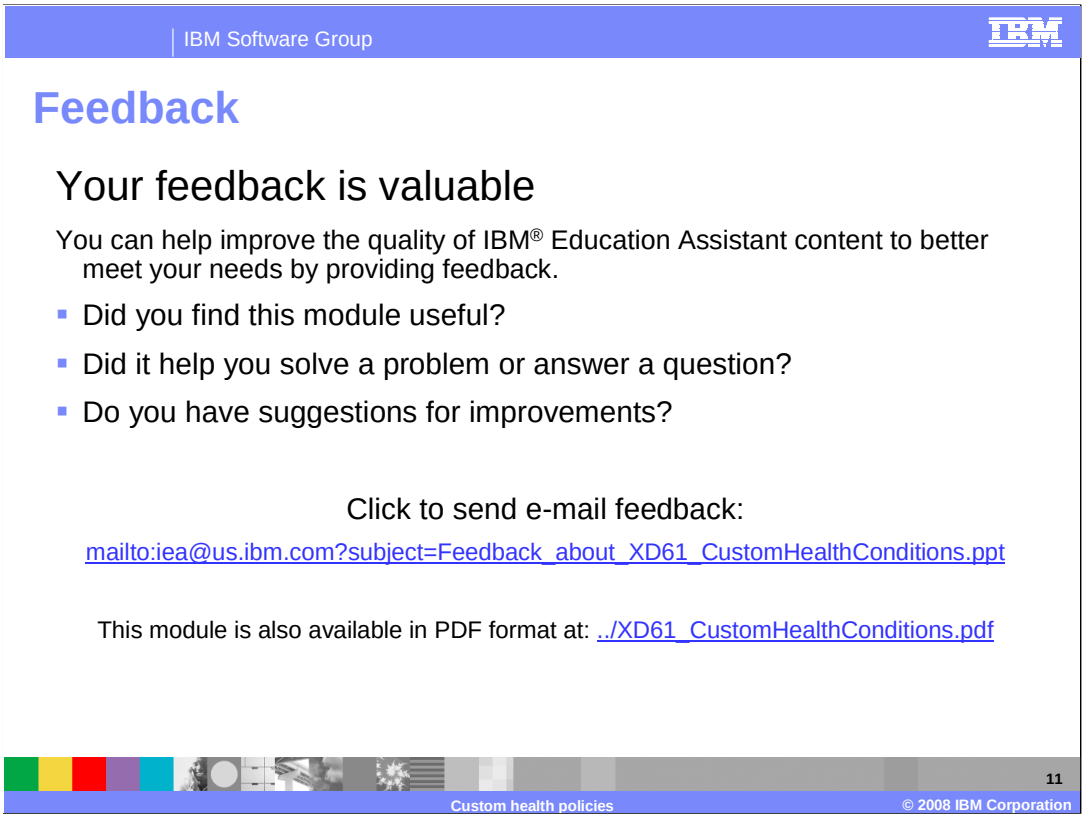

You can help improve the quality of IBM Education Assistant content by providing feedback.

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 WebSphere

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.