



IBM Software Group

IBM WebSphere Partner Gateway V6.1

Performance improvements



@business on demand.

© 2007 IBM Corporation
Converted to video August 5, 2014

This presentation provides an overview of the performance improvements in WebSphere® Partner Gateway V6.1

Goals and agenda

- Describe performance related changes
 - ▶ Caching
 - ▶ Alerts
 - ▶ File input/output operations reduction
- Processing of large files



The goal of the presentation is to provide you details on the new performance-related improvements like improved caching, changes in the way the alerts are setup, reduction in the input and output file operations, and processing of large files.

Caching of configuration data

- Document processing requires queries to database to retrieve configuration data.
 - ▶ Retrieved information is cached
- Cached data is updated
 - ▶ Using a change event from the console
 - ▶ Based on a time interval (in V6.0)
- WebSphere Partner Gateway V6.1
 - ▶ All caching is now event-based
 - ▶ Support caching for all processing steps
 - When Business Processing Engine reaches a steady state, no need for further DB query access
 - Alerts
- Still some database access
 - ▶ State engine
 - ▶ Activity logging



When processing a document, WebSphere Partner Gateway queries the database for configuration data specific to that document type. For example, what connection is being used and any attributes configured on that connection. All the data that is retrieved is cached and the cached data is updated when a change event is triggered from the community console or at fixed time intervals. In V6.1 all the caching is based on events. WebSphere Partner Gateway V6.1 supports caching of data required for all the processing. When the Business Processing Engine reaches a steady state, no further database queries are required. Although V6.1 minimizes the database queries it does not completely eliminate them; you still need to access the database for state engine information and for activity logging.

Alerts

- In previous releases, every event generated during document processing is sent to the Alert engine
 - ▶ Example: AS2 pass-through generates 4 events
 - ▶ 10000 docs and 10000 MDNs mean 80000 events
 - ▶ High traffic on queue
 - ▶ Alert engine queries DB to see if event is alertable

- For WebSphere Partner Gateway V6.1
 - ▶ All events are non-alertable by default
 - The reverse of previous releases
 - Manually set an event to alertable
 - ▶ Only alertable events are forwarded to alert engine
 - Reduce queue traffic and database queries by orders of magnitude



Previously all events were sent to the Alert engine. The Alert engine in turn did a database query on each event to see if there was an alert defined for that event.

WebSphere Partner Gateway V6.1 makes all events not alertable by default. The administrator needs to specifically set an event as alertable.

This means that only events that truly have alert subscriptions will be sent and processed by the Alert engine.

This reduces the queue traffic and number of database queries significantly.

File input/output operations reduction

- In previous releases, when a document is processed
 - ▶ 3 files are created:
 - Payload (.vcm) file
 - Transport header (.vhd) file
 - Metadata (.vmd) file
 - ▶ These files are also moved and copied multiple times

- WebSphere Partner Gateway V6.1
 - ▶ Handles the .vmd and .vhd file data as database data



WebSphere Partner Gateway V6.1 reduced the number of files that are written to the file system. In previous versions, document processing required the creation of the .vcm, .vhd and .vmd files for payload, transport header and metadata information of the documents. These three files are copied and moved to several locations in the file system as the different components of WebSphere Partner Gateway process the document. In V6.1 the data in these three files is made part of the database and this reduces the file operations and reduces the time taken for the disk access.

File input/output operations reduction (2)

- Decreased file system requirements
 - ▶ Observed percentage decreased in number of files created from 65% (AS2 protocol) to 75% (RosettaNet)
- Decreased inodes requirements
- Reduction in disk access contention



Handling .vmd and .vhd files as database data means a reduction in file system requirements. A 65% reduction in the number of files created is observed for AS2 and 75% for RosettaNet. Fewer files created and accessed from the file system also increases the performance due to reduction in disk access contention.

Section

Processing of large files



The next section will discuss the limitations of processing large files in WebSphere Partner Gateway V6.1

Limitations

- Encrypted, signed documents
 - ▶ Can consistently process documents of about 19 MB
 - ▶ IBM Public-Key Cryptography Standards package requires the entire document to be in memory
- ebXML
 - ▶ Can handle documents up to 1.5 gigabytes
- EDI
 - ▶ Transformation/Validation, up to 15 MB
 - ▶ Splitter, up to 70 MB
- RosettaNet Implementation Framework
 - ▶ Validation/Transformation, up to 30 MB



WebSphere Partner Gateway can process a file of up to 19 megabytes when encrypted and signed. The IBM Public-Key Cryptography Standards package requires the entire document to be in memory. When using ebXML, WebSphere Partner Gateway can handle file sizes up to 1.5 gigabytes. For EDI documents, WebSphere Partner Gateway can consistently process 15 MB documents when using the validation and transformation functions, and the splitter can handle files up to 70 MB. For Rosetta Net documents, WebSphere Partner Gateway can validate and transform data up to 30 MB.

Observations

- Most large documents limitations are due to having to load the entire document into memory
 - ▶ Consider increasing the Java™ heap size for the Document Manager
 - Windows operating systems limit each process to about 2 gigabytes of memory
 - Take advantage of greater memory addressability on 64-bit systems
 - ▶ Use compression
 - ▶ Configure the number of threads using by WebSphere Partner Gateway components
 - There might not be enough memory to process multiple documents in parallel
 - Serialize document processing by reducing the number of threads



Most large documents limitations are due to having to load the entire document into memory. You can optimize WebSphere Partner Gateway by reducing the memory footprint by using compression or processing fewer documents in parallel. Another option is increasing the Java heap size.

Section

Summary

The next section covers the summary and references.

Summary

- WebSphere Partner Gateway V6.1 performance improvements
 - ▶ Reduces file input/output operations by moving .vmd, .vhd files data to the database
 - ▶ Reduces database access by caching configuration data and events



In summary, this presentation covered details of performance improvements in V6.1, which include caching of all the data required for processing the documents and reducing the file input and output operations. The limitations when using large documents were also discussed.

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

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

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

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.

© Copyright International Business Machines Corporation 2007. All rights reserved.

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.