



Universal CMDB

Software Version: 10.33

Data Flow Management Best Practices

Document Release Date: July 2017

Software Release Date: July 2017



Legal Notices

Disclaimer

Certain versions of software and/or documents ("Material") accessible here may contain branding from Hewlett-Packard Company (now HP Inc.) and Hewlett Packard Enterprise Company. As of September 1, 2017, the Material is now offered by Micro Focus, a separately owned and operated company. Any reference to the HP and Hewlett Packard Enterprise/HPE marks is historical in nature, and the HP and Hewlett Packard Enterprise/HPE marks are the property of their respective owners.

Warranty

The only warranties for products and services of Micro Focus and its affiliates and licensors ("Micro Focus") are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Except as specifically indicated otherwise, a valid license from Micro Focus is required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© 2002 - 2017 Micro Focus or one of its affiliates.

Trademark Notices

MICRO FOCUS and the Micro Focus logo, among others, are trademarks or registered trademarks of Micro Focus (IP) Limited or its subsidiaries in the United Kingdom, United States and other countries. All other marks are the property of their respective owners.

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

To check for recent updates or to verify that you are using the most recent edition of a document, go to: <https://softwaresupport.softwaregrp.com>.

This site requires that you register for a Software Passport and to sign in. To register for a Software Passport ID, click **Register for Software Passport** on the Micro Focus Support website at <https://softwaresupport.softwaregrp.com>.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your Micro Focus sales representative for details.

Support

Visit the Micro Focus Support site at: <https://softwaresupport.softwaregrp.com>.

This website provides contact information and details about the products, services, and support that Micro Focus offers.

Micro Focus online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up Micro Focus support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as a Software Passport user and to sign in. Many also require a support contract. To register for a Software Passport ID, click **Register for Software Passport** on the Micro Focus Support website at <https://softwaresupport.softwaregrp.com>.

To find more information about access levels, go to: <https://softwaresupport.softwaregrp.com/web/softwaresupport/access-levels>.

Integration Catalog accesses the Micro Focus Integration Catalog website. This site enables you to explore Micro Focus Product Solutions to meet your business needs, includes a full list of Integrations between Micro Focus Products, as well as a listing of ITIL Processes. The URL for this website is <https://softwaresupport.softwaregrp.com/km/KM01702731>.

Contents

Data Acquisition Best Practices with Data Flow Management	4
Federation or Population and Discovery	4
When to Use Federation or Population	4
Conduct an MDR Capacity Plan	6
Determine a Data Acquisition Strategy	6
Hybrid Data Acquisition Approach	7
Send documentation feedback	9

Data Acquisition Best Practices with Data Flow Management

Federation or Population and Discovery

This section includes the following topics:

When to Use Federation or Population	4
Conduct an MDR Capacity Plan	6
Determine a Data Acquisition Strategy	6
Hybrid Data Acquisition Approach	7

When to Use Federation or Population

In many Configuration Management Systems, configuration data is scattered over multiple Management Data Repositories (MDR) around the organization. Each MDR is part of a management application provided by an IT system vendor. MDRs may include very rich configuration data for the systems they manage. In the most optimistic scenario, a UCMDB instance does not copy or 'populate' any of the data stored in the MDRs, rather the UCMDB gains access to all CI data via federation.

In some cases, a pure federated approach might be possible, but in other cases, this scenario becomes impractical. The following issues may contribute to the poor performance of a distributed Configuration Management System:

- It is not clear whether all relevant MDRs are available at the time of TQL query execution, particularly if some of the MDRs lack High Availability capability. If an MDR which is necessary for the execution of a query is not available, UCMDB fails to return results for the query because the completeness and correctness of the query result is unknown.
- The response time of the execution of such a federated TQL query may depend on slow MDRs and WAN interfaces.
- When a group of UCMDB users try to access a federated data source from MDRs scattered over multiple, distant, geographical locations, it can overload the network (WAN) of a distributed organization. This is especially true for large groups of users such as a group of Service Manager (SM) Incident Manager users.

- Large organizations often have a large group of CMS or SM users. Their MDRs may become overloaded by UCMDB since they are usually built to serve a limited set of administrative users. Some of the MDRs may be vital for the organization. For example, if an LDAP server is used as a federated data source for personal data, there is a danger that the server will be overloaded by federated queries. Applications using the LDAP server for authentication requests could be delayed or completely blocked.

Federation vs Population Factors

It is often necessary to populate some portion of the MDRs' data into UCMDB instances. The balance of the data can be federated. The percentage of populated versus federated data is dependent on many factors such as:

- **Dynamic nature of the data.** If the data in the MDR is updated frequently and the changes need to be visible to the user, then federation is preferred. Examples of this could include the failover of a Cluster Resource Group or the v-motion of a virtual machine. In these cases, the dynamic component is the relationship that connects the VM or CRG to the physical node or cluster node and only this data should be federated. It is not necessary to federate all the detailed data about the node, VM or CRG, only the reconciliation data for these entities and the relationship that connects them.
- **Importance of history.** If you need to store history of dynamic data, then federation is not the correct approach as history is not captured for federated data. If the MDR can publish change events, consider migrating these change events into UCMDB via an integration infrastructure such as your organization's Enterprise Service Bus (ESB). The ESB integration component could use the UCMDB Java SDK to push the change events into UCMDB.
- **Size of the user group that consumes the data.** If the user group that consumes the data via federated queries is large, the MDRs may be overloaded with too many queries. In this case, population of a higher percentage of the data is preferred.
- **Scope of federated queries and MDR response time.** If a federated TQL query needs access to many data sources to execute, then the response time of the query may be long. Replication is preferred, especially when the MDR response time is slow.
- **High Availability of MDRs.** If the MDRs are not highly available but the data they store is required for processes with HA requirements, then population is preferred because UCMDB is highly available. This is because if the data is federated and one MDR needed for the execution of a TQL query is not available, the UCMDB fails to return the query results even if other MDRs necessary for the query execution are available.

- **Network bandwidth.** If frequent access to remote MDRs involves access via an expensive WAN, then the percentage of populated data from such MDRs should be increased.

Conduct an MDR Capacity Plan

In general, before you decide to federate data from an MDR, you should conduct a thorough capacity planning process, answering the following questions:

- How many users should an MDR serve before it becomes a Federated Data Source?
- How many user requests will an MDR receive after becoming a Federated Data Source?
- Does the MDR have enough resources to support federated requests: CPUs, Memory/Cache, IO Bandwidth?
- Is it possible to optimize the MDR for federated queries, for example, by adding an index for a column of an MDR table?

Determine a Data Acquisition Strategy

Whenever an organization needs configuration data for a new IT Domain, the CMS administrator needs to determine the strategy for data acquisition:

- Data for the new IT Domain can be acquired from an MDR via federation or population.
- Data for the new IT Domain can be collected via Discovery (via direct access to managed entities).

For example, an organization decides to initiate a configuration management process for all Oracle instances. The CMS administrator needs to decide how the configuration information for Oracle instances will be acquired. The following are the available options:

- Federate or populate data from the MDR of the Oracle Management Application which maintains information regarding all Oracle instances in the organization.
- Acquire the configuration data via Discovery, via direct access to each Oracle instance (by running one of the Oracle discovery jobs in DFM).

Note: The UCMDB Configuration Management Application cannot manage compliancy for federated data. If compliancy is required for all data, then federation is not an option. The UCMDB

Configuration Management Application can only manage compliancy for CIs which are physically stored in the UCMDB.

Acquiring Data via an MDR Using Federation

The following are examples of cases where collecting data via access to an MDR is preferable:

- There is a relevant MDR whose level of detail and data accuracy meet the requirements of the configuration management process.
- The owners of managed entities refuse to provide the credentials of the managed entities.
- Direct access to the managed entities is limited to times when the managed entities are not in use by the production environment. Access to the MDR is not limited to these times and is more flexible.

Acquiring Data via Discovery (Direct Access to Managed Entities)

The following are examples of cases where collecting data via direct access to managed entities is preferable:

- There is no relevant MDR.
- It is not possible to trust the data accuracy of the MDR or the detail level of the data in the MDR does not meet the requirements for the implementation of the ITIL process.

Hybrid Data Acquisition Approach

In many cases, dependency data is not maintained by an MDR. For example, a management application of a messaging solution such as IBM MQ might include information about any MQ-Server in the organization and the queues it realizes. However, the MDR of the MQ management application does not include information about business applications that use any of the MQ Servers. To discover dependency data of a business application with MQ, you need to write a specific discovery pattern that extracts this configuration data from each specific application. A Federation or Replication Adapter should be written to federate or populate the detailed configuration data from an MDR. In addition, a

Discovery Adapter should be written to access the managed entities directly and extract dependency data.

If the MDR does not include the required level of details to support the necessary ITIL process, you can select one of the following approaches:

- Access both the MDR and the managed entities. First, import the initial, non-detailed data from the MDR and then access managed entities directly to import more detailed configuration data.
- Do not access the MDR at all. Import all the data from the managed entities.

Send documentation feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Data Flow Management Best Practices (Universal CMDB 10.33)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to cms-doc@microfocus.com.

We appreciate your feedback!