

HP Universal CMDB

for the Windows and Solaris operating systems

Software Version: 8.03

HP Universal CMDB–Storage Essentials (SE) Integration Guide

Document Release Date: September 2009

Software Release Date: September 2009



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP 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. Valid license from HP 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 Notices

© Copyright 2005 - 2009 Mercury Interactive (Israel) Ltd.

Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Intel®, Pentium®, and Intel® Xeon™ are trademarks of Intel Corporation in the U.S. and other countries.

Java™ is a US trademark of Sun Microsystems, Inc.

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

Oracle® is a registered US trademark of Oracle Corporation, Redwood City, California.

Unix® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates, or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign-in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

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

Support

Visit the HP Software Support web site at:

<http://www.hp.com/go/hpsoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software 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 web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP 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 an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Table of Contents

- Chapter 1: Storage Essentials (SE) Integration
with HP Universal CMDB7**
- SE Integration – Overview8
- Discover the SE Oracle Database9
- The Storage Essentials Integration Packages11
- Discovered CITs12
- Views.....17
- Correlation Rules21
- Reports23

Table of Contents

1

Storage Essentials (SE) Integration with HP Universal CMDB

This chapter explains how to integrate Storage Essentials (SE) with HP Universal CMDB (UCMDB).

This chapter includes:

Concepts

- ▶ SE Integration – Overview on page 8

Tasks

- ▶ Discover the SE Oracle Database on page 9

Reference

- ▶ The Storage Essentials Integration Packages on page 11
- ▶ Discovered CITs on page 12
- ▶ Views on page 17
- ▶ Correlation Rules on page 21
- ▶ Reports on page 23

SE Integration – Overview

Integration involves synchronizing devices, topology, and the hierarchy of a customer storage infrastructure in the UCMDB database (CMDB). This enables Change Management and Impact Analysis across all business services mapped in UCMDB from a storage point of view.

You integrate SE with UCMDB using the Discovery and Dependency Mapping (DDM) application.

When you activate the **Integration – Storage Essentials** module, DDM retrieves data from the SE Oracle database and saves CIs to the UCMDB database. Users can then view SE storage infrastructure in UCMDB.

The data includes information on storage arrays, fiber channel switches, hosts (servers), storage fabrics, logical volumes, host bus adapters, storage controllers, and fiber channel ports. Integration also synchronizes physical relationships between the hardware, and logical relationships between logical volumes, storage zones, storage fabrics, and hardware devices.

Note: DDM version 8.00 includes a module for discovering SE. No additional deployment is necessary.

Supported Versions

The integration procedure supports DDM version 8.00 or later and SE version 6.x.

SE Installation Requirements

The minimum VM installation requirements for SE integration are:

- 4 GB memory
- 50 GB hard drive space

Discover the SE Oracle Database

This task includes the steps to run the SE/UCMDB integration jobs.

This task includes the following steps:

- “Prerequisites” on page 9
- “Network and Protocols” on page 10
- “Activate the Discovery Job” on page 10
- “Topology Map” on page 11

1 Prerequisites

In DDM, in the Run Discovery window, verify that the following CIs have been discovered before running the SE discovery:

- **Network – Basic > Class C IPs by ICMP or Range IPs by ICMP:** discovers the IP address of the Oracle database server
- **Database – Oracle > Database TCP Ports:** discovers TCP ports on the IP address discovered previously
- **Database – Oracle > Oracle Connection by SQL:** discovers Oracle server instances
- **Network – Storage > SE Integration by SQL:** discovers storage infrastructure

For details on activating a job, see “Discovery Modules Pane” in *Discovery and Dependency Mapping Guide*. For an explanation of a discovery job, see “Jobs” in *Discovery and Dependency Mapping Guide*.

Note:

- ▶ For the **Oracle Connection by SQL** job, it is recommended to use the **REPORT_USER** Oracle user name, since this user has privileges necessary to run SQL queries on the APPIQ_SYSTEM tables.
 - ▶ This discovery pattern queries Oracle Materialized Views, and the views may be in the process of being refreshed when the discovery job is executed. This could result in an error message identifying the problem and a request to run the pattern later.
-

2 Network and Protocols

SE uses the **SQL protocol**. For details, see “SQL Protocol” in *Discovery and Dependency Mapping Guide*.

3 Activate the Discovery Job


The **SE Integration by SQL** job is included in the **Integration – Storage Essentials** module.

The **SE Integration by SQL** job runs queries against Oracle materialized views, installed and maintained by Storage Essentials in the Oracle database. The job uses a database CI as the trigger.

For details on activating a job, see “Discovery Modules Pane” in *Discovery and Dependency Mapping Guide*.

4 Topology Map

The following diagram illustrates the CITs discovered by the **SE Integration by SQL** job:



The Storage Essentials Integration Packages

The integration includes two UCMDDB packages:

- ▶ **SE_Discovery.zip**. Contains the trigger TQL for SE discovery, discovery script, discovery pattern, and discovery job. The discovery pattern has no parameters and requires no configuration.
- ▶ **Storage_Basic.zip**. Contains the new CI Type definitions, views, reports, and correlation rules. This package is common to all Storage Management integration solutions.

Tip: You can include the SE job in the DDM schedule. For details, see “Discovery Scheduler Dialog Box” in *Discovery and Dependency Mapping Guide*.

Discovered CITs

The following CITs represent SE storage entities in UCMDB:

- ▶ **Fiber Channel Connect.** This CIT represents a fiber channel connection between fiber channel ports.
- ▶ **Fiber Channel HBA.** This CIT has change monitoring enabled on parameters such as state, status, version, firmware version, driver version, WWN, and serial number. A Fiber Channel HBA inherits from the Host Resource CIT.
- ▶ **Fiber Channel Port.** This CIT has change monitoring enabled on parameters such as state, status, WWN, and trunked state. Since a Fiber Channel Port is a physical port on a switch, it inherits from the Physical Port CIT under the Network Resource CIT.
- ▶ **Fiber Channel Switch.** A switch falls under the Host CIT since SE maintains an IP address for each switch. Parameters such as status, state, total/free/available ports, and version are change monitored.

This package retrieves Fiber Channel Switch details from the `mvc_switchsummaryvw` and `mvc_switchconfigvw` views. The discovery retrieves detailed information about Fiber Channel Ports on each switch from the `mvc_portsummaryvw` view.

A switch inherits from a Host CI Type in UCMDB. Since UCMDB uses the IP address of a host as part of its primary key, this discovery pattern attempts to use an IP address from SE for this purpose. If an IP address is not available, the discovery job attempts to resolve the switch’s IP address using a DNS name (also maintained by SE). If neither an IP address nor a DNS name is available, the switch is discarded.

- ▶ **Logical Volume.** This CIT represents volumes on Storage Arrays and hosts with change monitoring on availability, total/free/available space, and storage capabilities.
- ▶ **Storage Array.** This CIT represents a Storage Array with change monitoring on details such as serial number, version, and status. Since a storage array may not have a discoverable IP address, it inherits from the Network Resource CIT.


This package retrieves Storage Array details from the `mvc_storagesystemssummaryvw` view. The discovery pattern retrieves detailed information on Storage Processors and HBAs from the `mvc_storageprocessorssummaryvw` and `mvc_cardsummaryvw` tables respectively.

The SE database may possibly not be able to obtain IP address information on Storage Arrays for a variety of technical and policy related reasons. Since a Storage Array is a host as far as UCMDB is concerned, the discovery assumes that the serial number of a Storage Array is unique and uses this as the primary key. The CI is then manually set as a complete host. If the serial number of a Storage Array is not available, the array is discarded.

Since Fiber Channel Ports may be present on a Storage Array, Storage Processor, or HBA, the discovery pattern uses three separate queries to retrieve Fiber Channel Ports for each Storage Array. Detailed information about Fiber Channel Ports on each array are retrieved from the `mvc_portsummaryvw` view. Since this view uses a container ID as the key, the discovery pattern queries the view by container ID for each Storage Array, each Storage Processor on a Storage Array, and each HBA on a Storage Array.

The discovery pattern retrieves detailed information about Logical Volumes on each Storage Array from the `mvc_storagevolumessummaryvw` view.


Results from these queries populate a map as shown below:



- ▶ **Storage Fabric.** This CIT inherits from the Network Resource CIT and represents a storage fabric. This CIT has no change monitoring enabled.
- ▶ **Storage Processor.** This CIT represents other storage devices such as SCSI controllers, and inherits from the Host Resource CIT. A Storage Processor CIT monitors change on parameters such as state, status, version, WWN, roles, power management, and serial number.
- ▶ Storage Pool.

Storage Pool information is also collected from each Storage Array using the query below.

Results from this query populate a map as shown below:



Host Details

This package retrieves Host details from the `mvc_hostsummaryvw` view. The discovery pattern retrieves detailed information on HBAs from the `mvc_cardsummaryvw` view.


SE maintains information on Operating Systems, CPUs, Memory, IP address, and DNS name on each host. UCMDB uses this information to create host CIs of type UNIX or Windows, and adds CPU and Memory CIs for each host as available.

Since UCMDB uses the IP address of a host as part of its primary key, this discovery pattern attempts to use the IP address from SE for this purpose. If an IP address is not available, the pattern then attempts to resolve the hosts IP address using a DNS name. If neither an IP address nor a DNS name is available, the discovery pattern ignores the host.

Similar to Storage Arrays, a host may have Fiber Channel Ports directly associated with itself or on HBAs on the host. The discovery pattern uses three separate queries to retrieve Fiber Channel Ports for each host. The job retrieves detailed information about Fiber Channel Ports on each host from the `mvc_portsummaryvw` view. Since this view uses a ContainerID as the key, the discovery pattern queries the view by container ID for each host, and each HBA on a host.

Finally, the discovery pattern retrieves detailed information about Logical Volumes on each host from the `mvc_hostvolumesummaryvw` and `mvc_hostcapacityvw` view. The `mvc_hostcapacityvw` view maintains capacity information for each volume over multiple instances in time, and the discovery pattern uses only the latest available information.


Results from these queries populate a map as shown below:



SAN Topology

SAN Topology consists of the Fiber Channel network topology and includes (fiber channel) connections between Fiber Channel Switches, Hosts, and Storage Arrays. SE maintains a list of WWN's each Fiber Channel Port connects to, and this package uses this list of WWN's to establish Fiber Channel Connection links.


Results from these queries populate a map as shown below:



Storage Topology

Storage topology consists of relationships between Logical Volumes on a host and Logical Volumes on a Storage Array. This package uses multiple tables to identify this relationship as shown in the query below. This view is a simple summary of all of the above information.

Results from these queries populate a map as shown below:




Views

This package contains views that display common storage topologies. These are basic views that can be customized to suit the integrated SE applications.

Storage Array Details


This view shows a Storage Array and all its components including Logical Volumes, HBAs, Storage Processors, and Fiber Channel Ports. The view shows each component under its container Storage Array and groups Logical Volumes by CI Type.

Storage Array does not require all components in this view to be functional. Container links stemming from the Storage Array have a cardinality of zero-to-many. The view may show Storage Arrays even when there are no Logical Volumes or Storage Processors.




FC Switch Details

This view shows a Fiber Channel Switch and all connected Fiber Channel Ports.



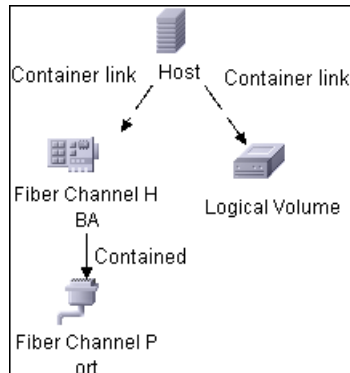
Storage Pool Details

This view shows Storage Pools with associated Storage Arrays and Logical Volumes.




Host Storage Details

This view shows only Hosts that contain a Fiber Channel HBA or a Logical Volume. This keeps the view storage-specific and prevents Hosts discovered by other discovery jobs from being included in the view.




SAN Topology

This view maps physical connections between Storage Arrays, Fiber Channel Switches, and Hosts. The view shows Fiber Channel Ports below their containers, and groups the Fiber Channel Connect relationship CIT to prevent multiple relationships between the same nodes from appearing in the top layer.



Storage Topology

This view maps logical dependencies between Logical Volumes on Hosts and Logical Volumes on Storage Arrays. There is no folding in this view.



Correlation Rules


This package contains basic correlation rules to enable the Impact Analysis and Root Cause Analysis features in UCMDB. These correlation rules are templates for more complex rules that you can define based on business needs.

All correlation rules fully propagate both Change and Operation events. For details on impact analysis, see “Correlation Manager Window” and “Correlation Manager Overview” in *Model Management*.

Note: Correlation events are not propagated to Fiber Channel Ports for performance reasons.


Storage Array Devices to Storage Array

This correlation rule propagates events between Logical Volumes, Storage Processors, Fiber Channel HBAs, and Storage Arrays.




Host Devices to Host

This correlation rule propagates events between Fiber Channel HBAs and Hosts, and Logical Volumes on the Host.




Logical Volume to Logical Volume

This correlation rule propagates events on a Logical Volume contained in a Storage Array to the dependent Logical Volume on the Host.




FC Switch Devices to FC Switch

This correlation rule propagates events from a Fiber Channel Port to and from a Switch. The event is propagated to the associated Storage Fabric as well.



FC Port to FC Port

This rule propagates events on a Fiber Channel Port to another connected Channel Port.



In an example scenario, consider an HBA on a Storage Array going down. The event propagates from the HBA to the Storage Array and the Logical Volumes on the Array because of the Storage Devices to Storage Array rule. The correlation event on the Logical Volume then propagates to other dependent Logical Volumes through the Logical Volume to Logical Volume rule. Hosts using those dependent Logical volumes see the event next because of the Host Devices to Host rule. Depending on business needs, you can define correlation rules to propagate events from these hosts to applications, Business Services, Lines of business, and so on. This enables end-to-end mapping and impact analysis using UCMDB.


Reports

This package contains basic reports that can be customized to suit the integrated SE applications.

In addition to the system reports, Change Monitoring and Asset Data parameters are set on each CI Type in this package, to enable Change and Asset Reports in UCMDB. For details see “Storage Array Configuration” on page 23, “Host Configuration” on page 24, “Storage Array Dependency” on page 24, and “Host Storage Dependency” on page 25.


Storage Array Configuration

This report shows detailed information on Storage Arrays and its sub-components including Fiber Channel Ports, Fiber Channel Arrays, and Storage Processors. The report lists Storage Arrays with sub-components as children of the Array.




Host Configuration

This report shows detailed information on Hosts that contain one or more Fiber Channel HBAs, Fiber Channel Ports, or Logical volumes. The report lists Hosts with sub-components as children of the Host.




Storage Array Dependency

This report is an attempt to map dependencies on a Storage Array. The report also displays information on Switches connected to it.



Host Storage Dependency

This report shows detailed information on storage infrastructure dependencies of a Host. The report lists Hosts and dependent components.



Storage Pool Configuration

This report shows detailed information on Storage Pool configuration.

