



TIDAL ENTERPRISE SCHEDULER RELEASE NOTES V 6.0.2

These release notes highlight new features, enhancements and changes within Tidal Enterprise Scheduler™ (TES) v6.0.2. If you encounter any problems or have any questions, call either 650-475-4600 and select option "3" for support. You will be routed to the Cisco Customer Interaction Network, or CIN:

<http://www.cisco.com/CustAdv/ts/tso/globalops/gsde/CIN.shtml>

Upgrading from 6.0.0, 6.0.1 to 6.0.2

When upgrading from 6.0.0, 6.0.1 to 6.0.2, the data cache of the TES plugin (also called DSP) needs to be deleted and re-created. The Client Manager will automatically re-create the cache if the cache folder is deleted from the file system. The cache folder is named "cache" and is located under the TES plugin's root folder. The TES plugin's root folder is located under the **plugins** folder of the Client Manager's root folder. The name of the TES plugin's root folder depends on the name assigned to that plugin. For example, the cache folder of a plugin named "production" would be located under **{ClientManager Root Folder}/plugins/production/cache**.

New Features

The following new features are included in this release of Tidal Enterprise Scheduler.

New Adapters

JD Edwards

The Tidal Enterprise Scheduler Adapter for JD Edwards facilitates the automation and scheduling of JD Edwards as part of cross-platform, cross-application enterprise process orchestration. Highlights include:

- ◆ Designed using the same user interface approach as other types of jobs, the JD Edwards adapter provides seamless integration of Universal Batch Engine (UBE) management into existing operational processes
- ◆ Integrated security management features of the scheduler can be applied to the implementation of JD Edwards connections and jobs to ensure secure management and use of required security credentials

You can use the JD Edwards adapter to add complex scheduling functions to an existing JD Edwards solution from a centralized console. You can schedule, launch, monitor, and control JD Edwards Universal Batch Engine (UBE) programs, which are used to execute JD Edwards defined processes. The JD Edwards adapter also provides easy to use access for connectivity setup, job definition, and monitoring.

The JD Edwards adapter includes the following features:

- ◆ JD Edwards connection management to monitor system status with a live connection to the Enterprise Server
- ◆ JD Edwards job management to do the following:
- ◆ Schedule and monitor JD Edwards UBEs from a centralized work console with Scheduler
- ◆ Specify parameter overrides for data selection and processing options used for job execution
- ◆ Print, save, and delete report output
- ◆ Associate defined dependencies and events with Scheduler for scheduling control
- ◆ Provide job execution control including terminating, holding, and releasing scheduled UBEs.

JDBC

The Tidal Enterprise Scheduler Adapter for JDBC provides automation of database operations as part of cross platform process organization between Tidal Enterprise Scheduler and the enterprise. The consolidation of scheduling automation is accomplished as follows:

- ◆ Designed using the same user interface approach as other types of jobs, seamlessly integrating JDBC data management into existing operational processes.
- ◆ Integrated security management features of the scheduler can be applied to the implementation of connections using JDBC and associated jobs to ensure secure management and use.

Scheduler's centralized console brings complex scheduling functions to an existing database management system. Users can define, launch, control, and monitor SQL commands, and stored procedures submitted to a database via JDBC on a scheduled basis. The JDBC adapter also provides easy to use access for connectivity setup, job definition, and monitoring.

The JDBC adapter includes the following features:

- ◆ Connection management to monitor system status with a live connection to the database via JDBC.
- ◆ JDBC job and event management includes the following:
 - ❖ Schedule and monitor SQL queries and stored procedures from a centralized work console with Scheduler
 - ❖ Provide dynamic runtime overrides for SQL parameters and values passed to stored procedures
 - ❖ Provide output formatting options to control the results, including table, XML, and CSVs associate defined dependencies and
 - ❖ A events with Scheduler for scheduling control
 - ❖ Provide job execution control including terminating, holding, and releasing scheduled SQL

VMWare

Workload automation extends beyond managing the execution of business application tasks to automating the management of the infrastructure that supports these tasks as well. As both services and infrastructure become virtualized through SOA and virtualization technology, the scheduling solution needs to proactively manage the loose coupling between business process and infrastructure.

Host servers often support dozens to hundreds of virtual machines. This creates a challenge for coordinating the availability of hosts through power management, preserving and restoring machine states through snapshots, and configuring servers for specific tasks.

For example, consider a business process that requires a dedicated machine to produce quarterly tax reports with a specific configuration of CPU, memory and applications, but only for a limited period of time four times a year. If this dedicated machine were virtualized using VMware, it could remain powered off until the schedule demanded that it be ready for this process, and only then powered on. Knowing when to power on or power off a machine and/or when to take a snapshot for checkpoint backups is the job of a scheduler. Therefore, the scheduler needs to automate these tasks along with the more typical application tasks.

The Tidal Enterprise Scheduler VMware Adapter connects to its Virtual Infrastructure to provide a new type of job to manage virtual machines and their hosts. The adapter enables you to perform the following tasks:

- ◆ Power Management – Power on, suspend or resume, shutdown, and power off virtual machines and their hosts
- ◆ Snapshot Management – Capture and restore snapshots on virtual machines
- ◆ Host Management – Power host servers up and down, reboot host servers, and move host servers into and out of maintenance

Release Notes

- ◆ Configuration Management – Configure CPU, memory, disk, and the network on virtual machines, and migrate virtual machines using VMotion

The adapter also supports events that can be used to automate a response to changes in power state and/or host and guest performance conditions.

Changes from 6.0.1 to 6.0.2

- ◆ For brand new installations of 6.0.2, the default security policies have been changed from 6.0.1. The Workgroup VIEW/ADD/EDIT rights have been removed from the User security policy.
- ◆ Client Manager support for SSL connectivity to Active Directory (CSCtn38008)
- ◆ A super user cannot "demote" themselves to become a non-super user (CSCto78772)
- ◆ Only a super user is allowed to promote another user to be super user
- ◆ Non-super is not allowed to modify the security of a super user

Known Issues

CSCtq68888

Problem

VMware Adapter in Enterprise Scheduler 5.3.1 was designed with the assumption that display names of all objects within the VMware infrastructure (folders, virtual machines, hosts, resource pools, etc.) are unique. When connection to a VMware infrastructure (vCenter or VMware ESX server) is first established (at server start up or when a new target is added), the adapter examines the VMware inventory and constructs a cache of various objects indexed by object's display name.

The following figure displays a virtual machine, *HOU-TEST-1*, and a host object, *HOU-ESX-02.tidalsoft.local* identified by their display names.




Figure 1 Example VMware Job Definition Dialog from 5.3.1

The assumption made about display name uniqueness was incorrect. There are no guarantees that the display names of objects within a single vCenter will be unique. Moreover, when multiple datacenters are created within a single vCenter, name collisions are likely. All activities and events exposed by the VMware adapter have one or more parameters that refer to these objects.

For example, Clone VM activity exposes the following configuration properties that refer to objects in the inventory by display name:

- ◆ Source VM
- ◆ Resource Pool to use for the clone
- ◆ Folder to put the clone into
- ◆ Datastore to use for the clone

Since the internal cache stores the objects indexed by display name, the Enterprise Scheduler 5.3.1 VMware Adapter implementation uses this display name to obtain the object on which the task or event need to be performed. If there are duplicates for any objects above, the activity might fail with an unexpected error message (e.g. that the resource pool is not from the correct datacenter) or the activity might succeed placing the final VM into the wrong resource pool or folder.

Solution

To resolve this issue, Enterprise Scheduler 6.0.2 VMware adapter implementation uses this inventory path to find the objects in the internal cache. This eliminates the possibility of using a wrong object for a task or event.

For an object, the inventory path will be a string in the format "a/b/c".

For example, the paths for the two "circled" VMs in following figure are "SJC-BL-ENV/SJC-UCS/ TEO/TEODEV-W2K8-64-02" and "SJC-BL-ENV/SJC-UCS/ BGD01".

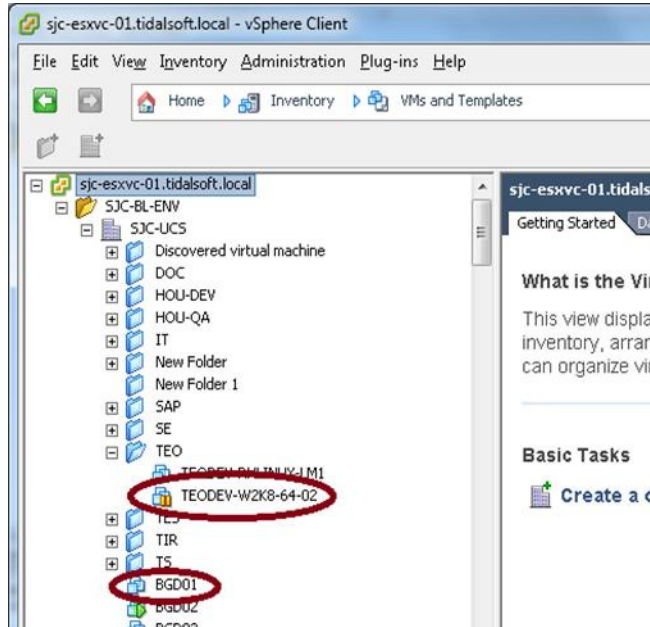


Figure 2 Inventory Path Example

The following graphic displays a virtual machine (ESX-41/Clones/TestVM) and Host (ESX-41/SJC-41esx/172.25.7.46) object identified by their inventory path.

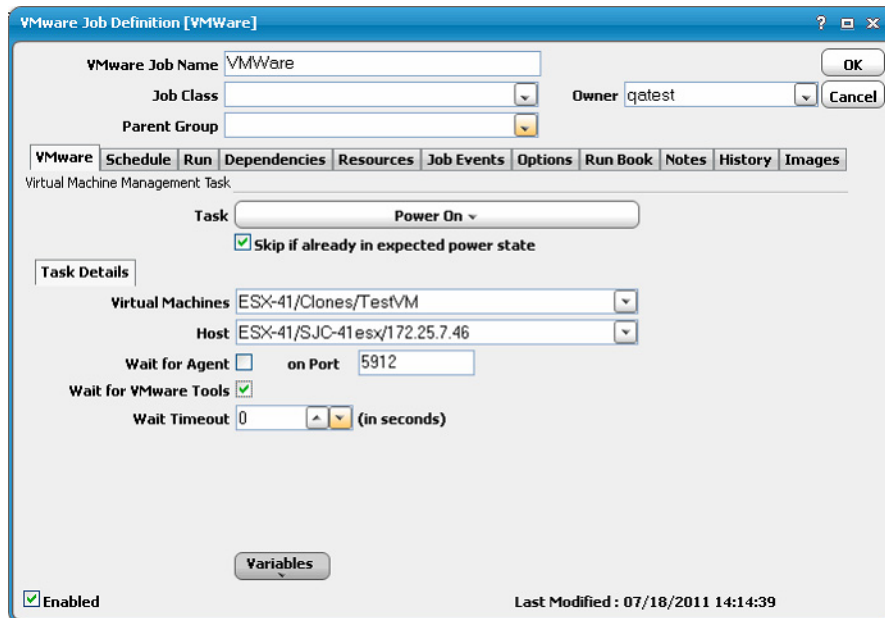


Figure 3 Virtual Machine and Host Identified by Inventory Path

Workaround

Although VMware jobs can be upgraded from 5.3.1 are upgraded to 6.0.2, these jobs still use the display name for inventory objects.

To execute the upgraded jobs successfully, edit the upgraded VMware job and change the display name of the objects to inventory path. You can make this change by selecting the correct inventory path from the list for the corresponding object.

Failing to change the display name to inventory path will result in unsuccessful job execution.

CSctq84729

VMWare tasks complete abnormally because a task in in progress.

Problem

VMware Issue "Operation Failed Since Another Task Is In Progress"

This error message can be seen when VM is actually frozen but it will not reflect the status and display the VM as running.

Any attempts to force a VM shutdown or restart results in the error "Operation Failed Since Another Task Is In Progress".

Same if attempted to Vmotion the machine.

Workaround

To fix this without rebooting the ESX server, kill the VM process via command line.

CSctq83806

Performance metrics inconsistencies for VMware adapter

Problem

Inconsistency in entries for Performance metrics in VMware adapter

Performance metrics for ESX server Shown in Virtual Center Connection:

Release Notes

Connection Definition (Edit Mode) [VMware VC Connection[VMwa...]]

VMware Adapter
 Name VMware VC Connection

General VMware Connection Virtual Machines History Performance Options Outages Description

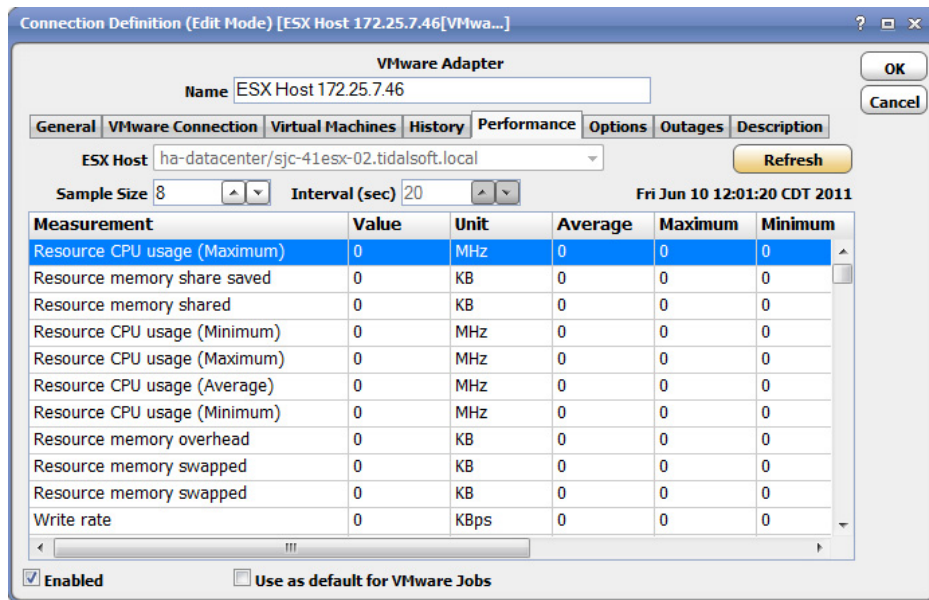
ESX Host [Dropdown] Refresh

Sample Size 10 Interval (sec) 20 Fri Jun 10 19:56:06 CDT 2011

Host	Measurement	Value	Unit	Average	Maximum
ESX-41/SJC-41esx/172.25.7.46	Usage	1293	Percent	1293	1294
	Usage in MHz	28	MHz	68	162
	Highest latency	5	Millisecond	6	6
	Usage	35	Percent	84	202
	Consumed	1075940	KB	1076429	107735
	Balloon	0	KB	0	0
	Swap out rate	0	KBps	0	0
	Swap in rate	0	KBps	0	0
	Usage	110	KBps	179	233
	Overhead	0	KB	0	0
	Usage	0	KBps	0	0

Enabled Use as default for VMware Jobs

Performance metrics for ESX server shown in Esx Server Connection displays more metrics than Virtual center Connection:



Measurement	Value	Unit	Average	Maximum	Minimum
Resource CPU usage (Maximum)	0	MHz	0	0	0
Resource memory share saved	0	KB	0	0	0
Resource memory shared	0	KB	0	0	0
Resource CPU usage (Minimum)	0	MHz	0	0	0
Resource CPU usage (Maximum)	0	MHz	0	0	0
Resource CPU usage (Average)	0	MHz	0	0	0
Resource CPU usage (Minimum)	0	MHz	0	0	0
Resource memory overhead	0	KB	0	0	0
Resource memory swapped	0	KB	0	0	0
Resource memory swapped	0	KB	0	0	0
Write rate	0	KBps	0	0	0

CSctq33569

Under certain circumstances, VMware can "freeze".

Problem

VMware adapter displays "Cannot perform the task because another task is in progress".

Workaround

Following links provide a workaround for this issue:

<http://communities.vmware.com/thread/60671?tstart=0&start=30>

<http://sysadminhell.blogspot.com/2008/03/couple-of-vmware-issues.html>

Fixed Bugs

Bug # CSCto7873

Problem

Resolve issues where a TES workgroup security policy may extend into a users base rights.

Bug # CSCzo31133

Problem

TES variables are not always resolved within the Job Details (Run Info and Notes)

Bug # CSCtn57493

Problem

Business Views may automatically re-arrange contents when adding a new job into the view.

Bug # CSCtq55537

Problem

Informatica Adapter does not support a Enterprise Schedule job with more than 5 parameters.

Bug # CSCtq50289

Problem

Informatica Adapter consuming the system memory and causing the Master to shutting down.

Bug # CSCtr26960

Problem

Enterprise Scheduler 6.0.2 needs support for the \$PREFIX group variable. This group variable was missing in early versions of TES 6.

Bug # CSCtr49633

Problem

In the Job Definition dialog, the Job Owner list is unfiltered. This issue has now been resolved and only lists valid owners that pertain to the rights of the job creator.

