

**IBM Workload Automation
Overview
Version 10.2.7**

Note

Before using this information and the product it supports, read the information in [Notices on page cxlviii](#).

This edition applies to version 10, release 2, modification level 7 of IBM® Workload Scheduler (program number 5698-T09) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

List of Figures.....	vii	User Object Credential Management - Ensuring operational continuity and security.....	33
Chapter 1. Summary of enhancements.....	8	Monitor certificate expiration - Total certificate control.....	34
IBM Workload Scheduler version 10.2.7 enhancements.....	9	Changed features, improvements, and RFEs in version 10.2.5.....	34
Advanced reliability for File Transfer integration.....	11	IBM Workload Scheduler version 10.2.4 enhancements.....	36
Automated database maintenance and optimization suite.....	12	Maximize workflow efficiency with advanced job stream submission and management.....	36
Stay in sync with workspace merge.....	13	Filtering with Orchestration Query Language in the Orchestration Monitor.....	37
Run cycle preview: visualize your schedule.....	14	Operator instructions within the Orchestration Monitor.....	38
Multiple job streams with dependencies submission.....	14	Discover the new and improved integration with password vaults.....	39
Successors & nested job streams submission.....	15	Improved security for password encryption.....	40
Run cycle based job stream submission.....	15	Unlocking the enhanced capabilities of Certman.....	41
Advanced operational oversight in Orchestration Monitor.....	16	Workstation management within the Graphical Designer.....	42
New Send Event action for ServiceNow integrations.....	17	Run cycle groups available as triggers in the Graphical Designer.....	42
Enhanced shadow job integration.....	18	Resources available in the Graphical Designer.....	43
Elevate your automation: the definitive guide to IBM® Workload Scheduler best practices.....	19	Workflow details available in the Orchestration Monitor.....	43
Accelerate job resolution and improve job stream visibility in Orchestration CLI	19	Enhanced job control post-submission.....	44
Changed features, improvements, and RFEs in version 10.2.7.....	20	Performance enhancements.....	45
IBM Workload Scheduler version 10.2.6 enhancements.....	22	Changed features, improvements, and RFEs in version 10.2.4.....	45
Introducing proactive observability: monitor IBM® Workload Scheduler like never before with Instana.....	22	IBM Workload Scheduler version 10.2.3 enhancements.....	46
Multiple workspaces: enhancing design flexibility.....	24	Graphical Designer: a brand-new way to design your orchestration.....	47
Remote-to-Remote File Transfer integration.....	25	Elevate your data management with PostgreSQL database.....	47
Centralized JWT Signing Key Management.....	25	Unleash the power of the cloud using Amazon RDS for MSSQL.....	48
Certman security enhancements: automated validation & OpenSSL integration.....	26	Customize your interfaces with Private Label Settings.....	48
Enhanced OQL filtering: the <code>STARTED</code> alias.....	27	Certman: the new tool to manage certificates.....	48
New in Orchestration Monitor: update agent certificates remotely.....	27	Power your Orchestration CLI model commands with the newly supported items.....	48
Referencing operator instructions in event rules.....	28	Explore the new items now available with the scalable and flexible REST APIs V2.....	49
Job stream view enhancements.....	29	Embedded job definitions for enhanced customization.....	49
Changed features, improvements, and RFEs in version 10.2.6.....	30	Changed features and feature capabilities in version 10.2.3.....	50
IBM Workload Scheduler version 10.2.5 enhancements.....	30	IBM Workload Scheduler version 10.2.2 enhancements.....	51
Advancing security with FIPS 140-3 compliance.....	31	Installation without root privileges extended to all IBM® Workload Scheduler components	51
Streamlined engine reporting configuration in administration panel.....	32		
Navigating OQL Queries in the Orchestration Monitor.....	33		

Orchestration CLI extended to include model commands.....	52	Advanced password encryption.....	67
Easier configuration for Orchestration CLI	52	Installing the agent on IBM i with a user different from QSECOFR.....	68
Automation as code: object definitions in JSON and YAML formats.....	53	Changed features and feature capabilities in version 10.1 Fix Pack 1.....	68
Observability with OpenTelemetry.....	53	Satisfying Requests for Enhancements (RFEs).....	69
CyberArk and Kerberos integrations automatically installed on dynamic agents.....	54	IBM Workload Scheduler version 10.1.0 enhancements.....	70
Serviceability enhancements.....	54	Managed file transfer.....	70
Changed features and feature capabilities in version 10.2.2.....	55	Workload Designer: a new concept to manage your workload.....	71
Satisfying Requests for Enhancements (RFEs).....	56	Detecting anomalies in your workload with AI Data Advisor (AIDA).....	72
IBM Workload Scheduler version 10.2.1 enhancements.....	56	Automatic encryption at rest for key product files.....	72
Enhanced security for default certificates.....	56	Automatic SSL configuration for fault-tolerant agents.....	73
Certificates downloadable to fault-tolerant agents at installation time or later.....	58	Support for Google Cloud SQL for SQL server.....	73
New OCLI plan commands supported.....	58	IBM® Workload Scheduler events now logged for monitoring tools.....	73
Support for OpenSSL 3.0.x libraries from UNIX operating systems.....	58	Changed features and feature capabilities.....	74
New actions available in the Orchestration Monitor.....	59	Satisfying Requests for Enhancements (RFEs).....	77
Customizing the login page.....	59	IBM Z Workload Scheduler V10.2 SPE April 2026.....	78
Changed features and feature capabilities in version 10.2.1.....	59	IBM Z Workload Scheduler V10.2 SPE December 2025.....	80
IBM Workload Scheduler version 10.2.0 enhancements.....	60	IBM Z Workload Scheduler V10.2 SPE July 2025.....	81
Enhanced monitoring.....	60	IBM Z Workload Scheduler V10.2 SPE April 2025.....	82
Self-Service Catalog Approval Flow.....	61	IBM Z Workload Scheduler V10.2.....	83
New enhanced integrations.....	61	IBM Z Workload Scheduler V10.1 SPE December 2023.....	84
Enhanced AI Data Advisor.....	61	IBM Z Workload Scheduler V10.1 SPE July 2023.....	86
Changed features and feature capabilities in version 10.2.....	62	IBM Z Workload Scheduler V10.1 SPE November 2022.....	87
IBM Workload Scheduler version 10.1.0 Fix Pack 4 enhancements.....	62	IBM Z Workload Scheduler V10.1.....	89
IBM Workload Scheduler version 10.1.0 Fix Pack 3 enhancements.....	62	Chapter 2. IBM Workload Automation and ITUP.....	94
Exporting variable tables with multiple lines	63	The ITUP processes.....	94
Use of base folder in Orchestration CLI commands.....	63	Service execution and workload management.....	94
Run cycle preview.....	63	Managing workload with IBM Workload Automation.....	95
IBM Workload Scheduler version 10.1.0 Fix Pack 2 enhancements.....	64	Chapter 3. IBM Workload Scheduler.....	97
IBM Workload Scheduler version 10.1.0 Fix Pack 1 enhancements.....	64	Overview.....	97
Self-Service Catalog enhancements.....	64	What is IBM Workload Scheduler.....	97
Orchestration CLI.....	65	The IBM Workload Scheduler network.....	97
Introducing REST API V2	65	Manager and agent types.....	98
Orchestration Query Language (OQL).....	65	Topology.....	101
Enhancing authentication using API Keys.....	66	Networking.....	101
Enhancing agent authentication using JSON Web Tokens.....	67	IBM Workload Scheduler components.....	103
		IBM Workload Scheduler scheduling objects.....	104
		The production process.....	107
		Scheduling.....	108

Creating job streams.....	108
Setting job recovery.....	109
Defining and managing mission-critical jobs.....	109
Scheduling workload dynamically.....	110
Running production.....	111
Running the plan.....	111
Running job streams.....	112
Monitoring.....	113
Reporting.....	114
Auditing.....	115
Using event-driven workload automation.....	115
Options and security.....	116
Setting the IBM Workload Scheduler options.....	116
Setting security.....	117
Using time zones.....	118
Using the access methods of IBM Workload Scheduler.....	118
Supported agents.....	119
Chapter 4. IBM Z Workload Scheduler.....	121
How your production workload is managed.....	121
Structure.....	121
Concepts.....	121
Using Plans in IBM Z Workload Scheduler.....	126
Automatically controlling the production workload.....	127
Integration with IBM Tivoli Monitoring.....	132
Interfacing with other programs.....	132
Manual control and intervention.....	133
Management of critical jobs.....	134
Security.....	135
Configurations of IBM Z Workload Scheduler.....	137
The controlling system.....	137
Controlled z/OS® systems.....	138
Remote panels and program interface applications.....	139
Scheduling jobs that are in IBM Workload Scheduler.....	139
Chapter 5. Dynamic Workload Console.....	141
Chapter 6. End-to-end scheduling.....	143
End-to-end scheduling with fault tolerance capabilities.....	143
End-to-end scheduling with z-centric capabilities.....	145
Distributed agents.....	145
Chapter 7. Hybrid scheduling environment.....	147
Notices.....	cxlviii
Index.....	152

List of Figures

Figure 1: IBM Workload Scheduler network made up by two domains.....	98
Figure 2: Supported agents configuration.....	119
Figure 3: Automatic recovery and restart.....	129
Figure 4: Production workload restart and hot standby....	130
Figure 5: Security.....	136
Figure 6: IBM Z Workload Scheduler configurations.....	138
Figure 7: End-to-end with fault tolerance capabilities configuration.....	144
Figure 8: End-to-end with z-centric capabilities configuration.....	145

Chapter 1. Summary of enhancements

IBM Workload Automation provides the following enhancements:

IBM Workload Scheduler

Version 10.2.7

[IBM Workload Scheduler version 10.2.7 enhancements on page 9](#)

Version 10.2.6

[IBM Workload Scheduler version 10.2.6 enhancements on page 22](#)

Version 10.2.5

[IBM Workload Scheduler version 10.2.5 enhancements on page 30](#)

Version 10.2.4

[IBM Workload Scheduler version 10.2.4 enhancements on page 36](#)

Version 10.2.3

[IBM Workload Scheduler version 10.2.3 enhancements on page 46](#)

Version 10.2.2

[IBM Workload Scheduler version 10.2.2 enhancements on page 51](#)

Version 10.2.1

[IBM Workload Scheduler version 10.2.1 enhancements on page 56](#)

Version 10.2

[IBM Workload Scheduler version 10.2.0 enhancements on page 60](#)

Version 10.1 Fix Pack 4

[IBM Workload Scheduler version 10.1.0 Fix Pack 4 enhancements on page 62](#)

Version 10.1 Fix Pack 3

[IBM Workload Scheduler version 10.1.0 Fix Pack 3 enhancements on page 62](#)

Version 10.1 Fix Pack 2

[IBM Workload Scheduler version 10.1.0 Fix Pack 2 enhancements on page 64](#)

Version 10.1 Fix Pack 1

[IBM Workload Scheduler version 10.1.0 Fix Pack 1 enhancements on page 64](#)

Version 10.1.0

[IBM Workload Scheduler version 10.1.0 enhancements on page 70](#)

IBM Z Workload Scheduler

Version 10.2


- [IBM Z Workload Scheduler Small Programming Enhancements released in April 2026 on page 78](#)
- [IBM Z Workload Scheduler Small Programming Enhancements released in December 2025 on page 80](#)
- [IBM Z Workload Scheduler Small Programming Enhancements released in July 2025 on page 81](#)
- [IBM Z Workload Scheduler Small Programming Enhancements released in April 2025 on page 82](#)
- [IBM Z Workload Scheduler version 10.2 enhancements on page 83](#)

Version 10.1

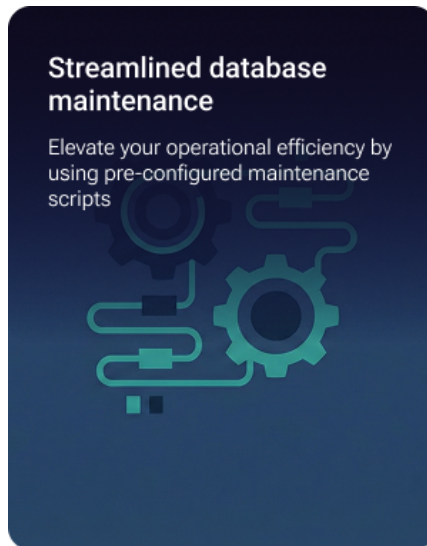
- [IBM Z Workload Scheduler Small Programming Enhancements released in December 2023 on page 84](#)
- [IBM Z Workload Scheduler Small Programming Enhancements released in July 2023 on page 86](#)
- [IBM Z Workload Scheduler Small Programming Enhancements released in November 2022 on page 87](#)
- [IBM Z Workload Scheduler version 10.1 enhancements on page 89](#)

IBM Workload Scheduler version 10.2.7 enhancements


Learn about the IBM Workload Scheduler version 10.2.7 features.



on page 11



on page 12



on page 13

Multiple job streams with dependencies submission

You can now submit multiple job streams simultaneously



on page 14

Successors & nested job streams submission

You can now include successors and nested job streams when submitting workloads.



on page 15

Run cycle based job stream submission

You can now apply existing scheduling logic to job stream submissions for better precision



on page 15

Send event now available for ServiceNow

You can now configure HCL Workload Automation to send events directly to a ServiceNow instance when a specific rule is triggered



on page 17

Run cycle preview: visualize your schedule

The Run cycle preview provides an immediate graphical representation of your scheduling rules within the Graphical Designer



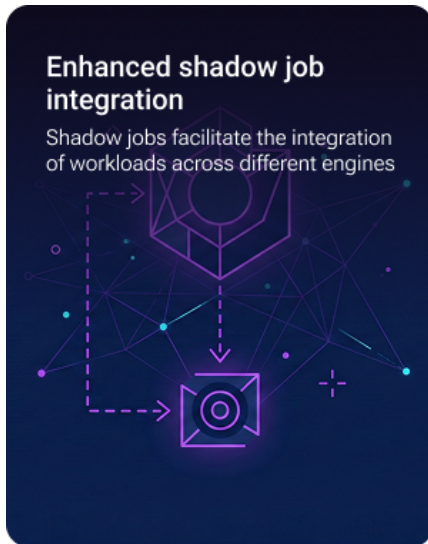
on page 14

Advanced monitoring and dependency control

Unify end-to-end oversight with centralized manual checkpoints, file-based prerequisites, and enriched process context




on page 16



Enhanced shadow job integration

Shadow jobs facilitate the integration of workloads across different engines

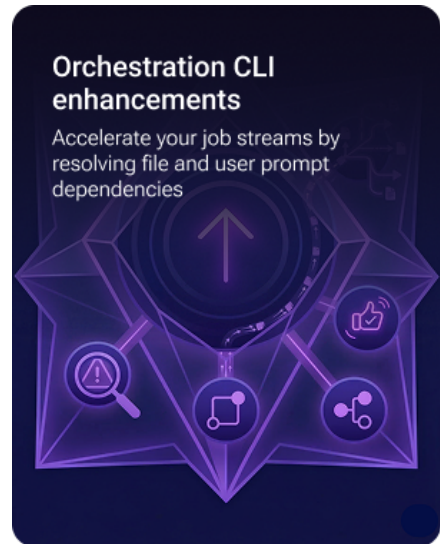
[on page 18](#)



Best practices for an efficient automation

A comprehensive blueprint for operational excellence in the installation, configuration, and optimization

[on page 19](#)



Orchestration CLI enhancements

Accelerate your job streams by resolving file and user prompt dependencies

[on page 19](#)

Learn about the IBM Workload Scheduler version 10.2.7 enhancements.

[Changed features, improvements, and RFEs on page 20](#)

Advanced reliability for File Transfer integration

The latest enhancements to the File Transfer integration ensure every transfer is complete and accurate.

In complex enterprise environments where data integrity is paramount, these new capabilities provide the oversight needed to manage files without the risk of transfer failures.

By implementing intelligent file-status checks, the integration now protects your workflows from the common pitfalls of automated data movement. You can ensure that your downstream processes never encounter corrupted or incomplete data, significantly reducing the need for manual troubleshooting or a request for a new transfer.

Enhanced transfer reliability and data Integrity

The new settings provide granular control over the transfer lifecycle, ensuring that only files that respect specific conditions move through the pipeline:

Dynamic write-completion monitoring:

By monitoring for changes over a specified timeout period, the integration ensures a transfer only begins when the file is stabilized and complete. This effectively prevents the transfer of files that are still being edited.

Minimum file size validation

To further protect against data corruption, you can now define a minimum size threshold for transfers. This prevents the download of truncated or empty files, ensuring that only complete datasets are processed by your business-critical applications.

Automated error prevention

By combining size validation with write-monitoring, the integration significantly reduces the risk of transfer failures. This provides a more resilient architecture that supports high-volume data flows without the risk of processing partial information.

Post-transfer checksum verification

By using the SHA-256 algorithm to validate the file after the transfer is complete, the integration ensures that data has not been altered or corrupted during transit. This provides an additional layer of certainty that the destination file is an exact match of the source.

Transfer smarter and with total confidence using the enhanced capabilities of the File Transfer integration.

For further information about the integration, browse to [Automation Hub](#).

Automated database maintenance and optimization suite

IBM® Workload Scheduler includes a robust Database Maintenance Suite designed to automate the lifecycle management of data storage and query performance. By implementing systematic **Table Reorganization** and **Statistical Sampling**, the system ensures that the underlying RDBMS operates at peak efficiency.

The suite targets all critical application schemas, including Model (MDL), Event (EVT), Planning (PLN), and Dynamic Workload Broker (DWB), ensuring comprehensive coverage across all functional modules.

Elevate your operational efficiency by using pre-configured maintenance scripts. These tools provide a robust solution for managing the health of your data layer, ensuring that your environment remains optimized for high-demand workloads.

By automating routine maintenance, you can reduce manual intervention and focus on core business logic.

Key technical features

Deep fragmentation clean-up

Utilizes high-level maintenance commands to resolve table bloat and reclaim unused disk space, significantly reducing the storage footprint.

Heuristic statistics update

Refreshes internal data distribution metadata, providing the Query Optimizer with the necessary insights to select the most efficient access paths for complex joins and filters.

Schema-wide coverage

Systematically processes over 100 core tables to prevent performance bottlenecks in high-concurrency environments.

Strategic benefits

Consistent peak performance

Removes the natural performance drift over time, ensuring that end-user interactions and batch processing remain fast and responsive.

Storage efficiency

Prevents unnecessary infrastructure costs by maintaining a compact database size and reclaiming storage from deleted or legacy records.

Reduced operational risk

Proactively mitigates the risk of sudden query degradation or storage-related outages, ensuring higher availability for enterprise workloads.

Scalability ready

Prepares the environment for increased data volumes, making the application more resilient as your business grows.

To learn more and understand in detail how to use the functionality, as well as the steps required to properly maintain your database and increase performance, see the administrative tasks for each database in *Administration Guide*.

Stay in sync with workspace merge

The new version of IBM® Workload Scheduler introduces the possibility to merge workspaces.

Maintaining synchronization between your local item edits and the database definitions is now effortless. The new workspace merge feature eliminates the friction of collaborative editing, ensuring that your work is always built on the most current version of a job stream. By bridging the gap between your local environment and the database, workspace merging ensures that collaboration no longer comes at the cost of speed or data safety.

Stay in control of your updates

You no longer have to worry about accidentally overwriting the work of someone else. If a discrepancy is detected during deployment, the **Merge Conflicts Found** panel guides you through the process of aligning your work with the latest version. By merging, you ensure your workspace items are perfectly aligned with the database, even if colleagues have made updates in the meantime.

Proactive conflict prevention

The merge feature is not just for resolving conflicts; it is a proactive tool. Integrated directly into the **Deploy** menu, it encourages a "check-first" workflow that protects the integrity of your automation logic.

Total creative control

You decide how to integrate changes. Whether you want to overwrite your local work with the latest database updates or prioritize your own edits, the tool puts the final decision in your hands.

Better navigation

- **Focus on what matters:** Use the **Show conflicts only** toggle to hide items that do not require your attention.
- **Step-by-step review:** Use the **Previous** and **Next** buttons to move through each discrepancy systematically.

To discover more about the workspace merge, see *Managing your workspaces*.

Run cycle preview: visualize your schedule

The Run cycle preview provides an immediate graphical representation of your scheduling rules within the **Graphical Designer**.

This feature generates a calendar view showing exactly when a job stream or run cycle group will run based on your current configuration.

Immediate visual feedback

Avoid guesswork when configuring complex scheduling rules.

Calendar visualization

Instantly see which days are included or excluded based on your run cycles and calendars.

Interactive updates

The preview automatically refreshes whenever you modify properties such as validity dates, time zones, or non-working day rules.

Detailed insights

Hover over specific dates in the calendar to view tooltips detailing why a specific run instance was included or excluded, including the specific run cycle rule responsible.

Find out more about the run cycle preview in [Run cycle preview](#)

Multiple job streams with dependencies submission

You can now submit multiple job streams simultaneously, significantly reducing the time spent on manual sequencing and dependency management.

This enhancement ensures that complex workstreams are added to the plan accurately, respecting their original definitions and business logic. This feature introduces a more intelligent way to submit workloads into the current plan, moving away from repetitive manual tasks toward a system-guided approach.

Core benefits of the multiple submission enhancement

- **Automated dependency handling:** Eliminate the risk of manual errors. When submitting multiple job streams, the system manages dependencies automatically, ensuring they enter the plan in the correct order.
- **Draft submission:** Gain the flexibility to run `Draft` job streams without changing their permanent status. This is ideal for ad-hoc troubleshooting or ticket resolution where a pre-defined, non-active job stream must be processed immediately.
- **Efficiency at a glance** The new interface offers a guided workflow to match the urgency of your task. The step-by-step wizard enables you to modify properties (such as **Priority**), refine run cycle intervals, and use the instance accordion to review, expand, or remove specific instances before they are finalized.

Explore the benefits of the new wizard for submitting nested job streams and run-cycle-based job streams:

- [Successors & nested job streams submission on page 15](#)
- [Run cycle based job stream submission on page 15](#)

For more information, see [Submitting multiple job streams](#).

Successors & nested job streams submission

When you select a job stream to add to the plan, you can now include all its successors automatically.

This ensures that the entire downstream workflow is accounted for without searching for individual dependencies. This feature supports nested job streams up to 25 levels deep or 250 job streams and their successors, maintaining the integrity of the whole workflow.

Workflow integrity

Ensure the entire downstream workflow is accounted for in the plan by automatically including all successors.

Deep dependency support

Manage complex environments easily with support for up to 25 levels of nested job stream dependencies or 250 job streams and their successors.

Operational continuity

Prevent gaps in the plan by identifying and resolving logically linked workflows during the submission phase.

Explore the benefits of the new wizard for submitting multiple job streams and run-cycle-based job streams:

- [Multiple job streams with dependencies submission on page 14](#)
- [Run cycle based job stream submission on page 15](#)

For more information, see [Submitting multiple job streams](#).

Run cycle based job stream submission

You can now use pre-defined run cycles to determine exactly which job stream instances are added to the plan.

This feature allows you to bridge the gap between manual requests and automated scheduling logic. For more control, you can set a specific time interval to filter which instances are generated.

Scheduling precision

Apply exact run cycle definitions to your submission for more accurate workload generation.

Granular control

Use specific time intervals (for example, from 10:00 to 15:00) to filter and submit only the instances you need.

Efficiency at a glance

Use a guided wizard and instance accordion to review, expand, or remove specific instances before they are finalized in the plan.

Explore the benefits of the new wizard for submitting multiple job streams and nested job streams:

- [Multiple job streams with dependencies submission on page 14](#)
- [Successors & nested job streams submission on page 15](#)

For more information, see [Submitting multiple job streams](#).

Advanced operational oversight in Orchestration Monitor

The **Orchestration Monitor** has been extended to provide deeper visibility and centralized control over complex workloads.

These enhancements integrate new monitoring possibilities for prompt and file items, broader visibility for folder search and optimized job stream identification with description-based filtering. By consolidating these elements, operators can manage diverse dependencies and navigate folder hierarchies more efficiently within a single view.

Optimized job stream identification with description-based filtering

The **Description** column is now available in the Orchestration Monitor, providing a rich metadata layer for monitoring and filtering.

- **Extended metadata:** Store alternative names, support contacts, or recovery steps, overcoming the 16-byte naming limit.
- **OQL filtering:** Use Orchestration Query Language Language to instantly isolate job streams based on description keywords.
- **Formatted insights:** View a dedicated panel which supports Markdown formatting such as bold text, bulleted lists, and hyperlinks for advanced troubleshooting instructions.

Enhanced monitoring possibilities for prompt and file

The Orchestration Monitor now includes dedicated views for **Prompt** and **File** items, centralizing the management of manual and data-driven dependencies. By selecting these item types from the monitoring drop-down, operators can access specialized interfaces tailored to these prerequisites.

- **Manage manual checkpoints:** Use the **Prompt** view to monitor and respond to "Yes/No" confirmation requests, ensuring human oversight for sensitive tasks before they proceed.
- **Monitor file dependencies:** Use the **File** view to track real-time data availability (such as **Exists** or **Unavailable**) across the network to prevent schedule delays.
- **Actionable control:** Access object-specific commands, such as **Reply Prompt**, and apply granular OQL filters to isolate critical dependencies instantly.

Extended visibility for folder search

The **Include Subfolders** toggle enables broader visibility into directory hierarchies within the **Orchestration Monitor**.

- **Recursive retrieval:** View items from a parent folder and all nested subdirectories simultaneously.
- **Faster navigation:** Eliminate the need to browse individual directories, facilitating immediate access to deeply nested data.

For more information, see [Orchestration Monitor overview](#).

New Send Event action for ServiceNow integrations

In addition to opening incidents, you can now configure IBM® Workload Scheduler to send events directly to a ServiceNow instance when a specific rule is triggered.

This update enables you to connect IBM® Workload Scheduler events directly to ServiceNow triggers.

Key capabilities

Action variety

Choose to either open an incident or trigger a custom event, or even both, depending on your business needs.

Attribute overrides

When defining actions in the **Workload Designer** or `composer`, you can specify parameters that override the default values set in `optman`.

Automatic entry creation

Triggered events appear as new entries in ServiceNow, containing all parameters defined in the IBM® Workload Scheduler configuration.

Benefits

Faster incident response

By sending an event instead of just opening a generic incident ticket, you can trigger specific workstreams in ServiceNow. ServiceNow's automated event management can therefore map and route alerts to the correct team instantly. By removing the need for manual ticket sorting, you speed up the resolution process and ensure that the data needed to fix the issue is available immediately.

Total control from a single interface

You can define parameters directly within the **Workload Designer** or the `composer` command-line interface. These rule-specific definitions take precedence over the global settings maintained in `optman`. This local override capability ensures that specific event rules can be tailored for unique ServiceNow instances without requiring a global configuration change.

For further information about the integration, browse to [Automation Hub](#). For information about other available actions instead, see [ServiceNow actions](#).

Enhanced shadow job integration

Shadow jobs facilitate the integration and synchronization of workloads across different engines.

A shadow job is defined on a local workstation to map a job running on different engines which can be distributed or z/OS engines. This mapping provides several benefits for managing hybrid workload automation environments:

Workload integration

Connects distributed and z/OS engines, ensuring they function as a single, coordinated system.

Automated dependency management

Ensures local jobs do not start until the corresponding remote job completes successfully, reducing the risk of data inconsistency.

Enhanced monitoring

Provides real-time status updates of remote jobs within the local plan, eliminating the need to switch between different console connections.

Customizable matching

Enables the selection of specific remote job instances based on scheduling dates or time intervals, ensuring precise dependency tracking.

In addition to these advantages, folder support is now available. You can use folders to organize objects by department, line of business, or environment (for example, `/PROD/Finance/Payroll`). Shadow jobs now support mapping to remote jobs stored within these folder structures. Folder management ensures a number of advantages:

Cross-engine synchronization

Enables local job streams to wait for the successful completion of remote processes before starting.

Folder-aware mapping

Supports remote jobs stored in folders. This enables the use of full folder paths to identify remote jobs, reducing naming conflicts in complex environments.

Unified visibility

Updates remote job statuses within the local plan, providing a single monitoring point for hybrid workflows.

Security alignment

Respects folder-based security models by enabling shadow jobs to target specific departmental workloads.

Prevention of naming conflicts

You can have multiple remote jobs with the same name (for example, `CLEANUP`) in different folders. Shadow jobs enable you to target the specific instance you need by providing the full folder path.

Granular access control

Because folders are often tied to security roles, supporting folders in shadow jobs ensures that cross-engine dependencies respect your organizational security boundaries.

For a full description of the key features and major use cases, see the topic about shadow jobs in *Dynamic Workload Console User's Guide*.

Elevate your automation: the definitive guide to IBM® Workload Scheduler best practices

A comprehensive blueprint for operational excellence in the installation, configuration, and optimization of IBM® Workload Scheduler.

In today's complex computing environments, precision is not just a goal—it is a requirement. You can find all the information to give you a head start in the comprehensive, newly minted manual dedicated to the installation, configuration, and optimization of IBM® Workload Scheduler.

This is not just a technical document; it is a blueprint for operational excellence. Whether you are streamlining a fresh deployment or refining a legacy environment, this guide provides the best practice framework needed to ensure your automation remains resilient, scalable, and secure.

Contents of the new manual

Seamless installation strategies

Move beyond the basics with proven methodologies for setting up your environment correctly the first time.

Advanced configuration tuning

Expert advice on tailoring IBM® Workload Scheduler to meet the unique demands of your enterprise infrastructure.

Security and compliance

Integrated frameworks to ensure your scheduling environment adheres to the highest industry standards.

Why This Guide?

This manual represents the intersection of structured documentation and high-performance computing. By moving away from antiquated methods and embracing these modern standards, your team can reduce manual intervention and focus on innovation.

The structured motifs in our framework symbolize the protection and stability that these best practices bring to your data center.

You can find the new manual in [Best practices](#).

Accelerate job resolution and improve job stream visibility in Orchestration CLI

Accelerate your job streams and eliminate restrictions by using the new Orchestration CLI commands to resolve file and user prompt dependencies.

Managing complex job streams requires the ability to quickly identify and resolve dependencies. While users of the conman command-line tool have long enjoyed the ability to easily manage job stream dependencies and user prompts, the current release brings these same powerful capabilities directly to Orchestration CLI. These command enhancements are designed

to give you immediate visibility into your scheduling environment and provide the controls you need to keep your jobs and job streams running smoothly.

The newly added Orchestration CLI commands provide the following key advantages:

reply

Keeps your job streams moving without delay. It confirms a specific prompt or multiple prompts in one step to instantly resume your associated jobs and job streams. It also bypasses manual, instance-by-instance permissions by confirming multiple prompt instances at once. For more information, see [reply](#).

showfile

Improves the visibility into file dependencies to prevent stalled operations. The command lists all associated text files and filters your results by workstation or file status to pinpoint exactly which job streams are waiting for file confirmations. For more information, see [showfile](#).

showprompt

Instantly identifies where user interaction is required. The command lists all associated prompts and filters them based on their status to see exactly which job streams are paused and awaiting a reply. For more information, see [showprompt](#).

nop/unnop

Temporarily disable jobs without deleting them or changing their dependencies. If a job faces a service outage or other disruptions, that job can be paused to facilitate the running of all other associated jobs without any delay. For more information, see [nop/unnop](#).

Changed features, improvements, and RFEs in version 10.2.7

A list of features changed and improved since the previous release and major improvements.

Changed and improved features

The following features have changed or have been introduced or removed since the previous release:

Support for Oracle AI Database 26ai

Oracle AI Database 26ai embeds artificial intelligence directly into the core data and development stack, enabling enterprises to run AI-powered workloads across operational systems and analytic data systems without complex integrations. You can now use Oracle AI Database 26ai with all server components and the Dynamic Workload Console. To improve operational performance, see the settings explained in the topic about optimizing performance for Oracle AI Database 26ai in *IBM Workload Scheduler: Planning and Installation*.

Automatic creation of PostgreSQL database at installation time

Manual creation of the PostgreSQL database before a fresh installation is no longer required. The database is now automatically created at installation time, without intervention on your part. For more information, see the topics about creating the PostgreSQL database tables for the master domain manager and the Dynamic Workload Console in *IBM® Workload Scheduler: Planning and Installation*. If you are using a Dynamic Workload Console in a z/OS environment, the database is also automatically created at installation time, as described in

the topic about creating and populating the database for PostgreSQL for the Dynamic Workload Console in *IBM Z Workload Scheduler: Planning and Installation*.

Custom certificate aliases for WebSphere Application Server Liberty

You can now optionally specify the alias of the certificate that WebSphere Application Server Liberty uses to connect to IBM® Workload Scheduler by setting the **serverKeyAlias** property in the `ssl_config.xml` file. The default value is `server`. If you specify a different alias, WebSphere Application Server Liberty uses only the specified alias. For more information about WebSphere Application Server Liberty configuration files, see the topic about configuring IBM® Workload Scheduler using templates in *Administration Guide*.

New OCLI command to set/remove nop flags from jobs in plan

The `nop` and `unnop` commands temporarily disable and re-enable jobs without deleting them or changing their dependencies. If a job faces a service outage or other disruptions, it can be paused to let all other associated jobs run without delay.

New Documentation Portal for IBM® Workload Scheduler and IBM® Z Workload Scheduler

Effective April 24, 2026, the official documentation for IBM® Workload Scheduler and IBM® Z Workload Scheduler has moved to a redesigned, dedicated portal. The new portal introduces several major improvements, including:

- **AI-powered chat** to search, navigate, and process information across the documentation set. Available starting from April, 30, 2026.
- **A new documentation template** designed to improve clarity, readability, and overall usability.
- **An updated look and feel**, offering a more modern, consistent, and intuitive user experience.

Please update your bookmarks for all releases: [IBM Workload Automation](#).

Requests for enhancements (RFEs)

Requests for Enhancements (RFEs) give customers the opportunity to collaborate directly with the product development team and other users. The team prioritizes and develops new product features based on proposals made by customers.

IBM® Workload Scheduler V 10.2.7 delivers the following RFEs:

WAO-624 - Maintenance scripts are now available for all supported databases

You can now use the dedicated scripts to maintain your database and improve performance. For more information, see [Automated database maintenance and optimization suite on page 12](#).

WAD-I-114, WAO-I-655 - Enhanced password security

IBM® Workload Scheduler now masks passwords encrypted with the AES algorithm in job logs and job properties to guarantee higher security standards.

WAD-I-170 Enhanced Windows file monitoring

The filemonitor utility on Windows has been enhanced to ensure higher reliability during file-transfer operations. This improvement addresses scenarios where files are copied manually or by using third-party tools that do not follow the standard temporary-and-rename pattern.

WAO-I-656

In addition to creating incidents, you can now configure IBM® Workload Scheduler to send events directly to a ServiceNow instance when a specific rule is triggered. For more information, see [New Send Event action for ServiceNow integrations on page 17](#).

WAD-I-149 - Add OCLI Command to Set/Remove NOP Flag for Jobs in Current Plan

Users need the ability to mark jobs as **NOP (No Operation)** or remove the NOP flag via OCLI in the current plan (tentative or planned) to control execution without modifying the plan permanently.

WAD-I-152 - Shadow job integration supporting folder structure

The Shadow job integration now supports remote jobs stored in folders. This enables the use of full folder paths to identify remote jobs, reducing naming conflicts in complex environments. For more information, see the topic about shadow jobs in *Dynamic Workload Console User's Guide*.

WAD-I-176 - Enhanced job stream visibility: real-time job counts

The **Graphical Designer** now displays the total job count directly within the job stream box in the workspace, providing instant visibility into the workload of each stream.

To view a complete list of new, planned and delivered RFEs, see: [RFE online community](#).

IBM Workload Scheduler version 10.2.6 enhancements

Learn about the IBM Workload Scheduler version 10.2.6 features.

IBM Workload Scheduler version 10.2.6 includes the following enhancements:

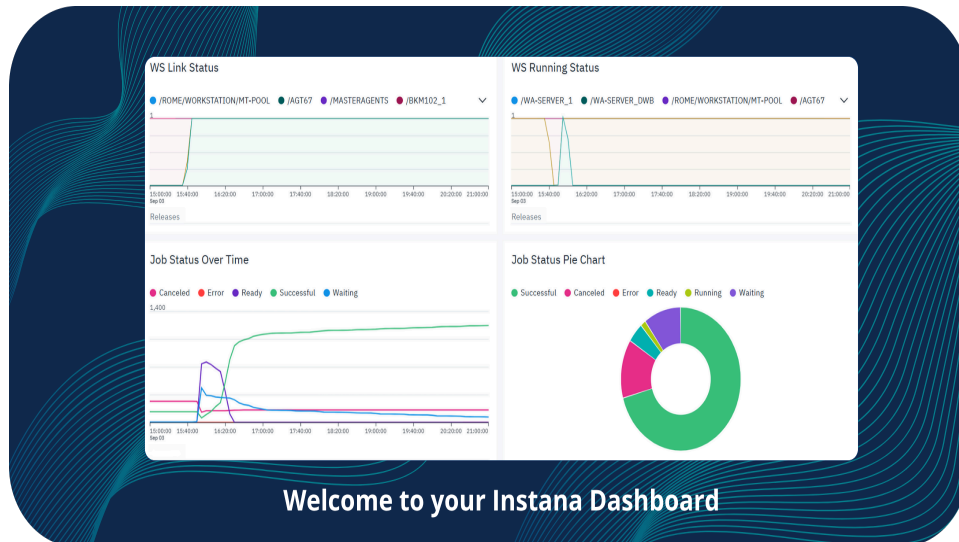
- [Introducing proactive observability: monitor IBM Workload Scheduler like never before with Instana on page 22](#)
- [Multiple workspaces: enhancing design flexibility on page 24](#)
- [Remote-to-remote File Transfer integration on page 25](#)
- [Certman security enhancements: automated validation & OpenSSL integration on page 26](#)
- [Centralized JWT Signing Key Management on page 25](#)[Enhanced OQL filtering: the STARTED alias on page 27](#)
- [New in Orchestration Monitor: update agent certificates remotely on page 27](#)
- [Referencing operator instructions in event rules on page 28](#)
- [Changed features, improvements, and RFEs in version 10.2.6 on page 30](#)

Introducing proactive observability: monitor IBM® Workload Scheduler like never before with Instana

Go beyond traditional monitoring and gain a complete, 360-degree view of your automation landscape.

You can now unlock unparalleled observability for your business processes by integrating IBM® Workload Scheduler with Instana, an observability platform designed for modern, containerized applications. This powerful combination provides the detailed analytics and historical data needed to identify patterns, rapidly troubleshoot issues, and continuously optimize performance.

You can finally correlate your workload jobs performance with the real-time health of your underlying infrastructure, turning complex data into clear, actionable insights. Now, you can easily identify patterns, troubleshoot issues with unprecedented speed, and optimize the performance of your most critical business processes.



This integration empowers you to:

Achieve end-to-end visibility

Track job processing across all systems and applications from a single, unified view, without even opening the Dynamic Workload Console. The dashboard provides the full context you need to understand dependencies, assess the true impact of issues, and make smarter, data-driven decisions.

Solve problems faster

By automatically correlating workload data with infrastructure traces, metrics, and logs, the dashboard can pinpoint the root cause of failures in minutes. This empowers your teams with all the necessary contextual data, ending time-consuming investigations.

Stay ahead of issues

Get real-time alerts for job delays, failures, and performance degradations before they impact your business objectives. Ensure service reliability with proactive notifications.

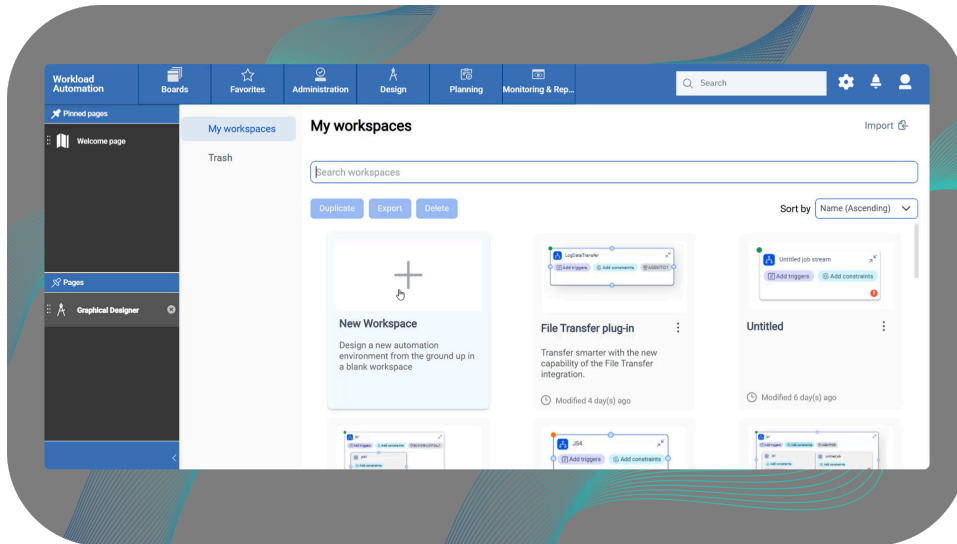
Visualize success instantly

Get immediate time-to-value with a purpose-built, pre-configured dashboard available directly from Automation Hub. The intuitive widgets are designed to distill complex information, like critical job risks, job status trends, and system health. This clarity can help any team member to gain valuable insights without specific Instana training.

Ready to turn complex metrics into clear, actionable insights? Dive into our documentation to learn how to configure this powerful new integration today. See [Leveraging Instana](#) for advanced observability.

Multiple workspaces: enhancing design flexibility

The multiple workspaces feature provides a flexible, structured approach to handle different projects and work streams. To support this, the new **My workspaces** page acts as a central interface to manage all your workspaces from a single location.



Multiple workspaces improve project organization by separating work streams. Each workspace functions as a self-contained environment, which helps you:

- **Leverage design flexibility:** Move freely from a design to another and keep up with fast-paced business needs.
- **Improve clarity:** By keeping projects separate, you can start a new design session or continue a specific process without clearing the canvas before starting a new workspace.
- **Organize by context:** Create distinct workspaces for different projects or scenarios.

Centralized oversight with the My Workspaces page

The new **My Workspaces** page is the central hub to easily manage your workspaces:

- **Centralized control:** The **My Workspaces** page displays a list of all existing workspaces, enabling you to find and manage workspaces using search and sort capabilities.
- **Streamlined operations:** Perform actions on individual workspaces or in bulk.
- **Safety Net:** A dedicated **Trash** page enables you to recover accidentally deleted workspaces, preventing permanent data loss.


Portable workspace definitions

You can export workspace definitions and import previously exported workspaces to continue working on them. The workspace definitions exported from the Dynamic Workload Console (JSON, YAML) are fully compatible with Orchestration CLI and the REST API. This enables seamless interoperability between visual design and programmatic management.

Find out more about multiple workspaces in [Managing your workspaces](#).

Remote-to-remote File Transfer integration

The new remote-to-remote capability within the File Transfer integration drastically simplifies how your enterprise moves critical data.



This new capability enables the secure, high-efficiency transfer of one or more files directly between two remote server endpoints. In the past, moving files between two external systems (Peer A and Peer B) often required a dynamic agent to be installed on at least one of those systems, or involved complex multi-step intermediary transfers. Now, a single job running on any designated agent can orchestrate the transfer between Peer A and Peer B, dramatically reducing complexity and removing the need for local agents on the source or destination machines.

You can now securely manage external data flow with maximum resource efficiency and significant architectural simplification. All transfer actions are run securely using the robust protocols already available and supported by the File Transfer integration: FTP, FTPS, FTPES, SFTP (previously referred to as SSH), and Windows protocols.

Advanced Features and Usability

The remote-to-remote capability fully supports the advanced functionalities already present in the File Transfer integration, ensuring complete control and reliability:

- **Unified Security:** All file transfers leverage the security and encryption capabilities of the underlying File Transfer protocols.
- **Backward Compatibility:** All existing File Transfer functionalities and job definitions are fully maintained and supported, guaranteeing a smooth transition.
- **Intelligent Workflow Details and Restart:** Retains the powerful ability to view the status of every single transfer and restart the transfer of individual files or groups of files without needing to restart the entire job stream from the beginning, ensuring maximum recovery and integrity.
- **Simplified File Selection:** Despite the systems being remote to the HWA network, the integration provides enhanced usability: you can easily browse and visualize the file system of both the source and destination remote servers directly within the job definition interface, enabling for rapid and effective file selection.

Transfer smarter with the new capability of the File Transfer integration.

For further information about the integration, browse to [Automation Hub](#).

Centralized JWT Signing Key Management

The signing key for JSON Web Tokens (JWTs) is now centrally stored in the database to ensure secure and consistent token validation across IBM® Workload Scheduler components.

Key storage and usage

The master domain manager now generates its own proprietary key, which it uses to sign certificates. This key is now centrally stored in the database and is therefore shared between the master domain manager and the backup master domain manager, ensuring business continuity in case you need to switch the manager. When you rotate the certificates on the master domain manager, the old tokens remain valid because the signing key remains unchanged.

The keystore includes a new certificate with the alias `mpjwtkey`, which contains the server public key (`tls.crt`). To ensure proper product functionality, this alias must remain unchanged and must not be modified under any circumstances.

Benefits of centralized storage

Centralizing the signing key in the database enables seamless distribution and accessibility across all IBM® Workload Scheduler components. This facilitates secure and uniform token validation throughout the environment.

Key validity and stability

The signing key stored in the database is configured with a validity period of 100 years, ensuring long-term stability and uninterrupted operations across your environment.

Upgrade considerations

If you are upgrading from a previous version, the signing key location is automatically managed by the system. No manual intervention is required—the upgrade process preserves and configures the key as needed.

Certman security enhancements: automated validation & OpenSSL integration

Two key enhancements have been made for more reliable and efficient security management:

Automated certificate verification

The `AgentCertificateDownloader` script now automatically verifies certificates using the `certman verify` feature. This change improves security and helps prevent deployment issues.

The script checks that:

- The certificate, private key, and CA are consistent.
- The certificate is not expired.

This verification runs during initial agent setup and automatic renewal, preventing the installation of invalid or inconsistent certificates.

Simplified certificate generation with OpenSSL syntax

Certman now natively supports OpenSSL syntax.

This enables you to use standard OpenSSL commands and parameters (such as `subject`, `expiration`, `key size`) directly when generating new certificates or a Certificate Authority (CA).

These enhancements provide more robust certificate verification and greater control over certificate generation.

Enhanced OQL filtering: the `STARTED` alias

OQL filtering capabilities are enhanced with the introduction of the new **STARTED** alias for jobs and job streams.

Job type	Workstation (Job)	Job stream	Workstation (Job stream)	Scheduled time	Earliest start	Actual start
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:54:00 AM PST	11/21/2025 2:54:04 AM PST	11/21/2025 2:54:04 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:53:00 AM PST	11/21/2025 2:53:03 AM PST	11/21/2025 2:53:03 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:52:00 AM PST	11/21/2025 2:52:06 AM PST	11/21/2025 2:52:06 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:51:00 AM PST	11/21/2025 2:51:05 AM PST	11/21/2025 2:51:05 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:50:00 AM PST	11/21/2025 2:50:05 AM PST	11/21/2025 2:50:05 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:49:00 AM PST	11/21/2025 2:49:04 AM PST	11/21/2025 2:49:04 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:48:00 AM PST	11/21/2025 2:48:04 AM PST	11/21/2025 2:48:04 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:47:00 AM PST	11/21/2025 2:47:52 AM PST	11/21/2025 2:47:52 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:42:00 AM PST	11/21/2025 2:47:52 AM PST	11/21/2025 2:47:52 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:40:00 AM PST	11/21/2025 2:47:52 AM PST	11/21/2025 2:47:52 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:38:00 AM PST	11/21/2025 2:47:52 AM PST	11/21/2025 2:47:52 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:41:00 AM PST	11/21/2025 2:47:51 AM PST	11/21/2025 2:47:51 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:37:00 AM PST	11/21/2025 2:47:51 AM PST	11/21/2025 2:47:51 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:43:00 AM PST	11/21/2025 2:47:51 AM PST	11/21/2025 2:47:51 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:45:00 AM PST	11/21/2025 2:47:51 AM PST	11/21/2025 2:47:51 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:46:00 AM PST	11/21/2025 2:47:51 AM PST	11/21/2025 2:47:51 AM PST
executable	/WA_DA	JS	/WA_DA	11/21/2025 2:44:00 AM PST	11/21/2025 2:47:50 AM PST	11/21/2025 2:47:50 AM PST

This feature significantly simplifies query writing and improves readability. It streamlines your filtering process by enabling a simple, case-insensitive keyword to replace more complex OQL filter logic.

The `STARTED` alias enables you to directly filter objects based on whether they have an *actual start value*, that is, whether they have started running.

For example, when monitoring jobs in the **Orchestration Monitor**, you can now use a query as simple as `STARTED = true` to find all jobs that have started. Additionally, you can create more refined queries by combining the alias with other OQL syntax keywords and fields for more in-depth filtering:

```
name LIKE '@job1' AND jobOptions.critical = true AND started = false
```

This provides a more intuitive way to check the run status of objects in the plan. For more information, see [Using Orchestration Query Language](#)

New in Orchestration Monitor: update agent certificates remotely

The **Update agent certificate** action is now available on the Orchestration Monitor, giving you a single, centralized way to refresh certificates across your environment.

This feature simplifies security maintenance and ensures your dynamic agents can quickly and reliably establish trust.

Key functionality

- **Quick access:** The **Update agent certificate** action is available when you select a dynamic agent workstation running version 10.2.6 or later, and click **Actions**. This enables users to initiate the remote certificate download.
- **Direct download from primary domain manager:** When triggered, this action performs the update of the remote dynamic agent certificates on all eligible workstations. The selected dynamic agent contacts

the primary domain manager to directly download new certificates from the master domain manager `<data_dir>/ssl/depot` folder, replacing the existing certificates in the dynamic agent local truststore.

- **Simplified trust for external services:** If you create the sub-folder `<data_dir>/ssl/depot/additionalCAs` on the master domain manager and populate it with public certificates for external services, the agent automatically downloads them and adds them to its local trust store. This ensures simple and immediate trust establishment for any services dependent on those auxiliary certificates.

Automatic recovery and rollback

To guarantee continuity of service, an automatic recovery mechanism has been implemented.

After the agent downloads and imports the new certificates, it attempts a primary connection to the reference primary domain manager. If the connection fails to establish within a period of 5 minutes, or if the downloaded certificates are expired or incorrectly formatted, the agent automatically reverts to the previous functional certificates. This process protects the agent from being permanently isolated due to a failed update.

Verification

You can verify the success of the agent certificate update process by checking the **Certificate expiration** column on the Orchestration Monitor. You can add the **Certificate expiration** column by modifying the table configuration.

For more information, see Orchestration Monitor overview

Referencing operator instructions in event rules

You can use operator instructions as variables for job and job stream events of `TWSObjectsMonitor` type.

By enabling dynamic data extraction from job and job streams definitions, this feature enhances event-driven workload automation (EDWA) processes, significantly improving event rules flexibility and reusability. This enhancement provides greater speed and efficiency in automation, substantially reducing the number of event rules otherwise required to cover multiple objects.

The feature supports the direct use of information contained within the **description** and **documentation** fields of jobs and job streams definitions as variables in EDWA actions. Such operator instructions can include runbook steps, support contacts, and recovery procedures.

The information is accessed via variables specifically available for `TWSObjectsMonitor` event actions related to jobs and job streams. You can choose to retrieve the entire set of instructions or specific, customized key-value pairs.

- **Retrieving the entire field content:** Retrieve the complete text from the operator instructions fields, typically used for sending full recovery procedures:

```
%{<event>.OpInstr}
```

- **Retrieving specific key-value pairs:** Extract custom metadata defined within the object description, for example support contacts:

```
%{<event>.OpInstr.<key>}
```



Note: The key-value pair must be present within the **description** or **documentation** fields of the job or job stream definition. The syntax requires a key, a separator (either = or :), and a value. For example:

```
support_team=help@wa.it
```

Key benefits

- **Improved efficiency:** Processes become more rapid due to automated, context-aware actions.
- **Rule consolidation:** The number of unique event rules required is drastically decreased, as one generic rule can cover multiple events across different jobs and job streams.
- **Enhanced recovery:** Provides direct access to operational data and support contacts at the moment an event occurs.

For more information, see TWSObjectsMonitor events.

Job stream view enhancements

Overview of job stream view

The latest enhancements to the job stream view prioritize user control, visual feedback, and operational efficiency. The updated interface provides a more stable and responsive environment for managing complex workloads.

- **View stability**
 - **Layout locking:** Freezes node coordinates to facilitate navigation and panning without causing disorienting layout shifts.
- **Operational efficiency**
 - **Multi-selection mode:** Toggles between single navigation and bulk selection modes. Selected items display with high-contrast highlighting for improved visibility at any zoom level. Users can perform simultaneous actions on multiple jobs with a confirmation summary of the operation status.
- **User feedback**
 - **Visual loading state:** Displays a distinct loading indicator during data-intensive processing. This visual confirmation alerts users that background calculations are active, reducing uncertainty during layout updates.
- **View export accuracy**
 - **SVG export functionality:** Generates scalable vector graphics that accurately reflect the current on-screen layout, operational status, and node positions.

Benefits of job stream view enhancements

These enhancements streamline daily operations and improve the usability of the graphical interface:

1. **Bulk processing:** Users can perform actions on multiple jobs simultaneously within the job stream view, removing repetitive manual tasks. The system provides immediate confirmation upon completion of the bulk action.
2. **Enhanced offline analysis:** The improved SVG export supports extensive graphs containing hundreds of jobs. Text, icons, and dependencies remain fully readable in the exported file, facilitating offline review and documentation.
3. **Navigation control:** Layout locking prevents accidental displacement of items or cells during navigation. Users can explore the job stream structure freely without altering the defined visual arrangement.

Changed features, improvements, and RFEs in version 10.2.6

A list of features changed and improved since the previous release and major improvements.

Changed and improved features

The following features have changed or have been introduced or removed since the previous release:

Support for Oracle PeopleTools 8.62.04

IBM® Workload Scheduler supports Oracle PeopleTools version 8.62.04

Support for Java Runtime Environment version 21 on IBM i agents

IBM i agents now support Java Runtime Environment 21, enabling applications to leverage the latest LTS Java runtime. This update provides enhanced performance, new language features, and improved security. Compatibility with Java Runtime Environment version 11 is seamlessly maintained in both fresh install and upgrade scenarios.

Requests for enhancements (RFEs)

Requests for Enhancements (RFEs) give customers the opportunity to collaborate directly with the product development team and other users. The team prioritizes and develops new product features based on proposals made by customers.

IBM® Workload Scheduler V 10.2.6 delivers the following RFEs:

WAO-I-417: Pre-installation disk space validation for all IBM® Workload Scheduler components

To ensure a smooth and successful installation experience, the program now performs an automatic disk space check before the installation begins. This proactive step helps prevent interruptions by verifying that your system meets the minimum disk space requirements upfront.

WAO-I-625: SSH protocol now renamed to SFTP protocol in the File Transfer integration

In the File Transfer integration, the **SSH protocol** has been renamed to **SFTP** (SSH File Transfer Protocol) for improved clarity. The core functionality is unchanged. For more information, see [File Transfer integration](#).

To view a complete list of RFEs, new, planned, and delivered, see: [RFE online community](#).

IBM Workload Scheduler version 10.2.5 enhancements

Learn about the IBM Workload Scheduler version 10.2.5 features.

IBM Workload Scheduler version 10.2.5 includes the following enhancements:

- [Advancing security with FIPS 140-3 compliance on page 31](#)
- [Streamlined engine reporting configuration in administration panel on page 32](#)
- [Navigating OQL Queries in the Orchestration Monitor on page 33](#)
- [User Object Credential Management - Ensuring operational continuity and security on page 33](#)
- [Monitor certificate expiration - Total certificate control on page 34](#)
- [Changed features, improvements, and RFEs in version 10.2.5 on page 34](#)

Advancing security with FIPS 140-3 compliance

FIPS is a U.S. government security standard that defines **security requirements for cryptographic modules** used to protect sensitive information. **FIPS 140-3** is the latest version of the U.S. and Canadian government security standard that defines security requirements for cryptographic modules in IT and telecommunications products. Its purpose is to ensure that products handling sensitive data via cryptography are secure and reliable. Federal agencies must use FIPS 140-3 validated modules. It is also widely adopted by defence contractors and financial institutions.

The introduction of FIPS 140-3 compliance ensures a robust and secure automation solution. This significant enhancement enables organizations, particularly those in regulated industries and federal sectors, to leverage IBM® Workload Scheduler while adhering to stringent U.S. government security standards for cryptographic modules.

IBM® Workload Scheduler now natively supports FIPS 140-3 by restricting cryptographic algorithms to the permitted subset, ensuring that data protection mechanisms align with these rigorous specifications, while carefully managing necessary, FIPS-permitted exceptions during transition.

A new and improved version of the secure command empowers administrators to activate or deactivate FIPS compliance post-installation. Key capabilities of secure include:

FIPS activation

Configures the environment for FIPS 140-3 operation. This includes verifying compatibility (for example with Windows users) and converting any existing passwords stored with the older TripleDES algorithm to the FIPS-approved AES encryption standard. It then activates FIPS flags across essential configuration files and restarts the affected components. The activation process can accommodate transitional phases for algorithms like TripleDES, as permitted, before full conversion.

FIPS deactivation

Reverts the FIPS-specific settings in the aforementioned configuration files and restarts the components, should FIPS mode need to be disabled.

The introduction of integrated FIPS 140-3 compliance offers significant advantages:

Enhanced Security Standards

Operate your environment in alignment with rigorous, federally mandated cryptographic standards.

Simplified Regulatory Adherence

Enable organizations in government, finance, healthcare, and other regulated sectors to confidently deploy IBM® Workload Scheduler.

Dedicated Compliance Management

Utilize the secure script for straightforward and centralized control over FIPS activation and deactivation.

Robust Data Protection

Ensure sensitive data, particularly passwords, are protected using FIPS-validated encryption algorithms like AES.

Flexible and Controlled Migration

Benefit from a clear process for enabling FIPS in new or existing environments, with guidance for migrating older FIPS-configured setups.

With FIPS 140-3 compliance, IBM® Workload Scheduler provides an even more secure and trustworthy platform, empowering you to meet demanding security requirements while automating your critical business processes.

To ensure FIPS compliance, all IBM® Workload Scheduler components must meet the following requirements:

- All components must be at version 10.2.5 or later.
- The certificates must employ at least a robust 2K RSA key and use encryption algorithms different from `MD5-RSA` and `SHA1-RSA`.
- The key must be in a format where the algorithms are supported by FIPS. For example, avoid the PKCS1 format with the MD5 algorithm or the PKCS8 format with the 3DES algorithm.

FIPS is supported on all supported operating systems with the exception of IBM i operating systems.

FIPS mode is primarily for meeting specific government security standards. If your organization does not need to enable it, the product continues to deliver robust security without it. If your organization does not require adhering to FIPS, you can install or upgrade as usual and FIPS will not be enabled in your environment.

When you perform a fresh installation, FIPS is disabled by default: you can enable it by setting the `enablefips` parameter to `true` when running the installation commands. When you upgrade from a previous version, IBM® Workload Scheduler checks your current FIPS settings and applies the same settings in the upgraded environment. Before you start the upgrade, ensure your certificates meet FIPS 140-3 standards. FIPS is supported on all supported operating systems with the exception of IBM i operating systems.

For more information about FIPS, see Enabling and disabling FIPS.

Streamlined engine reporting configuration in administration panel

Discover a more intuitive and controlled method for managing reporting credentials on a per-engine basis. The primary goal is to provide users with a dedicated and user-friendly interface to configure and update database access required for reporting functionalities tied to specific engines.

Previously, managing these specific credentials might have involved less direct methods. Now, administrators can easily handle this task: upon selecting an engine from the Administration Panel and clicking the edit icon, the familiar edit modal appears, now featuring a distinct "Reporting" tab. Within this tab, users can enable reporting for the selected engine and

securely input the necessary Database User ID and Password. Successful saving of these credentials is confirmed instantly via a toast notification, ensuring a clear and immediate feedback loop.

This refinement offers clear advantages:

Enhanced User Experience

Provides a dedicated, easy-to-navigate interface for reporting setup.

Centralized Control

Simplifies the management of reporting credentials for each engine directly within its configuration.

Improved Efficiency

Speeds up the process of configuring or updating reporting details.

This focused update improves the administrative experience by making the setup and maintenance of engine-specific reporting credentials more straightforward and secure.

Navigating OQL Queries in the Orchestration Monitor

The Orchestration Monitor OQL query functionality has been significantly enhanced to improve usability and performance. The new design empowers you with greater control and a more intuitive interface for managing your queries.

When interacting with the query input section, you can now:

- Manually trigger searches using a dedicated search icon or by simply pressing Enter. This removes the automatic running of queries while typing.
- Expand or collapse additional filters via an interactive visual icon, enabling a cleaner and more focused workspace.

The ability to precisely control query runs and filter options means you can now craft highly specific queries for workstations, job streams, jobs, and more without interruption. This leads to faster issue resolution and more efficient operational management.

User Object Credential Management - Ensuring operational continuity and security

The User Object Credential Management enables you to securely update the passwords stored within user objects directly from a central point of control. This ensures that changes are applied consistently across both the database and plan instances, guaranteeing operational continuity and better adherence to your security policies.

Previously, updating the credentials stored in user objects was often a manual, multi-step process. These password changes had to be applied separately to the database and then manually propagated to any active plans, a procedure prone to human error and delays. This disconnect could lead to authentication failures for in-flight tasks that rely on these objects, causing unexpected job failures, disrupting critical business processes, and creating potential security compliance gaps.

The introduction of this feature brings significant advantages:

Limit service disruptions

Update credentials for user objects across both the database and all running plan instances with a single action. This prevents authentication failures and ensures your job streams continue to run without interruption.

Strengthen security and compliance

Enforce password policies consistently for the credentials within your user objects and respond quickly to security mandates. Centralized control simplifies audits and demonstrates adherence to your organization's security standards.

Increase operational efficiency

Drastically reduce the time and effort required to manage user object credentials. What was once a complex, multi-system task becomes a single, swift operation, freeing up your team to focus on more strategic initiatives.

User Object Credential Management enables more resilient, secure, and efficient automation.

Monitor certificate expiration - Total certificate control

The Monitor certificate expiration provides a centralized view of your certificates directly within the Orchestration Monitor. This gives you immediate visibility into certificate validity, making it simple to track expiration dates and proactively manage the rotation process to maintain security and service availability.

The introduction of proactive certificate monitoring brings a number of benefits:

Prevent service outages

Gain a clear, at-a-glance view of all monitored certificates and their expiration dates. By receiving advance warning, you can plan and perform certificate rotations with zero downtime, preventing outages caused by expired credentials.

Enhance security and compliance

Ensure your environment is always protected with valid, trusted certificates. Proactive monitoring simplifies security audits and demonstrates a robust, compliant process.

Streamline certificate management

By centralizing certificate information within the Orchestration Monitor, you reduce administrative overhead and simplify the entire lifecycle management for your team.

By integrating certificate oversight directly into your daily monitoring, this feature transforms a manual task into a simple part of your operational routine, bolstering both the security and reliability of your automation.

Changed features, improvements, and RFEs in version 10.2.5

A list of features changed and improved since the previous release and major improvements.

The following features have changed or have been introduced or removed since the previous release:

New command for setting Dynamic Workload Console environment

You can now use the `dwc_env` command to set the environment on the Dynamic Workload Console. Launch the command as follows: `dl`

Windows On Windows operating systems

```
dwc_env.cmd
```

UNIX On UNIX operating systems

```
./dwc_env.sh
```

After setting the environment, you can launch the commands located on the Dynamic Workload Console (`secure` and `certman`) without specifying the command path.

The command is located in the installation directory of the Dynamic Workload Console.

The `certman` command can now verify the security level of certificates in your environment

The check on certificates is performed automatically at installation or upgrade time. If the current certificates are not secure, a warning message is displayed.

New databases supported

The following databases are now supported:

- Amazon RDS for Postgres
- Google Cloud SQL for Postgres

These databases offer a powerful combination of flexibility, scalability, ease of use, and performance together with high availability and reliability.

Support for master domain manager and Dynamic Workload Console now available on AIX operating systems

You can now install master domain manager and Dynamic Workload Console on supported AIX operating systems. For an up-to-date list of supported operating systems, see [Supported operating systems](#).

Requests for Enhancements (RFEs) give customers the opportunity to collaborate directly with the product development team and other users. The team prioritizes and develops new product features based on proposals made by customers.

IBM® Workload Scheduler V 10.2.5 delivers the following RFEs:

User object credential management

The User object credential management enables you to securely update the passwords stored within user objects directly from a central point of control in the Dynamic Workload Console. You can also perform the same action from the Orchestration CLI with a simple command, for more information see `altpass`.

Monitor certificate expiration WAD-I-118

The monitor certificate expiration provides a centralized view of your workstation certificates directly within the Orchestration Monitor. This enables you to track expiration dates and proactively manage the rotation

process to maintain security and service availability. For more information on the Orchestration Monitor, see [Orchestration Monitor overview](#)

Loop detection between jobs and job streams in current and trial plans WAO-I-553

Ensure robust detection of loops involving nodes (that is jobs and job streams) in both the current and trial plans. By enabling the new property `com.ibm.tws.planner.monitor.verboseLoopLogging=true` in the `TWSConfig.properties` file, the product logs loop-related notifications in the `trace.log` file whenever loops are detected.

To minimize log verbosity, notifications are only generated when there is a change in loop structure—such as a job or job stream entering or exiting the loop.

For more information, see [Detecting loops between jobs and job streams in current and trial plans](#)

To view a complete list of RFEs, new, planned, and delivered, see: [RFE online community](#).

IBM Workload Scheduler version 10.2.4 enhancements

Learn about the IBM Workload Scheduler version 10.2.4 features.

IBM Workload Scheduler version 10.2.4 includes the following enhancements:

- [Maximize workflow efficiency with advanced job stream submission and management on page 36](#)
- [Filtering with Orchestration Query Language in the Orchestration Monitor on page 37](#)
- [Operator instructions within the Orchestration Monitor on page 38](#)
- [Discover the new and improved integration with password vaults on page 39](#)
- [Improved security for password encryption on page 40](#)
- [Unlocking the enhanced capabilities of Certman on page 41](#)
- [Workstation management within the Graphical Designer on page 42](#)
- [Run cycle groups available as triggers in the Graphical Designer on page 42](#)
- [Resources available in the Graphical Designer on page 43](#)
- [Workflow details available in the Orchestration Monitor on page 43](#)
- [Enhanced job control post-submission on page 44](#)
- [Performance enhancements on page 45](#)
- [Changed features, improvements, and RFEs in version 10.2.4 on page 45](#)

Maximize workflow efficiency with advanced job stream submission and management

Transform the way you handle automated workflows with cutting-edge features designed to streamline submissions, enhance flexibility, and offer greater control over your orchestration processes. To achieve this, you can use the Job Stream Submission job, which is designed to automate the submission of job streams for processing. The new version of this job features the following improvements:

Hierarchical orchestration

Unlock advanced automation scenarios through nested job streams. This powerful feature enables you to build intricate, multi-layered workflows by embedding job streams within others. This enables precise orchestration of complex sequences of interdependent jobs and job streams, providing a structured approach to managing sophisticated processes.

Automated submission efficiency:

Streamline and optimize your operational flow by automating the submission of job streams for processing. This enhancement reduces manual intervention, accelerates workflow initiation, and improves overall efficiency by ensuring timely execution of your automated processes.

Real-time monitoring

Maintain comprehensive awareness of your active workflows with detailed, real-time monitoring accessible via the Workflow details view in the Dynamic Workload Console. This provides immediate insight into the status of each job stream, enabling proactive issue identification and resolution, ensuring smooth and predictable operations.

Flexible management options

Gain granular control over running job streams with enhanced drill-down capabilities within the Workflow details view. This enables you to investigate the status of individual components and, crucially, to restart the Job Stream Submission job as needed. This dynamic management capability ensures operational continuity and enables a swift recovery from unexpected issues.

Customizable Behavior

You can now configure whether the Job Stream Submission job actively monitors the run of the submitted job streams and adjusts its own status based on their final outcome. Alternatively, you can maintain the previous behavior where the submission job completes successfully regardless of the downstream job stream results, providing flexibility in how you track and manage workflow dependencies.

These advancements in job stream submission and management empower you with unparalleled precision, efficiency, and control over your automated workflows. By streamlining submissions, enabling sophisticated nested structures, providing real-time visibility, offering flexible management options, and allowing for customizable behavior, these features simplify even the most complex processes and significantly enhance your orchestration capabilities.

Filtering with Orchestration Query Language in the Orchestration Monitor

Harness the power of OQL for enhanced monitoring.

Experience the power of Orchestration Query Language (OQL) directly in the Orchestration Monitor of the Dynamic Workload Console.

Toggle the OQL button in the Orchestration Monitor and effortlessly filter workstations, jobs, job streams, and resources with comprehensive syntax, tailoring your view for superior monitoring and gaining precise insights for optimal workload management.

Operational efficiency through precise monitoring

By leveraging OQL in the Orchestration Monitor, you can filter workstations, jobs, job streams, and resources with comprehensive syntax. This precise filtering enables operations teams to quickly pinpoint specific areas of interest or concern in your environment, reducing the time spent sifting through irrelevant data and leading to faster identification and resolution of issues.

Faster issue resolution and reduced downtime

OQL empowers you to easily and quickly retrieve required results by writing targeted queries. This rapid access to specific information about the production plan environment enables operations teams to diagnose problems more efficiently, leading to quicker resolutions and minimizing potential downtime, which directly translates to increased system availability and business continuity.

Streamlined operations across multiple interfaces

OQL is not limited to a single interface; it can be used with REST API V2, Orchestration CLI, and the Orchestration Monitor of the Dynamic Workload Console. This consistency across different tools provides flexibility for various user roles and automation needs, streamlining operations and reducing the learning curve for different interaction methods, leading to increased team productivity.

Get access to OQL and enhance your monitoring capabilities, improve operational efficiency, reduce downtime, optimize workload management, and proactively address potential issues. For more information, see [Using Orchestration Query Language](#).

Operator instructions within the Orchestration Monitor

The Orchestration Monitor now features the **Operator instructions** panel.

Operators can now retrieve definition details for jobs and job streams with the new **Operator Instructions** panel within the Orchestration Monitor. The panel displays information from the **Description** and **Documentation** fields of the job or job stream definition, to facilitate a rapid comprehension of item parameters during monitoring activities.

Enhanced operational efficiency

Operators can quickly access crucial definition information about jobs and job streams directly within the Orchestration Monitor. This rapid access to details from the **Description** and **Documentation** fields facilitates a faster comprehension of job or job stream parameters during monitoring activities, leading to more efficient operations.

Reduced downtime

By providing operators with immediate access to relevant information, the **Operator Instructions** panel enables them to better understand the context and purpose of running jobs and job streams. This improved understanding can lead to more effective monitoring, quicker identification of potential issues, and ultimately, a reduction in system downtime.

Faster issue resolution

When an issue arises, operators can instantly retrieve definition details for the affected job, job instance, and the associated job stream. This readily available information can significantly speed up the troubleshooting

process, allowing operators to diagnose and resolve problems more quickly, minimizing business impact. This deeper understanding empowers operators to make more informed decisions during monitoring and when taking necessary actions.

Simplified sharing and onboarding

The **Operator Instructions** feature centralizes key information about jobs and job streams, making it easily accessible to all authorized personnel. This can streamline knowledge sharing within operations teams and facilitate the onboarding of new operators by providing them with immediate access to essential details.

The **Operator Instructions** panel gives operators the necessary details to improve their monitoring effectiveness, and achieve a faster comprehension of item parameters, ultimately contributing to a more efficient and informed operational environment.

For more information about OQL syntax, see [Orchestration Monitor overview](#).

Discover the new and improved integration with password vaults

Take password management to a new level with the new integrated capabilities for password storage and management.

You can now configure one or more **dynamic agents, pools, or dynamic pools** to interact with CyberArk or other password vaults and obtain the passwords required to run your jobs. Employing pools or dynamic pools guarantees high availability by automatically selecting the best-suited agent to contact the password vault.

Enhanced security

The new integration provides a centralized and secure solution for managing credentials, significantly reducing the risk of unauthorized access and security breaches.

Reduced operational overhead

Automating password retrieval through integration with password vaults removes the manual effort involved in managing and updating passwords across numerous jobs and systems, freeing up staff for more strategic tasks.

Increased efficiency and reliability

Jobs are less likely to fail due to incorrect or expired passwords, as the system automatically retrieves the current credentials from the vault at runtime. The high availability ensured by using pools or dynamic pools minimizes disruptions.

Flexibility and scalability

The support for multiple password vaults and custom scripts enables integration with existing security infrastructure and easy scaling of automation efforts.

Smooth integration

Dynamic agents, pools, and dynamic pools can be configured to interact with a chosen password vault to obtain necessary passwords at runtime.

Easy configuration

The ability to create multiple profiles on agents provides flexibility for different password vaults or specific queries.

Streamlined upgrade

Existing users on versions 10.2.1 and later benefit from an automatic migration of their CyberArk configurations, simplifying the adoption of the new feature.

This feature requires all IBM® Workload Scheduler components to be at version 10.2.4 or later.

In conclusion, the new password vault integration in IBM® Workload Scheduler offers a compelling value proposition by providing a secure, efficient, and flexible solution for managing passwords in automated workflows, directly addressing key security and operational challenges associated with managing a growing IT landscape.

For more information about configuring dynamic agents to function as proxies by creating a profile, see the topic about configuring the agent to work with a password vault in *Administration Guide*.

For more information about specifying the required information in the job definition, see Obtaining passwords from password vaults the topic about obtaining passwords from password vaults in *User's Guide and Reference*.

Improved security for password encryption

Ensure a higher security level with the AES algorithm. This feature is especially convenient for encrypting the passwords stored in `.opts` configuration files used for PeopleSoft and SAP.

The Triple DES encryption algorithm has been deprecated by NIST (National Institute of Standards and Technology) and is no longer recommended for new systems. It has been replaced with the AES (Advanced Encryption Standard) encryption algorithm, which offers the following advantages:

Enhanced Security and Reduced Risk

The latest version uses the Advanced Encryption Standard (AES) algorithm as the default for password encryption, which offers a higher level of security compared to the deprecated Triple DES (3DES). AES supports longer key lengths (128, 192, and 256 bits), making encrypted passwords more resistant to brute-force attacks. This significantly reduces the risk of unauthorized access to sensitive systems integrated with IBM® Workload Scheduler, protecting valuable data and the integrity of critical business processes.

Improved Compliance and Adherence to Modern Standards

NIST now recommends AES as the current encryption standard. By adopting the latest IBM® Workload Scheduler, clients ensure they align with modern security best practices and industry standards, potentially aiding in meeting regulatory compliance requirements.

Increased Efficiency and Reduced Operational Costs

AES is significantly faster and more efficient than 3DES, especially in software implementations. This leads to quicker password encryption and decryption processes, consuming less computational power and memory, potentially resulting in faster processing times and reduced expenses for resources and processing time.

Easier Transition and Minimal Disruption

The upgrade process is designed to be seamless. If clients were previously using 3DES, their encrypted passwords are automatically and seamlessly converted to the new AES algorithm during the upgrade, requiring no manual intervention. This minimizes disruption to existing operations and reduces the burden on staff.

Flexible and Versatile Password Management

A secure script that can be used to optionally encrypt passwords for SAP or PeopleSoft, even after the upgrade. This script offers flexibility, allowing administrators to encrypt passwords used during installation, upgrade, and management processes. Clients can choose to use a custom passphrase for potentially reusable passwords or leverage the standard encryption method with a unique key for each installation, further enhancing security by preventing decryption across different components.

The AES encryption provides stronger security, improved efficiency, better compliance, and a seamless upgrade process, ultimately contributing to a more secure and reliable automation environment for critical business operations. The availability of the secure script further empowers administrators with flexible password management options, reducing potential vulnerabilities and enhancing overall system security.

For further details, see the topic about automatic encryption for key product files in *Administration Guide*.

Unlocking the enhanced capabilities of Certman

Experience a more efficient certificate management, powered by Certman.

The new version of IBM® Workload Scheduler features extended Certman possibilities, including the `delete` command, which provides the capability to remove aliases from a keystore or truststore on a master domain manager, an agent (using the newly implemented `agentscope` keyword), or the Dynamic Workload Console.

Enhanced certificate management

The `certman delete` command enables the removal of aliases from keystores or truststores on various components (master domain manager, agents, Dynamic Workload Console). This streamlines certificate management tasks, saving administrators time and effort compared to potentially more complex manual procedures. Efficient certificate management translates to reduced operational overhead and faster response times for security-related tasks.

Improved security posture

The ability to specifically delete aliases, whether from the keystore, truststore, or both, on different IBM® Workload Scheduler components (including agents via the `agentscope` keyword) provides more precise control over security certificates. This granular control enables you to maintain a cleaner and more secure certificate environment by removing unnecessary or outdated aliases, reducing potential security vulnerabilities.

Simplified compliance

By enabling the easy removal of unused or expired certificate aliases, the `certman delete` command helps you maintain a more organized and up-to-date certificate inventory. This simplified management contributes to better compliance with security policies and regulations, making audits smoother and less time-consuming.

By upgrading to the latest version of IBM® Workload Scheduler, you can leverage these enhanced Certman capabilities to achieve more efficient, secure, and manageable certificate operations, ultimately benefiting from reduced operational costs, improved security, and better compliance.

For more information about the `certman delete` command, see [Remove an alias from the keystore or truststore](#). For further and more general Certman information, see [Managing certificates using Certman](#).

Workstation management within the Graphical Designer

You can now create and manage workstations directly within the Graphical Designer, with a unified and visual approach to your scheduling infrastructure. This provides significant operational advantages:

Centralized workstation control

By creating and managing workstations directly within the Graphical Designer, you gain a consolidated and visual perspective of your entire scheduling environment. This centralized approach minimizes the potential for errors and provides unparalleled control over your automation.

Simplified workstation lifecycle management

The Graphical Designer streamlines the entire workstation lifecycle. Upon initiating creation, the system guides you through selecting the workstation type, including Pool, Dynamic Pool, Extended Agent, and more. After entering the necessary details, the new workstation is added to the Assets list, making it immediately available for scheduling. Furthermore, the Assets tab provides consistent functionality to effortlessly edit, duplicate, or delete workstation definitions, simplifying ongoing maintenance and configuration.

The integrated workstation management within the Graphical Designer delivers a centralized and visual approach to administering your scheduling environment. By providing simplified creation and lifecycle management, this feature minimizes errors and grants you unprecedented control over your automation infrastructure, leading to a new level of operational efficiency.

See [Managing workstation definitions](#) to learn how to create workstations from the Graphical Designer interface.

Run cycle groups available as triggers in the Graphical Designer

You can now create and assign run cycle groups as **Triggers** to job streams directly in the Graphical Designer to design a precise, time-driven automation.

This intuitive integration provides significant operational advantages:

Enhanced scheduling precision

By visually defining and managing run cycle groups directly within the Graphical Designer, you gain unprecedented control over when your job streams run. This removes the complexities and potential errors associated with manual scheduling configurations. The visual interface makes it easy to define intricate scheduling patterns, ensuring that your automation runs exactly when needed, minimizing the risk of incorrect run times.

Streamlined design

Integrating run cycle group creation and assignment into the Graphical Designer simplifies the entire automation design process. You can now define both the tasks and their precise run schedule within a single, intuitive environment. This ease of use empowers you to build and deploy a time-driven automation more quickly and effectively.

The integration of Run cycle groups as **Triggers** within the Graphical Designer offers a powerful and intuitive approach to time-driven automation. By enhancing scheduling precision, streamlining workflow design and freeing up valuable resources, this feature enables you to achieve a greater operational efficiency and focus on strategic growth.

Resources available in the Graphical Designer

You can now associate resources with jobs and job streams directly within the Graphical Designer to build a responsive automation framework attuned to immediate resource accessibility.

This delivers notable operational benefits:

Improved automation responsiveness

By visually linking and managing resource dependencies directly within the Graphical Designer, you achieve superior control over when your jobs and job streams run. The visual interface facilitates the definition of clear resource needs, guaranteeing that your automation operates only when the necessary resources are available, preventing failures due to resource constraints.

Simplified workload creation

Integrating resources into the Graphical Designer streamlines the entire automation design process. You can now define jobs and job streams and their resource dependencies within a unified, user-friendly environment. This ease of use enables you to construct and implement resource-aware automation more rapidly and efficiently. Furthermore, the Orchestration Monitor provides a centralized view for tracking resource status and dependencies, enhancing visibility and control.

The integration of resource management in the Graphical Designer empowers you with a more intuitive, efficient, and agile approach to automation. It simplifies complex workloads, reduces errors, and improves overall productivity.

Workflow details available in the Orchestration Monitor

Elevate your operational efficiency with the introduction of Workflow Details in the Orchestration Monitor. This new feature provides unprecedented real-time visibility into the granular execution of your automation.

Enhanced operational insight and proactive management

Workflow details enable you to track the precise completion percentage of data transfers, monitor the status of individual steps within intricate business processes, and receive immediate updates on key activities. This granular insight empowers your operations teams to proactively manage their workload. Potential bottlenecks or delays become instantly apparent, enabling swift adjustments and preventing disruptions before they impact critical business outcomes. This proactive approach translates directly into improved service levels and reduced downtime.

Accelerated issue resolution

When issues do arise, Workflow details provide the necessary context for rapid diagnosis and resolution. Instead of sifting through logs, your teams can pinpoint the exact stage where a problem occurred, understand the specific details, and implement corrective actions swiftly. This dramatically reduces the impact of failures on business operations and ensures business continuity.

Optimized resource allocation:

You can identify processes that are taking longer than expected or consuming excessive resources, allowing for optimization efforts. This leads to increased efficiency, reduced operational costs, and the ability to handle a greater volume of work with existing resources.

By providing granular, real-time insights, Workflow details empowers businesses to enhance operational efficiency, accelerate issue resolution, optimize resource utilization, and gain unprecedented control over their integrated environment, ultimately contributing to improved business agility and resilience.

Enhanced job control post-submission

Quick job definition adjustments in the Orchestration Monitor and Job Stream View.

Gain unprecedented control over your jobs with the newly introduced monitoring actions. Directly within the Orchestration Monitor or Job Stream View, you can now control job processing by modifying the job definition without the need to go back to the original database definition. Simply use the following actions to modify the future instances of the selected job definition:

- Set hold for definition
- Remove hold for definition
- Set no operation for definition
- Remove no operation for definition

Increased operational agility and responsiveness

This feature enables you to make quick adjustments to job definitions directly from the Orchestration Monitor or Job Stream View. This removes the time-consuming process of navigating back to the original database definition, enabling faster responses to changing business requirements and immediate control over future job executions.

Improved efficiency and reduced administrative overhead

By enabling modifications from the monitoring interface, you can streamline the job definition adjustment process. Operations teams can save significant time and effort by avoiding the need to access and modify database records, leading to increased efficiency and lower administrative overhead.

Enhanced control

being able to set or remove **Hold** and **No operation** directly from the monitoring views provides unprecedented control over how future instances of jobs are processed. This enables proactive management of the workload based on real-time conditions.

For more information, see [Modifying job definitions from job instances](#).

Performance enhancements

The following improvements have been introduced with the new release:

Improved dynamic agent memory management

IBM® Workload Scheduler version 10.2.4 introduces significant improvements in dynamic agent memory handling. Under substantial workloads of up to 600,000 job submissions, the agent now exhibits negligible memory consumption and efficiently releases resources.

This enhancement provides greater stability and resilience for your workload automation environment, reduces the risk of memory-related issues, and enables higher job throughput without impacting system performance.

Orchestration Monitor query performance boost

The Orchestration Monitor now provides significantly faster access to critical workload information. Query results, based on tests with 250 entries, appear three times as quickly compared to the Workload Monitor.

This improved performance gives you quicker insights into your automation. Faster query times mean less waiting, enabling faster identification of potential issues, quicker decision-making, and improved overall operational efficiency. This enables more proactive monitoring and a more responsive monitoring experience.



Note: Data taken from test environments. The time represents the user process duration, starting from opening the monitor interface to obtaining a query result with 250 entries.

Changed features, improvements, and RFEs in version 10.2.4

A list of features changed and improved since the previous release and major improvements.

The following features have changed or have been introduced or removed since the previous release:

Drivers for Db2 for z/OS no longer shipped with the product

Whereas previously the drivers for Db2 for z/OS were included in the Dynamic Workload Console package, now they must be obtained from your Db2 for z/OS installation. If you use Db2 for z/OS with the Dynamic Workload Console version 10.2.4 or later, transfer the drivers in binary mode from the directory where you installed Db2 for z/OS to a directory of your choice. Specify this directory using the **dbdriverspath** parameter when running the `configuredb` or `dwcinst` scripts.

Requests for Enhancements (RFEs) give customers the opportunity to collaborate directly with the product development team and other users. The team prioritizes and develops new product features based on proposals made by customers.

IBM® Workload Scheduler V 10.2.4 delivers the following RFEs:

WAD-I-83

Request: CyberArk and password vault management need to be improved and made more effective.

Resolution: Several improvements have been added in password vault management. For more information, see [Discover the new and improved integration with password vaults on page 39](#).

WAO-466

Request: Remove the 1000-character limit from the `Comment` field of job stream definitions.

Resolution: The `Documentation` field, previously known as `Comment`, has been increased to 4000 characters.

WAO-514

Request: A solution to visually differentiate production and non-production environments.

Resolution: With the `Private Label Settings` page, customers can now personalize environments with custom images. For more information, see [Personalizing interfaces with custom images](#)

WAO-544

Request: Add **Operator instructions** for jobs and job streams, as already present on IBM Z Workload Scheduler.

Resolution: The Orchestration Monitor now features the **Operator instructions** panel. For more information, see [Orchestration Monitor overview](#)

WAO-574

Request: A solution that prevents IBM® Workload Scheduler stored audit files from containing not valid JSON.

Resolution: The content within the audit JSON files is valid JSON. Each individual line now constitutes valid JSON.

WAO-I-3, WAO-I-64, WAO-I-124, WAO-I-190, WAO-423

Request: Enhancing the Dynamic Workload Console with more flexible filtering options, such as combining multiple filter criteria simultaneously, specifying multiple values for a single filter, adding multiple values for the same filter within a widget and so on.

Resolution: The Orchestration Monitor now supports Orchestration Query Language filters to enhance monitoring activities. For more information, see [Using Orchestration Query Language](#).

ZWS-I-238

Request: Adding a new administrative role to the Dynamic Workload Console to determine which rights and activities are available to each user.

Resolution: The new **Administration Panel** in the Dynamic Workload Console allows for detailed and granular management of engine settings.

To view a complete list of RFEs, new, planned, and delivered, see: [RFE online community](#).

IBM Workload Scheduler version 10.2.3 enhancements

Learn about the IBM Workload Scheduler version 10.2.3 features.

IBM Workload Scheduler version 10.2.3 includes the following enhancements:

- [Discover the Graphical Designer: a brand-new way to design your orchestration on page 47](#)
- [Elevate your data management with PostgreSQL database on page 47](#)
- [Unleash the power of the cloud using Amazon RDS for MSSQL on page 48](#)
- [Customize your interfaces with White Label Settings on page 48](#)
- [Power your OCLI model commands with the newly supported items on page 48](#)
- [Explore the new items now available with the scalable and flexible REST APIs V2 on page 49](#)
- [Embedded job definitions for enhanced customization on page 49](#)
- [Certman: the new tool to manage certificates on page 48](#)
- [Changed features and feature capabilities in version 10.2.3 on page 50](#)

Discover the Graphical Designer: a brand-new way to design your orchestration

The Dynamic Workload Console version 10.2.7 comes with a brand-new Graphical Designer.

The Graphical Designer introduces drag and drop canvas capabilities that improve and simplify your automation design process. It uses a low-code, streamlined palette-based design and intuitive tabs, making editing your canvas workspace simple and efficient. A mini-map is also available to help you confidently navigate over workspaces that contain multiple complex job streams.

If you enjoy a smooth user experience, intuitive operations and an effective automation, the new Graphical Designer is everything you have been waiting for.

For more information, see [Graphical Designer overview](#).

Elevate your data management with PostgreSQL database

Improve database performance and reliability by implementing the PostgreSQL database in your IBM® Workload Scheduler environment.

You can now use the PostgreSQL database to create the database for the master domain manager and Dynamic Workload Console.

PostgreSQL offers an extensive list of features related to performance, security, and configuration options. It is free and open-source, and its high customizability supports tailoring it to the specific needs of your environment. It also provides scalability, availability, and functionality suitable for large volumes of data.

PostgreSQL replaces OneDB and Informix, which are no longer supported.

For more information about using PostgreSQL to create the database for the master domain manager and Dynamic Workload Console, see the topics about creating and populating the database for the master domain manager and Dynamic Workload Console in *IBM Workload Scheduler: Planning and Installation*.

Unleash the power of the cloud using Amazon RDS for MSSQL

You can now create the databases for the master domain manager and the Dynamic Workload Console using Amazon RDS for MSSQL.

Amazon RDS for MSSQL is a powerful, scalable managed database service that makes it easier to set up, operate, and scale SQL Server deployments in the cloud. It provides a powerful and convenient option for managing MSSQL databases in the cloud.

Amazon RDS for MSSQL also simplifies database management, enhances security, and provides the flexibility to scale up or down as your needs change.

For more information about creating the database using Amazon RDS for MSSQL, see the topics about creating the database for MSSQL cloud-based databases for the master domain manager and the Dynamic Workload Console.

Customize your interfaces with Private Label Settings

You can now use **Private Label Settings** to enhance your visual identity in IBM® Workload Scheduler.

Your brand visuals are your business signature. You can incorporate your brand visual identity in the Dynamic Workload Console and Self-Service Catalog interfaces by uploading your logo and icon images in just a few clicks from the **Private Label Settings** page.

Your logo is displayed in the login pages of both Dynamic Workload Console and Self-Service Catalog, while the icon is displayed in the navigation bar of the Dynamic Workload Console.

Start customizing your interfaces today to better match your brand personality.

For more information, see [Personalizing interfaces with custom images](#).

Certman: the new tool to manage certificates

Nowadays, having secure communications is paramount for businesses. To satisfy this crucial need, there is a new tool in IBM® Workload Scheduler named Certman.

Certman enables you to quickly generate, extract, verify, and import your custom certificates based on the Secure Socket Layer (SSL) protocol, making the whole process more intuitive and easier to manage. You only need to run a few simple commands, and everything is automatically set.

Simplify your certificate management with Certman, and your communication security will no longer be a chore!

For further information about the usage of Certman, see [Managing certificates using Certman](#).

Power your Orchestration CLI model commands with the newly supported items

You can now use Orchestration CLI model commands for new items to design your automation more efficiently.

The Orchestration CLI is a stand-alone command-line application you can download and install independently. It provides an efficient and streamlined interface from which you can design and monitor your automation environments.

The Orchestration CLI now supports model commands for the following items:

- Variables
- Variable tables
- Calendars
- Credentials
- Prompts
- Resources
- Run cycle groups
- Event rules
- Workstation classes
- Domains

You can now use the Orchestration CLI to design more sophisticated automation scenarios from a single point of control, increasing your productivity and efficiency without compromising on quality or scope.

Discover more about the Orchestration CLI and the new model commands.

Explore the new items now available with the scalable and flexible REST APIs V2

The powerful and flexible REST APIs V2 can now be used with more items, making more complex automation scenarios possible.

REST APIs enable the integration of workload scheduling capabilities with external products and solutions. With the implementation of our REST APIs V2, this integration is easier to configure, more powerful and flexible.

You can now use the new, more capable REST APIs V2 in conjunction with more items, including:

- Prompts
- Resources
- Run cycle groups
- Workstation classes
- Domains
- Event rules

The extended API V2 support to these items lets you easily create more complex, highly efficient automated environments, enabling IBM® Workload Scheduler to work seamlessly in concert with all the third party products and solutions you need.

Discover how to drive your automation with REST APIs and find out more about the capabilities of the new APIs V2.

Embedded job definitions for enhanced customization

A job stream embedded job definition is a self-contained job configuration that remains immune to external changes and gives precise control and customization within the job stream.

You can now create job definitions within job streams without referencing task templates. Embedded job definitions do not exist as standalone entities, they are solely manageable by editing job stream definitions: this provides greater control and customization capabilities. Embedded jobs contribute to a clutter-free environment by removing the necessity for single-use job templates. For more information, see [Creating an embedded job in a job stream](#).

Changed features and feature capabilities in version 10.2.3

A list of features changed since the previous release

The following features have changed or have been introduced or removed since the previous release:

Scheduling end-to-end with fault-tolerant capabilities is no longer available with IBM Z Workload Scheduler, version 10.2

This change impacts all components which connect to fault-tolerant agents. For instance, when you connect the Dynamic Workload Console to an engine running IBM Z Workload Scheduler, version 10.2, fault-tolerant agents are no longer available.

OneDB and Informix databases replaced by PostgreSQL

OneDB and Informix are no longer supported and support for PostgreSQL has been introduced to replace them. PostgreSQL provides a number of advantages, such as enhanced performance, robust features, and seamless scalability, as described in [Elevate your data management with PostgreSQL database on page 47](#). If you are using OneDB or Informix and want to upgrade to version 10.2.3, you can extract all your definitions from OneDB or Informix and import them into a supported database. For more information, see the topics about connecting the Dynamic Workload Console to a new node or database in *IBM Workload Scheduler: Planning and Installation* and connecting the master domain manager to a new database in *Administration Guide*.

Derby database is no longer supported

Derby database is no longer supported.

If you upgrade from an environment where you are using Derby, you have to install a supported database and extract all your definitions from Derby as described in the topic about connecting the Dynamic Workload Console to a new node or database in *IBM Workload Scheduler: Planning and Installation*.

FIPS compliance temporarily not supported

FIPS compliance is not supported in the current product version. This is because OpenSSL 3.0 libraries do not provide a FIPS-compliant validation algorithm for P12 certificates. Development is in progress with the aim of supporting FIPS in upcoming releases.

An optional parameter named **enablefips** is available in the `serverinst` and `twinst` scripts to check FIPS settings before you install or upgrade. The default value in a fresh installation is `false`.

If you are upgrading from an environment where FIPS compliance is supported, you should disable FIPS configuration. When upgrading, a check is performed on your settings to ensure FIPS is not enabled.

If it is enabled, the upgrade stops and an error message is displayed. To resume the upgrade, set the **enablefips** parameter to `false` when running the `serverinst` and `twinst` scripts. The upgrade proceeds and

the configuration files are updated accordingly. For more information, see the FAQ about FIPS compliance in section FAQ - Upgrade procedures in *Planning and Installation Guide*.

Dynamic Workload Console and master domain manager are no longer supported on AIX operating systems

Both Dynamic Workload Console and master domain manager are no longer supported on AIX operating systems. Support remains unchanged for all agent types. For a complete list, see [Supported operating systems](#)

Red Hat Enterprise Linux (RHEL) 7 out of support

Red Hat Enterprise Linux (RHEL) 7 went out of support in June 2024. As a result, it is not supported in IBM® Workload Scheduler version 10.2.3.

POWER System - Little Endian 64-Exploit, 64-Tolerate

POWER System - Little Endian 64-Exploit, 64-Tolerate is no longer supported.

Certificates now managed using the certman command

While in previous versions it was possible to generate certificates at installation time using the `serverinst` command, the new `certman` command can now manage certificates more efficiently and easily. For more information, see [Certman: the new tool to manage certificates on page 48](#).

IBM Workload Scheduler version 10.2.2 enhancements

Learn about the IBM Workload Scheduler version 10.2.2 enhancements.

IBM Workload Scheduler version 10.2.2 includes the following enhancements:

- [Installation without root privileges extended to all IBM Workload Scheduler components on page 51](#)
- [Orchestration CLI extended to include model commands on page 52](#)
- [Easier configuration for Orchestration CLI on page 52](#)
- [Automation as code: object definitions in JSON and YAML formats on page 53](#)
- [Observability with OpenTelemetry on page 53](#)
- [CyberArk and Kerberos integrations automatically installed on dynamic agents on page 54](#)
- [Serviceability enhancements on page 54](#)
- [Changed features and feature capabilities in version 10.2.2 on page 55](#)
- [Satisfying Requests for Enhancements \(RFEs\) in version 10.2.2 on page 56](#)

Installation without root privileges extended to all IBM® Workload Scheduler components

You can now install all server components (such as master domain managers, domain managers, dynamic domain managers and so on) and fault-tolerant agents using a user account that does not have root privileges.

On supported UNIX operating systems, you can now choose whether to install server components and fault-tolerant agents using a user account with or without root privileges. This is especially useful if security policies in your organization require not to share the password for root users. This installation method is already available since previous releases for dynamic

agents, IBM Z Workload Scheduler Agent (also known as the agent with z-centric capabilities), and Dynamic Workload Console.

Note that if you choose this installation method, known as **no-root installation**, only the user who performs the installation can use IBM Workload Scheduler.

This enhancement applies only to supported UNIX operating systems, while the installation method on Windows operating systems remains unchanged.

For more information, see the topics about user management, agent installation procedure, before upgrading, uninstalling the main components and uninstalling agents using the `twinst` script in *IBM Workload Scheduler: Planning and Installation*. Also see the **uname** parameter in the `twinst` script and the **wauser** parameter in the `serverinst` script.

Orchestration CLI extended to include model commands

Orchestration CLI is a stand-alone command-line application that you can download and install independently without requiring any other IBM® Workload Scheduler component. You can install Orchestration CLI on any workstation where you want to manage and control workflows. It is designed to replace the `composer` and `conman` commands, by providing a more modern, efficient, and versatile interface. By using Orchestration CLI, you can automate tasks efficiently, reducing manual effort and operational overhead. Orchestration CLI helps you streamline command-line interactions, enhance cross-platform compatibility, and build a more efficient workload automation process. It also simplifies maintenance, lowers costs, and minimizes IT requirements.

Orchestration CLI also provides a more modern and user-friendly interface, and is designed to be intuitive and efficient, making it easier for administrators and users to complete tasks. It combines modernity, compatibility, and enhanced functionality, when compared to the `conman` and `composer` command line.

In addition to the `plan` commands, already available in previous releases, `model` commands are now available in the Orchestration CLI. You can use these commands to create or modify the object definitions for folders, jobs, job streams, and other scheduling or security items in replacement of the corresponding `composer` commands.

See the list in the `Model commands` section in *User's Guide and Reference*.

For more information, see the topic about Orchestration CLI commands in *User's Guide and Reference*.

Easier configuration for Orchestration CLI

You can now easily change the configuration of your Orchestration CLI by switching between multiple configuration files. Use configuration files to establish one or more environments to run IBM® Workload Scheduler. Setting up different environments, with specific configuration settings, helps you tailor your environment, efficiently manage workloads, and utilize resources effectively.

You can create multiple contexts with different configuration settings using Orchestration CLI. This feature allows you to manage various workflow scenarios by changing the context. You can use context commands to manage or switch between multiple contexts and to perform actions in a specific context. The default context is the one that you create initially until you

change it using the commands. Unless you explicitly specify the context in the command, all commands are run using the default context.

See the list in the Context commands section in *User's Guide and Reference*.

For more information, see the topic about Orchestration CLI commands in *User's Guide and Reference*.

Automation as code: object definitions in JSON and YAML formats

In addition to the traditional scheduling language used in previous product versions, you can now create, display, and modify object definitions in JSON and YAML formats.

JSON and YAML formats guarantee a large number of advantages; for example, YAML is designed for readability, is intuitive, and easy to write. JSON, on the other hand, is designed to simplify parsing and is natively compatible with JavaScript. Both formats are very useful in DevOps and make manipulating object definitions easier and more immediate. This also enhances flexibility, performance optimization, compliance with standards, and integration ease.

The flexibility to support diverse serialization formats provides you with a range of choices. You can select the most appropriate format for specific use cases, aligning with the requirements of various services and platforms. You can select the one that best fits the requirements for a particular scenario. Some systems or services might use a specific serialization format as a standard. By supporting multiple formats, IBM® Workload Scheduler can integrate easily with different services and platforms. In some scenarios, older systems or external services might still rely on older serialization formats. The support for multiple formats ensures compatibility with diverse platforms and reduces friction in the integration process. For more information, see the topic about managing multiple formats in *User's Guide and Reference*.

Observability with OpenTelemetry

Gain a deeper insight into your environment with OpenTelemetry

Observability has become a key concept in computer science, particularly in the fields of systems and network management, and more recently, in cloud-native applications. It provides a large number of advantages to developers and system administrators, such as:

Improved System Understanding

Observability helps developers and system administrators gain a deep understanding of their environment. By observing the environment state and its changes over time, they can understand its behavior, identify patterns, and make informed decisions.

Enhanced Debugging

Debugging becomes more efficient with observability. It allows for the identification of anomalies or bugs in the environment. By observing the full environment over time, system administrators can pinpoint where the issue originated, what caused it, and how it affected the system.

Proactive Problem Solving

Observability supports proactive problem solving. Instead of waiting for a component to fail, observability allows issues to be detected and addressed before they cause a failure. This leads to increased reliability and availability.

Informed Decision Making

Observability provides the data necessary for informed decision-making. It provides insights into how changes in the environment affect its performance and behavior. This information is crucial when making decisions about upgrading, scaling, and allocating resources.

Efficient Resource Utilization

Observability can lead to more efficient resource utilization. By understanding the behavior of the environment, administrators can optimize resource allocation, reducing costs and improving system performance at the same time.

You can now use observability to gain insights on the performance and behavior of your IBM® Workload Scheduler environment for monitoring and analysis, and improve overall system reliability and user experience. For example, you can use metrics to analyze **traces**, metrics, and logs and filter the information you are capturing, thus reducing unnecessary noise and excess costs. Additionally, filtering makes it easier to also add custom tags to metrics for streamlined organization and searching.

For more information, see the topic about OpenTelemetry in *User's Guide and Reference*.

CyberArk and Kerberos integrations automatically installed on dynamic agents

CyberArk and Kerberos integrations are now available by default on each dynamic agent. The related documentation is available on [Automation Hub](#), together with the previous versions of the integrations.

This enhancement simplifies installation and configuration for these important integrations which strengthen security, compliance, and efficiency.

Serviceability enhancements

A number of serviceability enhancements have been inserted in the current release. The main items are as follows:

UNIX Support for populated *DATA_DIR*

On UNIX operating system, you can now specify at installation time a *TWA_DATA_DIR* and *DWC_DATA_dir* folder which is already populated. In previous releases, only an empty *DATA_DIR* was supported.

Improvements in the *wa_pull_info* script

The *wa_pull_info* script replaces the back-level *tw_inst_pull_info* script, which is no longer supported. The *wa_pull_info* script gathers information about your environment and can be used when requested by Software Support to diagnose and troubleshoot a problem. This script now features a number of usability improvements, such as:

- improved error handling
- higher number of information gathered
- improved video output, displaying only the most important steps and errors found, while the full output is stored in a file.

For more information about the script, see the topic about the `wa_pull_info` in *User's Guide and Reference*.

Changed features and feature capabilities in version 10.2.2

A list of features changed since the previous release

The following features have changed or have been introduced or removed since the previous release:

FIPS compliance temporarily not supported

FIPS compliance is not supported in the current product version. This is because OpenSSL 3.0 libraries do not provide a FIPS-compliant validation algorithm for P12 certificates. Development is in progress with the aim of supporting FIPS in upcoming releases.

A new optional parameter named **enablefips** is available in the `serverinst` and `twinsinst` scripts to check FIPS settings before you install or upgrade. The default value in a fresh installation is `false`.

If you are upgrading from an environment where FIPS compliance is supported, you should disable FIPS configuration. When upgrading, a check is performed on your settings to ensure FIPS is not enabled. If it is enabled, the upgrade stops and an error message is displayed. To resume the upgrade, set the **enablefips** parameter to `false` when running the `serverinst` and `twinsinst` scripts. The upgrade proceeds and the configuration files are updated accordingly. For more information, see the FAQ about FIPS compliance in section FAQ - Upgrade procedures in *Planning and Installation Guide*.

SSL FIPS enabled option removed

The **SSL FIPS enabled** option is no longer present in the `localopts` and `ita.ini` files and has been replaced with the new **SSL FIPS compliance** option.

New SSL configuration options available in `localopts` file

Two new options (**ssl_config_file** and **ssl_cipher_suites**) are now available for configuring SSL on dynamic agents and fault-tolerant agents. Use the new options to specify a configuration file for OpenSSL and a list of encryption algorithms for TLS v 1.3. For more information, see the section about setting local options in *Administration Guide*.

Changes in supported operating systems - version 10.2.2

Support for operating systems has been streamlined and rationalized. HP-UX, Sun Solaris, and Linux PPC are no longer supported. For a complete list of supported operating systems, see [Supported operating systems](#).

Changes in supported operating systems - version 10.2.3

In upcoming version 10.2.3, master domain manager and Dynamic Workload Console will no longer be supported on AIX operating systems.

twinstpullinfo command removed

The `twinstpullinfo` command was supported only on HP-UX and Solaris operating systems, which are no longer supported starting from version 10.2.2. The command has been replaced by the `wa_pullinfo` command on supported operating systems. For more information, see [Serviceability enhancements on page 54](#).

Satisfying Requests for Enhancements (RFEs) in version 10.2.2

IBM Workload Scheduler satisfies Requests for Enhancements (RFEs).

Requests for Enhancements (RFEs) give customers the opportunity to collaborate directly with the product development team and other users. The team prioritizes and develops new product features based on proposals made by customers.

IBM® Workload Scheduler V 10.2.2 delivers the following RFEs:

WAO-I-487 - Certificate rotation parameters now available in `values.yaml` file in cloud environments

You can now rotate certificates by setting the following new keys in the `values.yaml` file:

- **duration**
- **renewBefore**

This is useful for testing purposes, especially if the master domain manager and agents are in different clusters, and also in production environments, if a different rotation policy exists.

WAO-I-5 - Improvements in job management

With the recent enhancements in job management, when adding a job definition along with its description to a job stream, the description of the job definition remains visible.

To view a complete list of RFEs, new, planned, and delivered, see: [RFE online community](#).

IBM Workload Scheduler version 10.2.1 enhancements

Learn about the IBM Workload Scheduler version 10.2.1 enhancements.

IBM Workload Scheduler version 10.2.1 includes the following enhancements:

- [Enhanced security for default certificates on page 56](#)
- [Certificates downloadable to fault-tolerant agents at installation time or later on page 58](#)
- [New OCLI plan commands supported on page 58](#)
- [Support for OpenSSL 3.0.x libraries from UNIX operating systems on page 58](#)
- [New actions available in the Orchestration Monitor on page 59](#)
- [Customizing the login page on page 59](#)
- [Changed features and feature capabilities in version 10.2.1 on page 59](#)

Enhanced security for default certificates

Enhanced security for default certificates.

Certificates, either default or custom, are now required when installing IBM® Workload Scheduler. You can no longer install IBM® Workload Scheduler without securing your environment with certificates.

Default certificates are generated automatically when you install the master domain manager. To generate default certificates, define the password for the certificates using the **sslpassword** parameter when you run the `serverinst` script. The certificates are generated using a strong 4k encryption key and the password you specified. These certificates feature a long-term expiration date and are unique for each environment you install. Because the certificates are unique for each environment, if you install a new master domain manager and want it to communicate with an existing Dynamic Workload Console, you need to import the certificates from the master domain manager to the Dynamic Workload Console. For more information, see the FAQ about connecting a new or upgraded master domain manager to an existing Dynamic Workload Console.

The certificates are then stored on the master domain manager in the `installation_directory/defaultCerts` and `TWA_DATA_DIR/ssl/depot` directories. You can use the `TWA_DATA_DIR/ssl/depot` folder to retrieve the default certificates for the other product components.

Before you install the other server components (backup master domain manager, domain manager, backup domain manager, dynamic domain manager, backup dynamic domain manager) or the Dynamic Workload Console with the installation scripts, copy the certificates from the master domain manager to the workstation where you plan to install the component.

When you run the installation script, specify the **sslkeyfolder** and **sslpassword** parameters. These parameters indicate the path on the local workstation where the certificates are stored and the password you defined for the certificates when installing the master domain manager.

When you install a dynamic agent or fault-tolerant agent using the `twinst` script, specify the `wouser` and `wapassword` parameters. The agent uses these parameters to log in to the master domain manager and download the default certificates from the `TWA_DATA_DIR/ssl/depot` directory.

If you are upgrading all your components from earlier versions where you used default certificates, the certificates are updated automatically. Before you upgrade your environment, create an environment variable named `JKS_SSL_PASSWORD` on each server component (with the exception of the master domain manager) and on each Dynamic Workload Console workstation. In the `JKS_SSL_PASSWORD` environment variable store the password for the default certificates. For both fault-tolerant agents and dynamic agents, the updated certificates are downloaded automatically from the master domain manager without user intervention. You can find useful information about upgrading in the topic about FAQ on upgrade procedures in IBM® Workload Scheduler: Planning and Installation.

If you plan to upgrade a part of your environment, for example you plan to connect an agent at 10.2.x version with a master domain manager at an earlier version, convert the certificates to the new `.PEM` format before the upgrade, as described in the topic about upgrading in a mixed-version environment when using default certificates.

For more information, see the upgrade procedures in the section about upgrading from the CLI in *IBM Workload Scheduler: Planning and Installation*.

For more information about connection security, see [Connection security overview](#)the connection security overview in *Administration Guide*.

For more information about using default certificates, see [SSL connection by using the default certificates](#)the section about SSL connection by using the default certificates in *Administration Guide*.

For more information about all the installation commands, see the reference section about the installation commands in *IBM Workload Scheduler: Planning and Installation*.

Certificates downloadable to fault-tolerant agents at installation time or later

Download certificates to fault-tolerant agents using the `twinsinst` or `AgentCertificateDownloader` scripts.

You can now download certificates from the master domain manager to fault-tolerant agents, either at installation time with the `twinsinst` script, or at a later time with the `AgentCertificateDownloader` script. For more information, see the Reference section in *IBM Workload Scheduler: Planning and Installation*

New OCLI plan commands supported

New OCLI plan commands supported.

The following new plan commands are now available in the OCLI:

- `fence`. For more information, see the section about Orchestration CLI commands in *User's Guide and Reference*.
- `limit cpu`. For more information, see the section about Orchestration CLI commands in *User's Guide and Reference*.
- `limit sched`. For more information, see the section about Orchestration CLI commands in *User's Guide and Reference*.
- `altrpri`. For more information, see the section about Orchestration CLI commands in *User's Guide and Reference*.

Support for OpenSSL 3.0.x libraries from UNIX operating systems

If you install the master domain manager on recent UNIX operating systems, you can use the OpenSSL 3.0.x libraries provided with the operating system. The list of UNIX operating systems whose libraries you can use is as follows:

- Ubuntu 22
- AIX 7.3
- Red Hat 9

This is useful in case a vulnerability is discovered in a library: you can upgrade to a newer version and IBM® Workload Scheduler can work seamlessly with the updated library.

To ensure IBM® Workload Scheduler uses these libraries, always launch the installation or upgrade procedure from a brand new shell. You can also check the OpenSSL library currently in use with the `which openssl` command and check the OpenSSL version with the `openssl version` command.

If your operating system does not provide OpenSSL 3.0.x libraries, you can use the libraries packaged and installed with IBM® Workload Scheduler.

New actions available in the Orchestration Monitor

Brand-new actions are available in the Orchestration Monitor of the Dynamic Workload Console.

You can now quickly change the priority for jobs and job streams, modify the limit for job streams and workstations, and update the fence value for workstations directly from the Orchestration Monitor. In just few clicks, you can have a granular control on your environment.

Customizing the login page

You can now customize the login page of the Dynamic Workload Console and Self-Service Catalog with personalized labels and logo. For more information, see the topic about personalizing the login page in *Administration Guide*.

Changed features and feature capabilities in version 10.2.1

A list of features changed since the previous release

The following features have changed or have been introduced or removed since the previous release:

Improved performance with asynchronous calls

Asynchronous calls are now supported. This improves performance, responsiveness, and scalability because you can proceed with other tasks, while one or more calls are processed in the background.

Global Security Kit (GSKit) code removed from the product

GSKit is a set of libraries and utilities provided by IBM for SSL or TLS communication. Removal of GSKit does not impair communication in your environment, also in case some of your agents are still at a back-level version.

Communication between components is ensured by the TLS protocol, which leverages OpenSSL 3.0 and supports also agents at back-level versions. GSKit-related options have been removed from the `localopts` file and replaced with the new options listed below, which you can use to define the SSL version and level:

- **SSL Version**
- **SSL Ciphers**
- **CLI SSL Ciphers**
- **CLI SSL Version**

For more information, see the section about setting local options in *Administration Guide*.

Certificates can now be downloaded to fault-tolerant agents using the `twinst` and `AgentCertificateDownloader` scripts

You can now download certificates from the master domain manager to fault-tolerant agents, either at installation time with the `twinst` script, or at a later time with the `AgentCertificateDownloader` script. For more information, see the Reference section in *IBM Workload Scheduler: Planning and Installation*.

FIPS compliance no longer supported

FIPS compliance is no longer supported in the current product version. Development is in progress with the aim of supporting it again in upcoming releases. For this reason, the **ssl fips enabled** keyword in the `localopts` file is set to `no` by default. If you set it to `yes`, the setting is ignored and the value `no` is used instead.

IBM Workload Scheduler version 10.2.0 enhancements

Learn about the IBM Workload Scheduler version 10.2.0 enhancements.

IBM Workload Scheduler version 10.2.0 includes the following enhancements:

- [Enhanced monitoring on page 60](#)
- [Self-Service Catalog Approval Flow on page 61](#)
- [New enhanced integrations on page 61](#) [Enhanced AI Data Advisor on page 61](#) [Changed features and feature capabilities in version 10.2 on page 62](#)

Enhanced monitoring

Introducing the new and improved monitoring features.

Discover an enhanced monitoring experience that revolves around a single point of access to control your whole workload.

Tree view

Monitor your workload efficiently with the intuitive tree view navigation system. The tree view keeps the current query and the folder you are exploring synchronized, enabling you to explore the whole of your automated workload going from folder to folder effortlessly.

Saved queries

With just a few clicks, you can now save, edit, and share your queries, transforming the task of monitoring your business processes into a simple and efficient activity. You can also save queries generated while navigating the tree view, making designing your monitoring workflow a seamless experience.

Advanced query

The reworked advanced query tool gives you granular control over your workload, enabling you to find any item in no time. Advanced queries can also be saved, edited and shared.

Dependencies

You can now visualize all the dependencies of a job or job stream at a glance in a new dedicated page. Dependencies are displayed as information cards that you can interact with.

Look and feel

The new graphic interface is easier to read, easier to use and easier to manage. Processes are more streamlined and intuitive, and every activity flows into each other smoothly to guarantee a pleasant monitoring experience.

Self-Service Catalog Approval Flow

Introducing the Self-Service Catalog Approval Flow.

Make the services of your Self-Service Catalog more interactive with the new Approval Flow feature.

With the Approval Flow feature, you can decide the outcome of a request using a human-in-the-loop approach, leveraging the integration with Jira or ServiceNow.

You may need an approval for a specific service on the Self-Service Catalog. For example, a manager may need to approve a new license request for a software to check if such request fits in the budget. With the Approval Flow feature, you can specify the need for an approval in the definition of a specific service.

The Self-Service Catalog user can monitor the status of the service, which changes according to the status of the ticket on Jira or ServiceNow.

Thanks to this new feature, requests that would require many different steps to be approved can be processed in just few clicks!

New enhanced integrations

Empower your workflows with the new enhanced integrations.

IBM® Workload Scheduler V10.2 includes the following enhanced integrations:

- Jira
- ServiceNow
- RESTful
- JSONata
- IBM Robotic Process Automation (IBM RPA) integration

Thanks to Jira and ServiceNow integrations, you can open tickets or service requests on one of these two applications, monitor tickets and requests, or open a ticket and monitor it.

You can also use these integrations to enable human interaction to approve or reject requests or ask for more information.

Thanks to RESTful integration, you can quickly integrate any application that uses REST APIs, call an API and track it directly from the job avoiding the creation of several jobs.

Thanks to JSONata integration, you can easily manipulate data and pass the output properties from one job to another.

Thanks to IBM RPA integration, you can automate UI-based tasks by converting human-based actions in scripts and schedule them with IBM® Workload Scheduler.

Enhanced AI Data Advisor

Introducing AI Data Advisor (AIDA) enhancements.

New Machine Learning engine

Neural Prophet has been adopted as AIDA Machine Learning engine. It enables:

- More accurate predictions.
- Multivariate analysis of data to highlight correlation among KPIs.
- Improved time-to-value due to a faster training with less resource usage.
- Easy of extension to other Machine Learning models.

Metrics aggregation

Dynamic data compression depending on data variance. It enables a huge improvement in storage usage.

Customizable alert detection

Alert detection can be customized to better suit customer needs.

- In the `common.env` configuration file (or in the `value.yaml` file for Kubernetes deployments), you can use the `PROPHET_ORCHESTRATOR schedule_alert` parameter to set how often KPIs data must be checked against their prediction to detect anomalies.

Changed features and feature capabilities in version 10.2

A list of features changed since the previous release

The following features have changed or have been introduced or removed since the previous release:

Custom certificates enhancements

Default SSL certificates are no longer provided with the containers of the server, console, and agent components. For more information see [IBM Workload Automation](#)

Enabling installation of dynamic agents on container with a remote gateway

It is now possible to deploy a new dynamic agent and enable the communication directly with another agent gateway. For more information see and [IBM Workload Automation Dynamic Agent](#) and [IBM Workload Automation](#)

IBM Workload Scheduler version 10.1.0 Fix Pack 4 enhancements

Learn about the IBM Workload Scheduler version 10.1 Fix Pack 4 enhancements.

IBM Workload Scheduler version 10.1 Fix Pack 4 includes the following enhancements:

- [Changed features and feature capabilities in version 10.1 Fix Pack 4.](#)

IBM Workload Scheduler version 10.1.0 Fix Pack 3 enhancements

Learn about the IBM Workload Scheduler version 10.1 Fix Pack 3 enhancements.

IBM Workload Scheduler version 10.1 Fix Pack 3 includes the following enhancements:

- [Exporting variable tables with multiple lines on page 63](#)
- [Improved user experience with the new Run Cycle Preview on page 63](#)
- [Use of base folder in Orchestration CLI commands on page 63](#)
- Improved product quality, thanks to the resolution of many APARs and defects. For the complete list of APARs and defects, see: [Fix Pack readmes](#).

Exporting variable tables with multiple lines

You can now export variable tables with multiple lines.

You can now export a variable with multiple lines in a job stream.

This can be achieved in two ways:

- Add `\n` into the property value wrapping the value with quotes (works on Windows and Unix platforms).

```
For example:
jobprop VAR1 "line_1\nline_2\nline_3" -> is displayed as
line_1
line_2
line_3
jobprop VAR2 "\"line_1\nline_2\nline_3\"" -> is displayed as
"line_1
line_2
line_3
"
```

- Write the property value starting with quotes, and then press enter when a new line is needed (works only on UNIX platforms).

```
jobprop VAR2 "line_1
line_2
line_3"
-> is displayed as
line_1
line_2
line_3
jobprop VAR2 "\"line_1
line_2
line_3\""
```

Use of base folder in Orchestration CLI commands

Use of base folder in Orchestration CLI commands.

In the configuration of Orchestration CLI, the current folder field has been enhanced with an option to use base folder that helps the user to fetch the result of files specified in that field. For more information refer [Configuring Orchestration CLI](#)

Improved user experience with the new Run Cycle Preview

Improved user experience with the new Run Cycle Preview

The Run Cycle Preview is now available for job streams and run cycle groups, showing you the instances generated on a full year basis. With a renewed design, it provides you with accessible calendar and colors and a higher level of detail and control.

Watch a customer scenario that shows how easy is to use the Run Cycle Preview in this [video](#), on the [Workload Automation YouTube](#) channel.

IBM Workload Scheduler version 10.1.0 Fix Pack 2 enhancements

Learn about the IBM Workload Scheduler version 10.1 Fix Pack 2 enhancements.

IBM Workload Scheduler version 10.1 Fix Pack 2 improves the product quality thanks to the resolution of many APARs and defects. For the complete list of APARs and defects, see: [Fix Pack readmes](#).

IBM Workload Scheduler version 10.1.0 Fix Pack 1 enhancements

Learn about the IBM Workload Scheduler version 10.1 Fix Pack 1 enhancements.

IBM Workload Scheduler version 10.1 Fix Pack 1 includes the following enhancements:

- [Self-Service Catalog enhancements on page 64](#)
- [Orchestration CLI on page 65](#)
- [Introducing REST API V2 on page 65](#)
- [Orchestration Query Language \(OQL\) on page 65](#)
- [Enhancing authentication using API Keys on page 66](#)
- [Enhancing agent authentication using JSON Web Tokens on page 67](#)
- [Advanced password encryption on page 67](#)
- [Installing the agent on IBM i with a user different from QSECOFR on page 68](#)
- [Changed features and feature capabilities in version 10.1 Fix Pack 1 on page 68](#)
- [Satisfying Requests for Enhancements \(RFEs\) in version 10.1 Fix Pack 1 on page 69](#)

Self-Service Catalog enhancements

Self-Service Catalog: a business-oriented interface to submit on-demand business flow.

A new and improved version of the Self-Service Catalog is available. You can now launch your services quickly and easily and check on them at anytime by accessing the Self-Service Catalog from any device.

To use the Self-Service Catalog you do not need to be a Workload Automation expert, but you can leverage on services based on automation capabilities in no time, provided you are connected to the Dynamic Workload Console in Single Sign-On (SSO). For more information, see [Configuring the Dynamic Workload Console for Single Sign-On](#).

The IBM® Workload Scheduler scheduler or application designer can now define services directly from the Workload Designer marking the job streams as services and specifying service parameters. As part of the job stream definition, the service definitions now can be easily transferred to a different environment

When creating and editing SSC-ready job streams, it is recommended you use the Dynamic Workload Console.

For more information about defining SSC-ready job streams, see the online help for job stream definitions in the Dynamic Workload Console and section about job stream definition in *User's Guide and Reference* for editing the job stream from the command line.



Note: Upgrading to the latest version, the catalogs created in the previous Self-Service Catalog are not visible in the new interface. In the new Self-Service Catalog interface, you can see the catalogs created with the new Workload Designer from the V10.1 Fix Pack 1 or later. Use the previous interface to see the catalogs created with the previous Self-Service Catalog.

Orchestration CLI

The new command line interface to run jobs or job streams and to interact with IBM Workload Automation server.

The installation package for IBM Workload Automation contains the installation files for Orchestration CLI. After downloading the package, you can connect Orchestration CLI to any of the IBM Workload Automation servers by configuring the `config.yaml` file. Orchestration CLI is a stand-alone application that can be used on any desktop or server, or on dynamic agents and replaces `conman`.

You can use the same Orchestration CLI with multiple servers of IBM Workload Automation. Orchestration CLI is also compatible with later versions of IBM Workload Automation and no need to update each time when the server is updated. For more information, see Orchestration CLI.

Introducing REST API V2

IBM® Workload Scheduler REST API V2 to better integrate workload scheduling capabilities with external products and solutions.

A new version of REST APIs has been introduced to operate on the product from both User Interface and Command Line Interface.

The new features of REST API V2 are designed to make the user experience even smoother:

- Enhanced filtering opportunities: according to how specific your query needs to be, you can decide whether to use `planFilter`, which is similar to `conman` syntax and offers the same filtering capabilities, or OQL syntax, which is simpler and allows ordering the results. You can also decide to use both of them, using `planFilter` for filtering and OQL for ordering.
- Improved payloads for easier consumption: both the hierarchy and the structure of payloads have been revamped to improve their understanding and usage.
- Introduction of efficient multi-item endpoints: each action can be performed by ID and by filter. In this way, you can decide whether to operate on a single item or on multiple items.

Orchestration Query Language (OQL)

Orchestration Query Language: querying has never been so easy.

The Orchestration Query Language (OQL) is a new syntax that applies to REST API V2 and helps you monitoring your IBM® Workload Scheduler production plan environment. .

Creating queries and retrieving items in your database is now quicker and easier than before thanks to the different intuitive OQL keywords at your service.

For more information, see the section about Monitoring your environment in the Dynamic Workload Console User's Guide.

Enhancing authentication using API Keys

Use API Keys to authenticate a command line or application easily and quickly.

You can create both Personal and Service API Keys in the Dynamic Workload Console and easily assign them to either specific users or groups. A comprehensive API Keys monitoring tool gives you full control over every valid, expiring and expired API Key that have been associated with an engine. For more information, see the topic about authenticating the command-line client using API Keys in Dynamic Workload Console User's Guide

You can use API Keys to authenticate the command line. You can use an API Key to get authenticated when you launch composer, conman, wappman, and ocli commands, instead of having to provide username and password as in previous versions.

To use the API Key with these commands, you need to have a specific set of authorizations defined in the security file, so that you can generate and retrieve the key from the Dynamic Workload Console. To find out the required authorizations, see the topic about Object type - file in the Classic security model section of the Administration Guide.

To generate the key from the Dynamic Workload Console, perform the steps listed in the section about accessing IBM® Workload Scheduler from the command-line client using API Keys in *Dynamic Workload Console User's Guide*.

After generating the token, you can either specify it in the command line with the **-jwt** parameter, or add it in the `useropts` file.

For more information about adding JWT in the `useropts` file, see the topic about setting user options in *Administration Guide*.

For more information about using JWT with commands, see the topics about running conman and composer and about the wappman command in *User's Guide and Reference*.

You can also use the API Key to authenticate the master domain manager when installing the agents. This authentication allows the product to download the JWT or the certificates to be used for secure communication between master domain manager and dynamic agents. If you provide the API Key (with the **apikey** parameter), you no longer need to specify username and password (**wouser** and **wapassword** parameters) as in previous versions.

For more information about using the API Key for authentication purposes, see the topics about the `twsinst` and `AgentCertificateDownloader` scripts and about example installation commands in *IBM Workload Scheduler: Planning and Installation*.

Ensure there are no misalignments in date and time in your network nor significant network delays because this might impact JWT performance.

Enhancing agent authentication using JSON Web Tokens

Use JSON Web Tokens to enhance your agent authentication standard.

A JSON Web Token (JWT) is a standardized, self-contained access token which makes it possible for two parties to securely exchange data. Authentication information, expiry time information, and other user-defined claims are digitally signed, so that no database queries are required and the session does not need to be stored on a server.

JWT is especially suited for authentication purposes. Its short messages can be encrypted and securely convey who the sender is and whether they have the necessary access rights. It is also very useful in REST applications, because it ensures stateless protocols, since the information for the authentication is sent with the request.

JWT ensures mutual authentication between master domain manager and dynamic agents. Using JWT is easier and more immediate than downloading and maintaining certificates and, in a containerized environment, you no longer need to configure the ingress controller for SSL passthrough. For more information about JWT on containers, see the **Ingress controller** section in [IBM Workload Automation Server](#).

To download the JWT on your dynamic agents at installation time, use the **jwt** parameter as explained in the section about the `twsinst` script in *IBM Workload Scheduler: Planning and Installation*. You can also download the JWT at a later time as explained in the section about `AgentCertificateDownloader` script in *IBM Workload Scheduler: Planning and Installation*.

You can find some installation examples in the section about example installation commands in *IBM Workload Scheduler: Planning and Installation*

You can also revoke a JWT simply by deleting the workstation definition where the JWT is installed. For more information about deleting a scheduling object from the command line and Dynamic Workload Console, see the section about revoking and reissuing a JSON Web Token in *User's Guide and Reference*.

Ensure there are no misalignments in date and time in your network nor significant network delays because this might prevent JWT from working.

Advanced password encryption

New encryption method for increased security

Security is a major concern in today's interconnected world. Government organizations, financial institutions, healthcare providers, and insurance companies are just a few examples of the types of entities who are taking security seriously.

You can optionally encrypt the passwords that you will use while installing, upgrading, and managing IBM® Workload Scheduler. The `secure` command uses the AES method and prints the encrypted password to the screen or saves it to a file.

For more information, see the section about encrypting passwords and the `secure` script in *IBM Workload Scheduler: Planning and Installation*.

Installing the agent on IBM i with a user different from QSECOFR

You can now install dynamic agents on IBM i with a user different from QSECOFR.

You can now use a user different from QSECOFR to install IBM i agents. In this case, the new **allObjAuth** parameter is required when running the `twinst` command to indicate that the user has the required ALLOBJ authority. Ensure the user is existing and has ALLOBJ authorization. The agent is started after connecting to the system with the TWSUSER or the user defined at installation time.

When you upgrade or uninstall the agent, a check is performed to ensure you are using the same user that performed the installation. If you used **allObjAuth** parameter at installation time, specify it again when upgrading or uninstalling the agent.

The name of the user used to perform the installation is maintained in the `TWA_DATA_DIR/installation/instInfo/instUser` file.

For more information, see the sections about installing, upgrading, and uninstalling agents on IBM i systems and agent installation parameters on IBM i systems in *Planning and Installation Guide*.

Changed features and feature capabilities in version 10.1 Fix Pack 1

A list of features changed since the previous release

The following features have changed or have been introduced or removed since the previous release:

REST APIs documentation automatically integrated into the product documentation

Information about REST APIs is automatically extracted from the code and integrated into the documentation at [IBM Workload Scheduler API](#) for easier reading and increased retrievability.

File Proxies now selected randomly in the Managed File Transfer integration

By default, the integration uses the File Proxy specified by the dynamic domain manager. You can also optionally define a list of preferred File Proxies to be used instead of the default list.

The File Proxy is selected randomly in the list provided by the dynamic domain manager. If you provided a custom list, the random selection is performed only on your list.

The random selection mechanism ensures higher tolerance for peaks in workload and workstation failures. For more information, see [File Transfer Integration](#) on Automation Hub.

Default storing mechanism of job log changed

When you perform a fresh installation, by default the job log is stored in memory instead of being stored in a file, as in the previous releases. This ensures compliance with the PCI standard. For more information, see the section about log files and archived files in *Administration Guide*.

Enhanced security with the `conn_verify` property

You can now optionally set the **conn_verify** property to `1` or `2`, so that the agent must provide valid certificates in order to establish an SSL connection. The default value is `0`, which means that no verification is performed on certificates.

Hostname verification for File Proxy

You can now optionally enable a check on the hostname defined for the certificates used by File Proxy. To enable the verification, create the Java system property named `wa.fileproxy.enforceSSLhostnameVerify` and set it to `true`.

Hostname verification for the mail sender plug-in

You can now optionally define the **MAILPLUGIN_ENABLE_HOST_VERIFICATION** property in the `TWSConfig.properties` file to enable the verification of the hostname of the SMTP server. If you set the property to **yes**, hostname verification is enabled, if you set it to **no**, no verification is performed. The default value for fresh installations is **yes**.

Higher security for job integrations

The **Verify host name** check box is selected by default in some integrations. Starting from version 10.1 Fix Pack 1, to increase the security standard provided by the product, this selection starts a check on the data contained in the keystore file. Ensure you provide the correct path and credentials to the keystore file in the **Certificates** section for the following integrations:

- JSR 352 Java Batch
- Restful
- Job Management
- Job Stream Submission
- Variable Table

If the information in the **Credentials** section is not correct, this causes the job to fail with an error similar to the following:

```
AWKJMJ041E An error occurred during the Job submission.
The error was "AWKSPI006E Keystore authentication failed. The error message is
"No certificates provided, cannot verify the hostname."
```

If your integrations fail after upgrading to Fix Pack 1, you have the following options:

- deselect the **Verify host name** check box.
- provide the correct path and credentials to the keystore file.

Satisfying Requests for Enhancements (RFEs) in version 10.1 Fix Pack 1

IBM Workload Scheduler satisfies Requests for Enhancements (RFEs).

Requests for Enhancements (RFEs) give customers the opportunity to collaborate directly with the product development team and other users. The team prioritizes and develops new product features based on proposals made by customers.

IBM® Workload Scheduler V 10.1 Fix Pack 1 delivers the following RFEs:

RFE 149367 - Installation of dynamic agents on Windows operating systems using Local System Account

You can now install dynamic agents on Windows operating systems without providing a specific user. This means that you no longer need to provide **twuser** credentials to the operator performing the installation.

Moreover, the password no longer expires, ensuring a smooth running of your jobs. For more information, see the section about agent installation parameters in *IBM Workload Scheduler: Planning and Installation*.

RFE 69917, AHA 46 - Interactive jobs supported on dynamic agents installed on Windows operating systems

You can now determine the behavior of native and executable jobs starting interactive programs on dynamic agents installed on Windows operating systems. Interactive programs run only if the job user has an active session open when the job runs. If there is no active session for the job user, the job behavior is defined by the new **RunInteractiveJobOnInvalidSession** property in the `JobManager.ini` file. You can set the property to `true` to enable jobs to start interactive programs even if there is no active session for the job user. Alternatively, set the property to `false` to prevent jobs from starting interactive programs if there is no active session for the job user. This feature was already available for fault-tolerant agents and is now also available by default for dynamic agents. For more information about enabling and disabling this feature, see the section about configuring properties of the native job launcher [NativeJobLauncher] in *Administration Guide*.

For example, consider a job which consists of both interactive instructions (for example, using a text editor), and batch instructions, which run in background when the job is launched. If you set the **RunInteractiveJobOnInvalidSession** property to `true`, batch instructions can run even if the interactive part is not displayed because the user does not have an active session at the moment.

RFE 60060 - dynamic agents to track the workload on the local system

This capability is now available using the Orchestration CLI and REST APIs version 2, as described in [Orchestration CLI on page 65](#) and [Introducing REST API V2 on page 65](#).

To view a complete list of RFEs, new, planned, and delivered, see: [RFE online community](#).

IBM Workload Scheduler version 10.1.0 enhancements

Learn about the IBM Workload Scheduler version 10.1.0 enhancements.

IBM Workload Scheduler version 10.1.0 includes the following enhancements:

- [Managed file transfer on page 70](#)
- [Workload Designer: a new concept to manage your workload on page 71](#)
- [Detecting anomalies in your workload with AI Data Advisor \(AIDA\) on page 72](#)
- [Automatic encryption at rest for key product files on page 72](#)
- [Automatic SSL configuration for fault-tolerant agents on page 73](#)
- [Support for Google Cloud SQL for SQL server on page 73](#)
- [IBM Workload Scheduler events now logged for monitoring tools on page 73](#)
- [Changed features and feature capabilities on page 74](#)
- [Satisfying Requests for Enhancements \(RFEs\) in version 10.2 on page 77](#)

Managed file transfer

The File Transfer integration, already available on [Automation Hub](#), now provides full capabilities for managing and processing all your file transfers using IBM® Workload Scheduler with no need for third-party products. The main enhancements are as follows:

Support for File Proxy

A new component, named File Proxy, is now available by default on each master domain manager. You can use the File Proxy to manage securely files to be transferred with the File Transfer integration available on [Automation Hub](#). All transfer operations are performed by IBM® Workload Scheduler, with no need of third-party products.

You can also optionally install your File Proxy as a stand-alone component on a workstation different from the master domain manager, for example to reduce network traffic and resource usage on the master domain manager. You can configure your stand-alone File Proxy in high availability by defining the URLs of alternate file proxies or a load balancer in the **Broker.fileproxy.urls** property in the `BrokerWorkstation.properties` file. For more information, see the section about the `BrokerWorkstation.properties` file in *Administration Guide*.

Ensure the selected workstation runs a supported version of a Windows or Linux operating system and proceed with the installation and start of the service, as described in the sections about the `fileproxyinst` and `fileproxystart` commands in *IBM Workload Scheduler: Planning and Installation*.

Support for file transfers via the Workstation-to-Workstation internal protocol

You can now easily transfer files to and from agents connected to the same master domain manager. You no longer have to specify the address of the workstations involved in the file transfer, but you can simply define the workstations from the File Transfer integration. Ensure both the master domain manager and agents are at version 10.1. For more information about the integration, see the File Transfer integration on [Automation Hub](#).

Search on remote file systems

You can now perform a search operation on the file system of a remote workstation. On Linux systems, if you do not specify a path, the search is performed on the working directory. On Windows systems, the path is required.

Integrations with Managed File Transfer (MFT) tools as well as Robotic Process Automation (RPA) tools

More integrations with other MFT tools as well as RPA tools to support orchestration of all data transfer using IBM® Workload Scheduler.

Workload Designer: a new concept to manage your workload

Workload Designer: a new concept to manage your workload.

An innovative infrastructure and design have been thought to simplify the user experience and create a more responsive, fast and fluid user interface.

You do not need to switch among the Dynamic Workload Console pages to complete your business workflow anymore. The new Workload Designer contains everything you need. Automate business-critical processes from a single point of access and control.

Automate fast, automate better.

Detecting anomalies in your workload with AI Data Advisor (AIDA)

Starting from this release, a new component is available in IBM® Workload Scheduler - AI Data Advisor (AIDA) - based on Artificial Intelligence and Machine Learning, for early anomaly detection and analysis.

AIDA enables fast and simplified data-driven decision making, for an intelligent workload management. By analyzing historical data and metrics gathered by IBM® Workload Scheduler and predicting their future patterns, AIDA identifies anomalies in the trend of Key Performance Indicators (such as the number of completed jobs in the current plan, the job duration, the job end-time) and sends alerts immediately to anticipate and prevent problems and delays. Alerts show up on the Workload Dashboard and can be notified via email.

For more information, see *AI Data Advisor (AIDA) User's Guide*.

Furthermore, you can always take your environments under control by adding the AIDA widget to your custom dashboard. For further information, see *Creating a customized dashboard for monitoring*.

For instructions about how to install AIDA, see *Deploying AI Data Advisor*.

Automatic encryption at rest for key product files

Starting from this release, key IBM® Workload Scheduler files are automatically encrypted for all fresh installations using AES-256 or AES-128 cryptography.

Data breaches are becoming more and more common and pervasive in today's business world. Encryption is a key feature when it comes to protect sensitive data, such as the data at rest stored in your Symphony plan or message queues. For this reason, all fresh installations starting from this release automatically encrypt key product files using AES-256 or AES-128 cryptography.

Data at rest means data is not being accessed or used but instead stored on your computer, external hard drive, cloud storage, server, or database. Encryption at rest ensures that this data is protected and encrypted.

If you want IBM® Workload Scheduler to encrypt files such as the Symphony file, messages queues, and the `useropts` file at runtime, you do not need to take any actions. By default, the product is automatically encrypted without your intervention. You can also define the folder containing the certificates and the certificates password using the **sslpassword** and **sslkeyfolder** parameters when installing the master domain manager and agents, both fault-tolerant agents and dynamic agents.

You can optionally avoid the encryption at installation time, by setting the **useEncryption** parameter in the `serverinst` and `twsinst` commands to `false`.

For more information, see the section about automatic encryption in *Administration Guide* and the sections about master components installation - `serverinst` script and agent installation parameters - `twinst` script in *IBM Workload Scheduler: Planning and Installation*.

Automatic SSL configuration for fault-tolerant agents

Configure WebSphere Application Server Liberty Base and fault-tolerant agents in SSL mode

You can now configure WebSphere Application Server Liberty Base and fault-tolerant agents in SSL mode and provide your certificates at installation time using the command-line installation and setting the `sslkeyfolder` and `sslpassword` parameters. This ensures encryption in motion for all data moving within your environment.



Note: You are strongly recommended to customize the default certificates straight after the installation process has completed to ensure your environment is secure.

For more information about installation parameters, see the section about master components installation parameters - `serverinst` script and agent installation parameters - `twinst` script in *IBM Workload Scheduler: Planning and Installation*.

IBM® Workload Scheduler also generates certificates in GSKit format. To enable FIPS compliance, set the `ssl fips enabled` option to `yes` in the `localopts` file after the installation has completed and restart the product. As a result, options for GSKit are automatically used by IBM® Workload Scheduler.

For more information about the `localopts` file, see the section about setting local options in *Administration Guide*.

Support for Google Cloud SQL for SQL server

You can now install your on-premises IBM® Workload Scheduler environment also on Google Cloud SQL for SQL server, a fully-managed database service that helps you set up, maintain, manage, and administer your relational databases on Google Cloud Platform.

You only need to install your database, then proceed with the command-line installation for master domain manager and Dynamic Workload Console, as always, specifying MSSQL as the database type.

For the complete command-line installation procedure, see the section about the typical installation scenario in *Planning and Installation Guide*

For more information, see the section about Creating Azure SQL or Google Cloud SQL for SQL server databases for the master domain manager and Dynamic Workload Console in *Planning and Installation Guide*.

IBM® Workload Scheduler events now logged for monitoring tools

Increasing observability for IBM® Workload Scheduler events

You can now monitor the availability of your infrastructure and services, monitor performance, and analyze bottlenecks. IBM® Workload Scheduler features an improved mechanism for monitoring and auditing events. The monitoring engine

generates a .json file containing all WebSphere Application Server Liberty Base events, a file which is consumable by all applications using .json and meeting the OpenMetrics standard, for example AI Data Advisor (AIDA), Splunk and Fluentd.

By further analyzing these values through a data analytics tool, such as AI Data Advisor (AIDA), you detect anomalies and anticipate failure or degradations. For more information about AIDA and how to use it, see *AI Data Advisor (AIDA) User's Guide*.

You can also install Instana and monitor the whole infrastructure, both on-premises and cloud, on which you have installed IBM® Workload Scheduler.

On the cloud side, within the server and agent containers you have a number of sidecar containers which monitor .json log files. On the on-premises side, you can feed your monitoring tool the .json log files and monitor WebSphere Application Server Liberty Base in your environment by creating custom dashboards in the monitoring tool.

This enhancement grants an immediate deep dive into events generated in the IBM® Workload Scheduler environment.

For more information about the metrics IBM® Workload Scheduler exposes, see the section about monitoring IBM® Workload Scheduler in *User's Guide and Reference*.

To enable logging in .json format, add the **JSON** and **JSON_NO_UTF8** keys, depending on the output format you require, in the `BmEvents.conf` file. For more information, see the section about the `BmEvents` configuration file in *User's Guide and Reference*.

You can also find more information in the video [Get your focus right with observability!](#)

Changed features and feature capabilities

A list of features changed since the previous release

The following features have changed or have been introduced or removed since the previous release:

Scheduling objects defined using REST APIs

Starting from product version 9.5, Fix Pack 2, a back-level version of the composer command line was provided to create definitions of a subset of scheduling objects which did not support REST APIs. The back-level version was named `composer951`. This back-level command line is now no longer available because all scheduling objects support REST APIs.

JAVA APIs deprecated

JAVA APIs are deprecated and have been replaced with REST APIs. For more information about REST APIs, see *Driving IBM Workload Scheduler with REST API*.

Rationalization of the JobManager.ini file

The contents of the `JobManager.ini` have been trimmed and revised to remove some properties which are reserved for internal use. These properties have been moved to a file named `JobManagerInternal.ini`, which is located in the binary directory and is overridden at each product upgrade. This file is not documented because it is not relevant for the final user. For more information about the `JobManager.ini` file, see the section about configuring the agent in *Administration Guide*.

Java software prerequisite changed to version 11

Java is a prerequisite for Java job plug-ins and other Java-based functions. With this release, the supported version has been changed to Version 11. This version of Java is not available on HP-UX and Solaris SPARC operating systems. If you install on these operating systems, use the Java Version 11 provided with the operating systems. This change applies to the `twinsinst` script: before launching this command on HP-UX and Solaris SPARC operating systems, provide the path to Java. For more information, see the section about the `twinsinst` script in *IBM Workload Scheduler: Planning and Installation*.

Ensure higher security and safety of your data with TLS V1.2

IBM® Workload Scheduler no longer supports TLS V1.0 and TLS V1.1. Only TLS V1.2 is supported. This enhancement ensures a higher security level for SSL communication.

Audit file log moved to master domain manager and backup master domain manager

When you choose to keep track of changes to scheduling objects, you can store your audit trails in a file, in the IBM® Workload Scheduler database or in both. While in previous versions, the file was stored on agents, with the current version, the file has been moved to the master domain manager and backup master domain manager. For more information, see the section about enabling and storing audit trails in *Administration Guide*.

enPlanAudit global option is now immediately effective

While in previous versions the `enPlanAudit` global option required to run a `JnextPlan` command to make changes effective, starting from the current version, changes to `enPlanAudit` global option are immediately effective. For more information, see the summary about global options in *Administration Guide*.

Audit file format changed

The format of the file where you can optionally store the auditing information for your environment has been changed from `.csv` to `.json`. For more information, see the section about enabling and storing audit trails in *Administration Guide*.

Job Duration Predictor job deprecated

The Job Duration Predictor job is now deprecated and has been removed. Instead, a new IBM® Workload Scheduler component is available: AI Data Advisor (AIDA), based on Artificial Intelligence and Machine Learning, for early anomaly detection and analysis. For further information about AIDA, see the *AI Data Advisor (AIDA) User's Guide*.

New and improved command for password encryption

A new command is now available for encrypting and decrypting passwords you use with IBM® Workload Scheduler with the `xor`, `3DES`, and `AES` algorithms. For more information, see the section about encrypting passwords in *Planning and Installation Guide*.

Several integrations now available on Automation Hub

The following integrations are no longer available with IBM® Workload Scheduler and can be installed from [Automation Hub](#):

- Amazon EC2
- Apache Spark
- Enterprise Java Beans (EJB)
- Hadoop Distributed File System
- Hadoop Map Reduce
- IBM® BigInsights
- IBM® Cloudant
- IBM® Cognos
- IBM InfoSphere DataStage
- IBM® SoftLayer
- IBM® WebSphere® MQ
- Informatica PowerCenter
- J2EE (JMS)
- Job Duration Predictor. This job has been replaced by AI Data Advisor (AIDA). For further information about AIDA, see the *AI Data Advisor (AIDA) User's Guide*.
- Microsoft Azure
- MQTT
- Oozie
- Oracle E-Business Suite
-

TEC Event Integration Facility (EIF) Event Forwarder, also known as TEC Event Forwarder

- Salesforce
- SAP BusinessObjects BI
- SAP PI Channel
- Web Services

Support for several products and integrations no longer available

The products and integrations listed below are no longer supported and the related documentation has been removed from the library:

- Application Lab
- Application Performance Management (APM)
- Extended agent for Multiple Virtual Storage (MVS)
- IBM Tivoli Monitoring (ITM)
- Option Editor
- OSLC Automation
- OSLC Provisioning
- Provisioning BigInsights

- Remote EJB, also known as EJB APIs. These APIs have been replaced by restful APIs.
- SmartCloud Control Desk (SCCD)
- SOAP WebServices based on WSDL. These services have been replaced by restful APIs.
- Tivoli Business Service Manager (TBSM)
- Tivoli Enterprise Portal (TEP)
- Tivoli Provisioning Manager (TPM)
- Tivoli Storage Manager (TSM)
- Tivoli® System Automation for Multiplatforms (TSAMP)
- Tivoli NetView has been replaced with Simple Network Management Protocol (SNMP).
- workload broker jobs

As a result, `optman` parameters related to OSLC have been removed.

Also, the following **optman** options are no longer supported and are forced to their default value:

enTimezone

Deprecated and forced to `yes`.

enSwfaultTol

Deprecated and forced to `no`.

enRetainNameOnRerunFrom

Deprecated and forced to `no`.

ignoreCals

Deprecated and forced to `yes`.

For more information about **optman** options, see the section about global options in *Administration Guide*.

Satisfying Requests for Enhancements (RFEs) in version 10.2

IBM Workload Scheduler satisfies Requests for Enhancements (RFEs).

Requests for Enhancements (RFEs) give customers the opportunity to collaborate directly with the product development team and other users. The team prioritizes and develops new product features based on proposals made by customers.

IBM® Workload Scheduler V 10.2 delivers the following RFEs:

datecalc utility is now available also for dynamic agents

You can now use the `datecalc` utility to resolve date expressions and return dates in the format you choose also on dynamic agents. For more information, see the section about the `datecalc` utility in *User's Guide and Reference*.

The `init.cfg` file is no longer overwritten during the upgrade phase

The `init.cfg` file is a file available for customization and so it should not be overwritten during the upgrade phase from Version 9.4 Fix Pack x. The file has now been moved to the following paths:

UNIX

`TWA_DATA_DIR`

Windows

`twc_install_dir`

where `twc_install_dir` is the directory path where the agent is installed.

If you perform an upgrade from Version 9.4 Fix Pack *x* to Version 10.1, the file is automatically moved to the paths listed above.

For more information about the `init.cfg` file, see the section about rule operation notes in *User's Guide and Reference*.

Modified By and Last Modified have been reintroduced

With the new Workload Designer, the item information regarding who modified it and the last modify has been reintroduced.

To view a complete list of RFEs, new, planned, and delivered, see: [RFE online community](#).

IBM Z Workload Scheduler Small Programming Enhancements released in April 2026

The following features were provided with IBM Z Workload Scheduler small programming enhancements (SPEs) released for Version 10.2, in March 2026 (APAR PH68599).

New Documentation Portal for IBM® Workload Scheduler and IBM Z Workload Scheduler

Effective April 24, 2026, the official documentation for IBM Workload Scheduler and IBM Z Workload Scheduler has moved to a redesigned, dedicated portal. The new portal introduces several major improvements, including:

- **AI-powered chat** to search, navigate, and process information across the documentation set. Available starting from April, 30, 2026.
- **A new documentation template** designed to improve clarity, readability, and overall usability.
- **An updated look and feel**, offering a more modern, consistent, and intuitive user experience.

Please update your bookmarks for all releases: [IBM Workload Automation](#).

Configuring the communication between a z-centric agent and Z controller through a gateway (RFE ZWS-I-226)

In a simple configuration you configure the z-centric agents to connect directly with the Z controller, while in more complex network topologies you might want to prevent a direct communication (for example, when the agents are behind a firewall and need to communicate through the internet). In this latest case, you can configure one or more z-centric agents to communicate with the controller through a local or remote gateway. This unidirectional, polling-based architecture enables secure deployments in highly restricted or segmented network environments, simplifies firewall configuration, and supports scalable connectivity for multiple agents.

For detailed information, see the section about configuring the communication with the Z controller in *Scheduling End-to-end with z-centric Capabilities*. For a description of the possible gateway configurations, see the topic about the z-centric agent installation examples in *IBM Z Workload Scheduler: Planning and Installation*.

Resolution of mandatory dependencies in the daily plan and current plan

By setting the MPRESOL parameter of the `JTOPTS` statement, you can have mandatory dependencies resolved during the daily plan. Moreover, according to the value that you specify, the resolution process searches for a match by also taking into account the completed archived occurrences both in the current plan and daily plan.

For details, see the topics about the resolution of mandatory dependencies in the daily plan and current plan in *Managing the Workload*.

Intelligent JES2 job selection

With the JES symbol `JOBTOKEN`, you can now send to JES2 a token generated by the Z controller for each submitted operation, to be used by the Intelligent JES2 Job Selection feature.

For detailed information, see the OPCOPTS initialization statement in *Customization and Tuning*.

Registering IBM Z Workload Scheduler on the z/OS system

You can now register IBM Z Workload Scheduler through the System Management Facilities (SMF), to collect information about product usage on the z/OS system in the SMF record type 89.

Trigger WLM for a single target job's job chain (RFE ZWS-I-305)

IBM Z Workload Scheduler uses the WLM function to assign extra resource to critical jobs that are running late by promoting them to a higher performance WLM service class. However, you can keep the critical job from being promoted by setting the NOTDEFPROMOTION parameter in the OPCOPTS statement. For details, see *Customization and Tuning*.

Need to be able to customize TdwcGlobalSettings within the zconnector (RFE ZWS-I-308)

In the `connectionFactory.xml` sample file you can specify the interval of time that the Z connector waits for a response from the z/OS server started task before issuing an error. For details, see the `connectionTimeout` property in the topic about defining a z/OS engine in the Z connector in the *IBM Z Workload Scheduler: Planning and Installation*.

USS Module Traceability Enhancement (RFE ZWS-I-315)

To improve serviceability and diagnostic speed, USS-resident modules (LINKEDITED libraries) now include embedded identification of Version, Date, and PTF or APAR. This enables you to immediately verify the maintenance level following a system ABEND or error.

Add the IP address and port number values to the existing message EQQHT29E (RFE ZWS-I-321)

Message EQQHT29E now also provides you with the IP address and port number of the workstation with which the connection is not established.

Restart/Cleanup functionality with STEPRESCHK(YES) to allow more flexibility at step Restart (RFE ZWS-I-327)

If you specify a DD name in the DDNEVER parameter of the RCLOPTS statement, the related step is *never* re-executable regardless of the value that is set in the STEPRESCHK parameter of the same statement. For detailed information about RCLOPTS, see *Customization and Tuning*.

IBM Z Workload Scheduler Small Programming Enhancements released in December 2025

The following features were provided with IBM Z Workload Scheduler small programming enhancements (SPEs) released for Version 10.2, in December 2025 (APAR PH68409).

Specifying matching criteria for a z/OS shadow job to bind a z/OS remote job

When you define a z/OS shadow job to bind a job instance scheduled on another Z controller, the search for the remote job can be made dependent from any of the following matching criteria:

- Closest preceding
- Same scheduled date
- Within a relative interval
- Within an absolute interval

For detailed information, see the topic about how a z/OS shadow is bound with a remote job instance in *Managing the Workload* and watch this [video](#) on the Workload Automation YouTube channel.

From the Orchestration Monitor you can set the status of jobs and workstations

From the Orchestration Monitor, the Set Status action now provides more administrative control and flexibility by enabling you to modify the status of jobs and workstations.

RFE-I-78 CPBPLIM should not be allowed to exceed 50% and all deps nesting level default should be less than 99

To help you set the percentage of CP buffer pool size in the OPCOPTS CPBPLIM and CPDTLIM parameters, you are now provided with relevant information by message EQQN259I.

For details, see the topic about the OPCOPTS statement in *Customization and Tuning*.

WAO-I-415: A new field for occurrences with the original ADID

With the Workload Automation Programming Language you can list the occurrences that were cloned from the AD database into the current plan, based on their original application ID, by specifying the `TMPLADID` argument in the `LIST CPOC` request. For details, see the topic about the `LIST CPOC` request in the *Workload Automation Programming Language for Z User's Guide and Reference*.

ZWS-I-310 Add the JES job number and/or the JOBID to message EQQ8006I

Message EQQ8006I now also provides you with the job ID.

ZWS-I-312 Input file to update the XML used for dataset triggering should not open the file for output

The XML file used as input for the data set triggering function is now accessible only in read mode.

IBM Z Workload Scheduler Small Programming Enhancements released in July 2025

The following features were provided with IBM Z Workload Scheduler small programming enhancements (SPEs) released for Version 10.2, in July 2025 (APAR PH66793).

The Federator component of the Dynamic Workload Console supports Db2 for z/OS

The supported databases where you replicate (mirror) your current plan and run your queries through the Federator now include Db2 for z/OS.

Workflow details for plug-in jobs

You can now monitor the plug-in jobs from the Orchestration Monitor page of the Dynamic Workload Console. For details, see Orchestration Monitor overview.

Additional details related to successor shadow jobs

When you view the list of successor shadow jobs that are bound to a local job, you are provided with additional information related also to the distributed shadow jobs.

For details about viewing the successor shadow jobs, see the section about listing the shadow jobs bound to a job in the CP in *Managing the Workload*.

Enriched metrics for monitoring and analyzing your workload

Enhanced metrics now provide you with deeper insights on the performance and behavior of your environment for monitoring and analysis, thus improving overall system reliability and user experience.

New MODIFY command to collect statistics about records in the CP

By issuing the MODIFY command `/F procname,CPRECCNT` the message EQQE1711 is logged in the EQQMLOG data set to show the number of occurrences, operations, and workstations currently included in the current plan, along with their status.

Support of IBM® Semeru Runtime Certified Edition V21

The Dynamic Workload Console reporting feature now requires IBM® Semeru Runtime Certified Edition (formerly known as IBM® 64-bit SDK for z/OS®, Java Technology Edition) V17 or V21.

Support of FIPS for distributed components

A new and improved version of the secure command empowers administrators to activate or deactivate FIPS compliance post-installation. Key capabilities of secure include:

FIPS activation

Configures the environment for FIPS 140-3 operation. This includes verifying compatibility (for example with Windows users) and converting any existing passwords stored with the older TripleDES algorithm to the FIPS-approved AES encryption standard. It then activates FIPS flags across essential configuration files and restarts the affected components. The activation process can accommodate transitional phases for algorithms like TripleDES, as permitted, before full conversion.

FIPS deactivation

Reverts the FIPS-specific settings in the aforementioned configuration files and restarts the components, should FIPS mode need to be disabled.

ZWS-I-18 Enhance long duration alert message information

Messages EQQE028I, EQQE038I, and EQQ039I provide you with more detailed information for a more efficient assessment and analysis:

- Job start time
- Job estimated duration
- Current job duration
- Current job status
- Extended job name

ZWS-I-68 WAPL enhancement: List CPSAI

By issuing the LIST CPOPCOM request from Workload Automation Programming Language (WAPL) or the native program interface (PIF) you can now list any available System Automation Information for an operation, enabling you to speed up searching, modifying, and verifying the information itself.

For more details about the arguments added to Workload Automation Programming Language, see the section about LIST CPOPCOM in the *Workload Automation Programming Language for Z User's Guide and Reference*.

For more details about the arguments added to the program interface, see the section about the LIST CPOPCOM arguments in the *Developer's Guide: Driving IBM Z Workload Scheduler*.

IBM Z Workload Scheduler Small Programming Enhancements released in April 2025

The following features were provided with IBM Z Workload Scheduler small programming enhancements (SPEs) released for Version 10.2, in April 2025 (**APAR PH64943**).

Listing the successor shadow jobs bound to a job in your CP

From the ISPF panels you can now list the shadow jobs that are bound to a local job in your current plan, thus enabling you to evaluate the impact of performing the following actions on the local job instance:

- Browsing an occurrence external dependencies
- Browsing an operation dependencies
- Browsing an operation external dependencies
- Deleting an occurrence
- Deleting an operation
- Setting an occurrence to complete
- Rerunning an occurrence
- Setting an occurrence to waiting

For details, see the topic about listing the z/OS shadow jobs in *Managing the Workload*.

A more efficient flow of data between the Z controller and Federator

To enhance the processing of the events sent from the Z controller to the Federator, in the MIRROPTS statement you can now set THREADNUM to define the number of threads used by the DB Filler task and SOCKETNUM to define the number of sockets per thread. For details, see the topic about the MIRROPTS statement in *Customization and Tuning*.

ZWS-I-31 Update sample code in line with product changes

The EQQJVXIT sample exit stored in SEQQSAMP has been updated with the latest DSECT for the current plan record layout.

IBM Z Workload Scheduler version 10.2 enhancements

IBM Z Workload Scheduler version 10.2 includes the following enhancements.

Enhance the performance of the interfaces by mirroring the current plan on a database

Thanks to the Federator component which is automatically installed with the Dynamic Workload Console V10.2.3, you can now replicate the z/OS current plan on a database. This enables you to query your jobs, job streams, and workstations on the configured database through the Federator, without affecting the Z controller, which results in a performance enhancement of both the Dynamic Workload Console and REST APIs V2.

Find detailed information about activating the replication of data to a database in the section about mirroring the current plan to enable the Orchestration Monitor in the *Dynamic Workload Console User's Guide* and watch the [video](#) available in the Workload Automation channel on YouTube.

Have granular control over your workload through the Orchestration Monitor of the Dynamic Workload Console V10.2.3

By registering your z/OS engine to the Federator component of the Dynamic Workload Console V10.2.3, you are enabled to:

- Monitor jobs, job streams, and workstations from the new Orchestration Monitor.
- Perform actions on the monitored objects through the REST APIs V2, with which you can also run OQL queries.

The Event-filtering exit (User Exit 4) has been added to the User Exits Workbench tool

The Event-filtering exit (User Exit 4) has now been added to the list of exits that you can customize with the Workbench tool. For details, see the section about creating user exits through the User Exits Workbench in *Customization and Tuning*.

Elevate your data management with PostgreSQL database

Improve database performance and reliability by implementing the PostgreSQL database in your IBM® Workload Scheduler environment.

You can now use the PostgreSQL database for creating the database for the Dynamic Workload Console.

PostgreSQL offers an extensive list of features related to performance, security, and configuration options. It is free and open-source, and its high customizability supports tailoring it to the specific needs of your environment. It also provides scalability, availability, and functionality suitable for large volumes of data.

PostgreSQL replaces OneDB and Informix, which are no longer supported.

For more information about using PostgreSQL to create the database for the Dynamic Workload Console, see the topic about creating and populating the database for PostgreSQL for the Dynamic Workload Console in *IBM Z Workload Scheduler: Planning and Installation*.

Unleash the power of the cloud using Amazon RDS for MSSQL

Amazon RDS for MSSQL is a powerful, scalable managed database service that makes it easier to set up, operate, and scale SQL Server deployments in the cloud. It provides a powerful and convenient option for managing MSSQL databases in the cloud.

Amazon RDS for MSSQL also simplifies database management, enhances security, and provides the flexibility to scale up or down as your needs change.

For more information about using Amazon RDS for MSSQL to create the database for the Dynamic Workload Console, see the topic about creating and populating the database for MSSQL for the Dynamic Workload Console in *IBM Z Workload Scheduler: Planning and Installation*.

Derby database is no longer supported

Derby database is no longer supported.

If you upgrade from an environment where you are using Derby, you have to install a supported database and extract all your definitions from Derby as described in the topic about connecting the Dynamic Workload Console to a new node or database in *IBM Workload Scheduler: Planning and Installation*.

Deprecated functions

The following functions were delivered and supported until IBM Z Workload Scheduler V10.1, and are now deprecated:

- Open Services for Lifecycle Collaboration (OSLC) integration.
- Fault-tolerant agents
- Direct access storage device (DASD) connection method between the Z controller and trackers.

IBM Z Workload Scheduler Small Programming Enhancements released in December 2023

The following features were provided with IBM Z Workload Scheduler small programming enhancements (SPEs) released for Version 10.1, in December 2023 (APAR PH56381).

Integrating with IBM Z JCL Expert

By integrating with IBM Z JCL Expert, you can now stop wasting time searching for and fixing errors. You can check multiple JCLs simultaneously and automate the checking process.

From the Dynamic Workload Console V10.2.1 you can:

1. Invoke IBM JCL Expert when a job ends in error.
2. Verify the output of the JCL check in a new panel.
3. Speed up the debug process.
4. Automatically validate your JCL coding to ensure compliance with custom rules.

Keep your environment safe with the power of IBM Z JCL Expert! For detailed information, see the section about integrating with IBM Z JCL Expert in *Managing the Workload*.

Creating exits through the User Exits Workbench

The User Exits Workbench tool enables you to create and customize a number of user exits with the IBM z/OS Metal C language, instead of the Assembler language. You can quickly create, edit, compile, and link-edit the user exits from a single ISPF panel, as an alternative to creating them manually. For more details, see the section about creating user exits through the User Exits Workbench in *Customization and Tuning*.

Time-dependent operations hold on trackers

For trackers using JES2 for z/OS V3.1, or later: When a time-dependent operation scheduled on a tracker is ready to be run and you have set the JESSYMBOL HOLDUNTIL parameter in **OPCOPTS**, the operation will be sent to the tracker at the time resulting from the IA time minus the seconds set in JESSYMBOL HOLDUNTIL. On the tracker, JES2 holds the operation until the time dependency is resolved and optimizes the resources required for execution, then the operation is run.

This applies only to jobs that run on computer automatic workstations, except for virtual and started-task workstations.

Increase default MAXREL size for EQQADDS and EQQLTDS (RFE ZWS-I-40)

In EQQPCS01 sample, the maximum record length (MAXRECL) of EQQADDS and EQQLTDS data sets has been increased from 131072 to 400000.

EQQN131W no longer needed in recent releases of IWSz (RFE ZWS-I-159)

Message EQQN131W was introduced with IBM Z Workload Scheduler V8.6 to warn you about a change in the unit of measure for the JTOPTS MAXJSFILE parameter. Because versions earlier than V8.6 are no longer supported, the message is no longer meaningful and has been deleted.

Restore the CP at the moment just before a human error (RFE ZWS-I-178)

If manual actions were performed on some application occurrences by error but the new current plan still exists and is valid, you can have IBM Z Workload Scheduler to take over the new current plan again, until the date and time of the latest valid record. This prevents re-creating the CP including *all* the archived JT events, even those that could be invalid or inconsistent after the detected errors.

For more details about the procedure to follow, see Re-creating the current plan from the new-current-plan and EQQJTARC until a specific date.

Java runtime now installed automatically

Java runtime is now installed automatically when you install an agent. For this reason, the **-addruntime** parameter in the **twinst** script is ignored when set to `false` and Java runtime is always installed by default. For more information, see the section about the **twinst** script in *IBM Z Workload Scheduler: Planning and Installation*.

Improved security for your environment

Custom certificates in `p12` format are now used for ensuring SSL communication in your environment. `p12` certificates are secure and easy to use. Default certificates are no longer supported. For more information about configuring your environment with custom certificates, see the topic about installing a Dynamic Workload Console server in *IBM Z Workload Scheduler: Planning and Installation* and the topic about customizing the SSL connection between the agents and Z controller when using your certificates (SAF) in *IBM Z Workload Scheduler: Scheduling End-to-end with z-centric Capabilities*.

IBM Z Workload Scheduler Small Programming Enhancements released in July 2023

The following features were provided with IBM Z Workload Scheduler small programming enhancements (SPEs) released for Version 10.1, in July 2023 (APAR PH52702).

Opening an incident through ServiceNow

When an anomaly is detected, the faster you react the better! To have the Z controller open incidents through ServiceNow® and manage issues through a ticketing system, now you can:

- Integrate IBM Z Workload Scheduler with ServiceNow. For detailed information, see the section about opening incidents through ServiceNow in *Managing the Workload* or the [video](#) that summarizes the main integration steps on the [Workload Automation](#) YouTube channel.
- Use the Servicenow plug-in, available for z-centric agents on [Automation Hub](#).

Defining jobs outside of EQQJBLIB and job-library-read exit (EQQUX002)

Typically, you define jobs in the concatenated libraries referenced by EQQJBLIB or with the job-library-read exit (EQQUX002). To define operations whose jobs are located outside of both EQQJBLIB and EQQUX002, you are now provided with the JOBLIBRARY user field. By setting this user field, IBM Z Workload Scheduler dynamically allocates the specified library when required.

For more detailed information, see the section about defining jobs in *Managing the Workload*.

Provide EXECUTE command for WTO workstations (RFE ZWS-I-95)

You can now issue an Execute command on a WTO general workstation, even if the job submission is deactivated on the Z controller.

EQQAUDIT should include records if operation is manually held (RFE ZWS-I-167)

If you update the Manually Hold or NOP option by using the ISPF fast path 5.3 - Automatic Options, the modified value is displayed in the EQQAUDIT report as "CHANGED VALUE".

IBM Z Workload Scheduler Small Programming Enhancements released in November 2022

The following features were provided with IBM Z Workload Scheduler small programming enhancements (SPEs) released for Version 10.1, in November 2022 (APAR PH45777).

Self-Service Catalog enhancements

A new and improved version of the Self-Service Catalog is available. You can now launch your services quickly and easily and check on them at any time by accessing the Self-Service Catalog from any device.

To use the Self-Service Catalog you do not need to be a Workload Automation expert, but you can leverage on services based on automation capabilities in no time, provided that you are connected to the Dynamic Workload Console in Single Sign-On (SSO). For more information, see the section about configuring the Dynamic Workload Console for Single Sign-On on *IBM Z Workload Scheduler: Planning and Installation*.

The IBM® Workload Scheduler scheduler or application designer creates annotations in the Dynamic Workload Console and marks them as services, so that they are available for managing from the Self-Service Catalog interface.

To create and edit SSC-ready annotations, it is recommended that you use the Dynamic Workload Console. For more information about defining SSC-ready annotations, see the online help for annotation definitions in the Dynamic Workload Console.

Deploying z-centric and dynamic agents on Kubernetes to orchestrate your workload on cloud

A cloud deployment ensures access anytime anywhere and is a fast and efficient way to get up and running quickly. It also simplifies maintenance, lowers costs, provides rapid upscale and downscale, minimizes IT requirements and physical on-premises data storage.

To ensure a fast and responsive experience when using IBM Z Workload Scheduler, you can now deploy z-centric and dynamic agents on a Kubernetes cluster. With the sample Helm Chart provided, you install and configure the agent with a single yaml file.

Orchestrating your Kubernetes workload from a containerized agent optimizes allocations of resources in your cluster and reduces the chance of human errors. Find out more in the video [Orchestrating Your Workload on Cloud from IBM Z Workload Scheduler](#) on the [Workload Automation YouTube](#) channel.

For details, see the section about Deploying z-centric and dynamic agents on Kubernetes in *IBM Z Workload Scheduler: Planning and Installation*.

Better performance and stronger security with TLS V1.3

IBM Z Workload Scheduler supports TLS V1.3, which provides a faster and more responsive connection and more secure cipher suites for encryption. By supporting the Server Name Indication (SNI), an extension of the TLS protocol, IBM Z Workload Scheduler is enabled to securely connect any z-centric agent deployed on a cloud cluster.

For details about how to configure TLS, see the section about customizing TLS to connect components in *Customization and Tuning*.

Advanced password encryption

Security is a major concern in today's interconnected world. Government organizations, financial institutions, healthcare providers, and insurance companies are just a few examples of the types of entities who are taking security seriously.

You can optionally encrypt the passwords that you will use while installing, upgrading, and managing IBM Workload Automation. The encryption of the passwords performed by the `secure` command is based on the AES method. For details, see the section about encrypting passwords in *IBM Z Workload Scheduler: Planning and Installation*.

Installing the agents on IBM i with a user different from QSECOFR

You can use a user different from QSECOFR to install z-centric and dynamic agents on IBM i. In this case, the new `allObjAuth` parameter is required when running the `twsinst` command, to indicate that the user has the required ALLOBJ authority. Ensure that the user is existing and is assigned ALLOBJ authorization. The agent is started after connecting to the system with the TWSUSER or the user defined at installation time.

When you upgrade or uninstall the agent, a check is performed to ensure that you are using the same user who performed the installation. If you used `allObjAuth` parameter at installation time, specify it again when upgrading or uninstalling the agent.

The name of the user used to perform the installation is maintained in the `TWA_DATA_DIR/installation/instInfo/instUser` file.

For more information, see *Installing agents on IBM i systems*, *Upgrading agents on IBM i systems*, and *Uninstalling agents on IBM i systems*.

Enable SAF Keyring Support for TCPOPTS (RFE ZWS-I-99)

Reduce vulnerabilities with TCPIP connections by providing RACF key ring support (ZWS-I-115)

To strengthen security, preventing loss of data and service outage, you can now set a single SAF key ring for HTTP and TCP/IP communications. This applies to the communications between:

- Tracker and Z controller
- Data store and Z controller
- Server and remote interface

For details, see the `SSLKEYSTORETYPE` parameter of the `TCPOPTS` statement in *Customization and Tuning*.

IBM Z Workload Scheduler version 10.1 enhancements

IBM Z Workload Scheduler version 10.1 includes the following enhancements.

Managed File Transfer

The File Transfer integration, already available on [Automation Hub](#), now features three important enhancements:

- **Support for File Proxy:** A new component, called File Proxy, is now available for you to install. Through the File Proxy you manage securely files to be transferred with the File Transfer integration on any workstation that runs a supported version of Windows or Linux operating system. All transfer operations are performed by IBM Z Workload Scheduler with no need of third-party products.
- **Support for file transfers via the Workstation-to-Workstation internal protocol:** You can easily transfer files to and from z-centric agents. You no longer have to specify the address of the workstations involved in the file transfer, but you can simply define the workstations from the File Transfer integration. Ensure that the agents are at version 10.1. For more information about the integration, see [File Transfer integration](#).
- **Search on remote file systems:** You can perform search operations on the file system of a remote workstation. On Linux, if you do not specify a path, the search is performed on the working directory. On Windows, the path is required.

A renewed experience of Workload Designer

Workload Designer: a new concept to manage your workload. An innovative infrastructure and design have been thought to simplify the user experience and create a more responsive, fast and fluid user interface. You do not need to switch among the Dynamic Workload Console pages to complete your business workflow anymore. The new Workload Designer contains everything you need. Automate business-critical processes from a single point of access and control.

Automate fast, automate better.

Detecting anomalies in your workload with AI Data Advisor (AIDA)

Starting from this release, a new component is available in IBM Z Workload Scheduler: AI Data Advisor (AIDA), based on Artificial Intelligence and Machine Learning, for early anomaly detection and analysis.

AIDA enables fast and simplified data-driven decision making, for an intelligent workload management. By analyzing historical data and metrics gathered by IBM Workload Scheduler and predicting their future patterns, AIDA identifies anomalies in the trend of Key Performance Indicators (such as the number of completed jobs in the current plan, the job duration, the job end-time) and sends alerts immediately to anticipate and prevent problems and delays. Alerts show up on the Workload Dashboard and can be notified via email.

For more information, see the *AI Data Advisor (AIDA) User's Guide*.

Chatbot based solution to enhance incident management process

When a problem arises, reacting is crucial. Identifying the issue, gathering possible solutions, choosing the best way to proceed are the fundamentals of problem-solving. In this realm, rapid communication becomes critical. By integrating with IBM Z ChatOps, IBM Z Workload Scheduler provides you with a chat tool where you are notified about incidents and can share information with other team members. You are alerted through the chat platform of your choice (such as Microsoft Teams, Slack, or Mattermost) and communicate with the other chat users to share data and perform actions. Collaboration becomes easy, immediate, and effective for promoting teamwork and addressing daily issues.

For detailed information, see the section about posting incidents and sharing information through a chat tool in *Managing the Workload*.

For an overview of this feature, watch the video [Chatbot based solution to manage incidents](#) on the [HCL Workload Automation](#) YouTube channel.

Exposing metrics to monitor your workload

IBM Z Workload Scheduler now exposes metrics about the status of your jobs by using OpenMetrics standard. The data collected can be processed by multiple data analytic tools, creating an easy and unified way to control and monitor your workload.

For details, see the section about exposing metrics to monitor your workload in *Managing the Workload*.

Ensure highest security and safety of your data with TLS V1.2

To ensure highest safety of your data, IBM Z Workload Scheduler now supports the more secure and modern TLS V1.2.

TLS V1.0 and V1.1 are no longer supported. To communicate with the Dynamic Workload Console, master domain manager, and dynamic domain manager through TLS it is required that you configured secure ciphers, such as ECC or AES-GCM ciphers. For details about how to configure TLS, see .

Job running on a z-centric agent is set to Error if query job status is unknown

When a connection problem occurs between the Z controller and a z-centric agent where a job is running, if the query status for that job is unknown, the job is *not* set to Ready but it is set to Error with UNKN.

Rationalization of the JobManager.ini file

The contents of the `JobManager.ini` file have been trimmed and revised to remove some properties which are reserved for internal use. These properties have been moved to a file named `JobManagerInternal.ini`, which is located in the binary directory and is overridden at each product upgrade. This file is not documented because it is not relevant for the final user. For more information about the `JobManager.ini` file, see the section about configuring the agent in *Scheduling End-to-end with z-centric Capabilities*.

New and improved command for password encryption

A new command is now available for encrypting and decrypting passwords that you use with IBM Z Workload Scheduler with the `xor`, `3DES`, and `AES` algorithms. For more information, see the section about encrypting passwords in *IBM Z Workload Scheduler: Planning and Installation*.

Java software prerequisite changed to version 11

Java is a prerequisite for Java job plug-ins and other Java-based functions. With this release, the supported version has been changed to Version 11. This version of Java is not available on HP-UX and Solaris SPARC operating systems. If you install on these operating systems, use the Java Version 11 provided with the operating systems. This change applies to the **twinsinst** script: before launching this command on HP-UX and Solaris SPARC operating systems, provide the path to Java. For more information, see the section about the **twinsinst** command in *IBM Z Workload Scheduler: Planning and Installation*.

SEQSAMP JCL samples have incorrect or obsolete REGION size (RFE 66985)

The JCL samples with a region different from 0 MB have now a region of 64 MB.

DBGGRANT script for DELETE authorization (RFE 103251)

By granting a user access to the reporting tables and views with the **dbgrant** script, users are now enabled to both edit and delete them.

Display USER LEVEL at controller startup (RFE 115863)

At controller and server startup, message EQQZ286I is issued to show the maximum user level supported by the system.

Increasing the limitation of potential successors in the current plan (RFEs 144541, 144596)

Product scalability has been improved by increasing the limitation of potential successors in the current plan from 1000 to 3000.

User-defined JCL skeleton should be available for EQQEMAIL solution (RFE 145320)

You can define your own JCL to be used when an email is sent upon an alert condition, instead of using the predefined EQQEMAIL members. You can use any text for the JCL, no parsing is performed by IBM Z Workload Scheduler. For more details, see the section about sending an email when an alert condition occurs in *Managing the Workload*.

DYNAMICDEL(YES) should be the default (RFE 148581)

In *Customization and Tuning*, all users who significantly use the dynamic addition of special resources are recommended to set BATCHOPT DYNAMICDEL(YES). This setting prevents the size of the data space from increasing in time, with an initial performance degradation that worsens until the space available is exhausted and the batch job terminates with Abend 01D. For this reason, the default value of the DYNAMICDEL parameter in the BATCHOPT statement has been changed from NO to YES.

Pool dynamic workstation shows offline status if resources are unavailable to schedule jobs (RFE 150599 Workstation should show offline for agent pools tied to a DDM)

If a pool dynamic workstation is found without available resources, it is set offline until resources become available again; in that case, the workstation becomes active again.

When the PULSEIVL parameter of the HTTPOPTS or ROUTOPTS statement is set to 0 (which is the default), the heartbeat checking:

- Is run for dynamic domain managers according to the policy established by the Z controller.
- Is not run for z-centric agents.

Automatic health check performed periodically. A message is issued if vital services are no longer available. (RFE 151638 Startup message of the controller is needed when every service is available)

On the Z controller and tracker, a periodic check is performed to ensure that the following vital subtasks are active:

For the Z controller

General Service, Normal Mode Manager, Workstation Analyzer, Event Manager, Submit Task, Event Writer

For the tracker

Submit Task, Event Writer

A message is issued every time one of these subtasks is started, ended, or abended.

Enhance description on the number of occurrences (RFE 151653)

When setting the MAXOCCNUM parameter in BATCHOPT statement, it is advisable not to set the parameter to a larger number than required by actual workload needs, because the increased overhead incurred could cause a noticeable performance degradation.

Several integrations now available on Automation Hub

The following integrations are no longer available with IBM® Workload Scheduler and can be installed from [Automation Hub](#):

- Amazon EC2
- Apache Spark
- IBM BigInsights
- IBM Cloudant
- IBM Cognos
- IBM InfoSphere DataStage
- Enterprise Java Beans (EJB)
- Hadoop Map Reduce
- IBM SoftLayer
- Informatica PowerCenter
- J2EE (JMS)
- Microsoft Azure
- MQTT
- Oozie
- Oracle E-Business Suite

- Salesforce
- SAP BusinessObjects BI
- SAP PI Channel
- Web Services
- IBM WebSphere MQ

Chapter 2. IBM Workload Automation and ITUP

This chapter explains where IBM Workload Automation is placed within the IBM® Tivoli® Unified Process (ITUP).

The IBM® Tivoli® Unified Process provides detailed documentation of IT Service Management processes based on industry best practices, to help users to significantly improve their organization's efficiency and effectiveness. ITUP helps users to easily understand processes, the relationships between processes, and the roles and tools involved in an efficient process implementation.

The processes described in ITUP are strongly aligned with the Information Technology Infrastructure Library (ITIL™) which is based on best practices observed within the IT industry. ITIL™ provides high-level guidance of what to implement, but not how to implement. ITUP contains detailed process diagrams and descriptions to help users understand processes and their relationships, making ITIL™ recommendations easier to follow.

ITUP is based on the IBM Process Reference Model for IT™ (PRM-IT), which was jointly developed by IBM® Global Services and Tivoli®. PRM-IT provides detailed process guidance for all activities that fall under the office of the CIO, including, but not limited to, IT Service Management.

The ITUP processes

About this task

ITUP describes a comprehensive set of processes within an IT organization. Each process is defined by:

Tool Mentors that describe best practice use of IBM® tools in a process context

Tool mentors help users identify which IBM® products and solutions can be used to perform specific process activities and details their appropriate use. By following this guidance users can reduce time, effort, and errors, and get the maximum value out of their investments.

Role definitions, responsibilities, and resources

ITUP describes the roles and responsibilities of all actors in the process model. Users can identify their roles and understand the activities they need to perform and the tools available to help them.

Work products and other information

ITUP describes all work products, often referred to as artifacts, produced as output or required as input by processes and activities. Other information such as key terms and concepts are also defined.

Scenarios describing common problems and best practice solutions

Scenarios help users understand how real world problems can be addressed with process improvement and integration, proper tool use, and defined roles and responsibilities.

Service execution and workload management

About this task

Among the processes supported by IBM Workload Automation, service execution is one of the key IT processes described by ITUP. Its mission is to deliver operational services to the IT infrastructure and to the enterprise.

The main service execution activities supported by IBM Workload Automation are:

- Deliver service
- Manage delivery resources
- Manage workload

Workload management has the target to maximize the utilization of task execution resources and to minimize the total time that is required to deliver the output of task processing. This activity operates at both a macro-level and micro-level to prepare work schedules and to pre-process work items where necessary so that the delivery resources can be matched to the demands of the flow of work in an optimal fashion.

The objectives of workload scheduling focus on:

- Managing the execution of activities according to business calendars, time constraints, and resource availability.
- Managing activities that have interactions between each other and have dependencies with each other and external entities and events.
- Enabling integration with business application environments like SAP R3 and PeopleSoft, and managing the running of activities in those environments.
- Managing the lifecycle for defining the activities to run and their running.
- Enabling monitoring and control on the running of these activities and collecting results and historical running data.

Managing workload with IBM Workload Automation

About this task

IBM Workload Automation is a portfolio of products provided by IBM® to automate all workload management tasks. The scheduling features of IBM Workload Automation help you plan every phase of your workload production. During the processing period, the production control programs manage the production environment and automate most operator activities. The schedulers prepare jobs for running, resolve interdependencies, and launch and track jobs. Because jobs start running as soon as their dependencies are satisfied, idle time is minimized, and throughput improves significantly. Jobs never run out of sequence, and, if a job fails, the schedulers can handle the recovery process with little or no operator intervention.

Workload management is based on a database that contains the definitions of the scheduling objects. There are two versions of the scheduling objects database depending on the placement of the main workload controller: it can be based on a mainframe computer (in this case, z/OS®) or on a distributed platform. Some of the scheduling objects can exist in both of the databases, some apply only to the distributed platform, and others might apply only to the mainframe platform.

The minimum set of object definitions that are required to produce a workload consists of a workstation, a job, and a job stream. Other required scheduling objects might be predefined and exist by default.

A workstation is a definition that represents a computer system or another entity that is capable of running specific tasks, and that has the ability to report the status of task execution to the scheduler. With the IBM Workload Automation interfaces, you can identify the physical resources associated with the workstations.

A job is the representation of a task (an executable file, program, or command) that is scheduled and launched by the scheduler. The job is run by a workstation and, after running, has a status that indicates if the run was successful or not. A job definition can specify information on what to do whenever its run was not successful. Jobs not included in a job stream do not have any attribute for running, and are only the description of a task with a definition on how to perform it in a form that is known to the specified workstation.

A job stream represents a container for related jobs and organizes them, in terms of run time, sequencing, concurrency limitations, repetitions, assigning priority or resources, and so on. Job streams are the macro elements of the workload that you manage.

The scheduling plan is the to-do list that tells IBM Workload Scheduler or IBM Z Workload Scheduler what jobs to run, and what dependencies must be satisfied before each job is launched. IBM Workload Scheduler or IBM Z Workload Scheduler builds the plan using the elements that are stored in the scheduling database.

The running of a plan requires tracking to identify possible problems that can impact the effective delivery of the work products. It is possible to perform the tracking from a Web-based Java™ interface, the Dynamic Workload Console, on either of the platforms (z/OS® and distributed). As an alternate interface to the Dynamic Workload Console on the z/OS® platform you can also use the ISPF panel interface, and on the distributed platforms you can use the command-line interface.

Chapter 3. IBM Workload Scheduler

The IBM Workload Scheduler scheduling features help you plan every phase of production. During the plan processing period, the IBM Workload Scheduler production control programs manage the production environment and automate most operator activities. IBM Workload Scheduler prepares jobs for execution, resolves interdependencies, and launches and tracks each job. Because jobs start running as soon as their dependencies are satisfied, idle time is minimized, and throughput improves significantly. Jobs always run in sequence and, if a job fails, IBM Workload Scheduler handles the recovery process with little or no operator intervention.

Overview

The next sections provide an outline of IBM Workload Scheduler.

What is IBM Workload Scheduler

IBM Workload Scheduler is composed of the following parts:

IBM Workload Scheduler engine

The scheduling engine. It runs on every computer of a IBM Workload Scheduler network. During installation, the engine is configured for the role that the workstation will play within the scheduling network, such as master domain manager, domain manager, or agent.

The Dynamic Workload Console

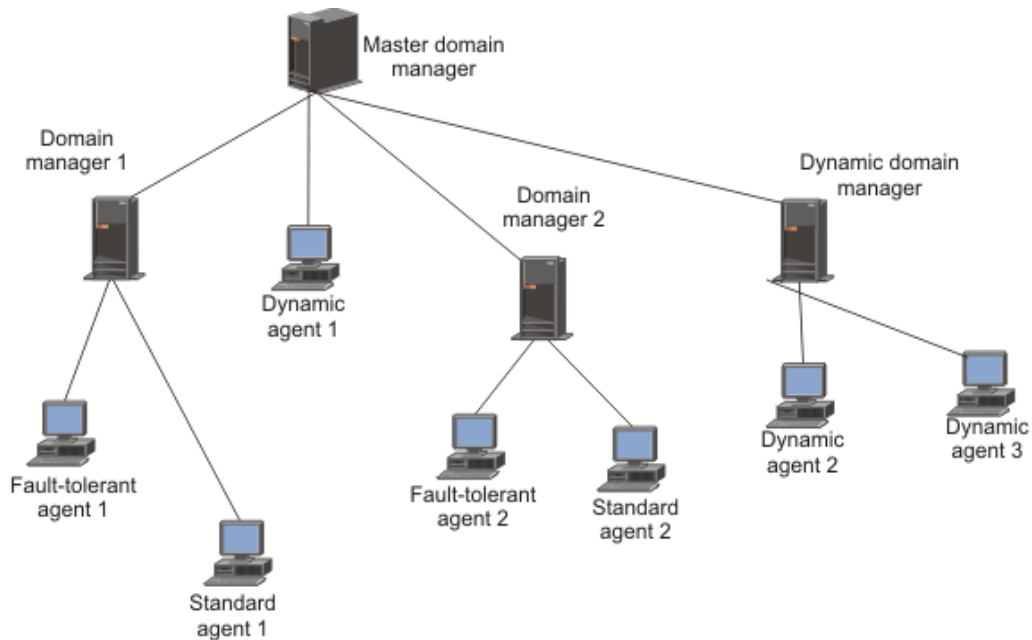
Is Web-based, light, powerful, and user friendly. It can be used on any computer that has a web browser and provides access to all the current IBM Workload Scheduler functions. It is the strategic graphical user interface for the entire IBM Workload Automation portfolio.

The IBM Workload Scheduler network

An IBM Workload Scheduler network is made up of the workstations, or CPUs, on which jobs and job streams are run.

An IBM Workload Scheduler network contains at least one IBM Workload Scheduler domain, the master domain, in which the master domain manager is the management hub. Additional domains can be used to divide a widely distributed network into smaller, locally managed groups.

Figure 1. IBM Workload Scheduler network made up by two domains



Using multiple domains reduces the amount of network traffic by reducing the amount of communication required between the master domain manager and other computers.

In a single domain configuration, the master domain manager maintains communication with all of the workstations in the scheduling network.

In a multi-domain configuration, the master domain manager communicates with the workstations in its domain and with the subordinate domain managers. The subordinate domain managers, in turn, communicate with the workstations in their domains and with their subordinate domain managers. Multiple domains also provide fault-tolerance by limiting the problems caused by losing a domain manager to a single domain. To limit the effects further, you can designate backup domain managers to take over if their domain managers fail.

Every time the production plan is created or extended the master domain manager creates a production control file, named Symphony™. IBM Workload Scheduler is then restarted in the network, and the master domain manager sends a copy of the new production control file to each of its automatically linked agents and subordinate domain managers. The domain managers, in turn, send copies to their automatically linked agents and subordinate domain managers.

Once the network is started, scheduling messages like job starts and completions are passed from the agents to their domain managers, through the parent domain managers to the master domain manager. The master domain manager then broadcasts the messages throughout the hierarchical tree to update the production control files of domain managers and fault tolerant agents running in Full Status mode.

Manager and agent types

About this task

Primarily, workstation definitions refer to physical workstations. However, in the case of extended and network agents, the workstations are logical definitions that must be hosted by a physical IBM Workload Scheduler workstation.

IBM Workload Scheduler workstations can be of the following types:

Master domain manager (MDM)

The domain manager in the topmost domain of a IBM Workload Scheduler network. It either contains or connects to the relational database that stores the scheduling object definitions. It creates or updates the production file when the plan is created or extended and distributes it in the network. It performs all logging and reporting for the network.

Backup master

A fault-tolerant agent or domain manager capable of assuming the responsibilities of the master domain manager for automatic workload recovery.

Domain manager

The management hub in a domain. All communications to and from the agents in a domain are routed through the domain manager.

Backup domain manager

A fault-tolerant agent capable of assuming the responsibilities of its domain manager.

Dynamic domain manager

An installed component in a distributed IBM Workload Scheduler network that is the management hub in a domain. All communication to and from the dynamic agents in the domain is routed through the dynamic domain manager.

Backup dynamic domain manager

A workstation which can act as a backup for the dynamic domain manager, when problems occur. It is effectively a dynamic domain manager, waiting to be activated. Its use is optional.

Fault-tolerant agent (FTA)

A workstation capable of resolving local dependencies and launching its jobs in the absence of a domain manager.

Dynamic agent

Dynamic agents run the same types of jobs as fault-tolerant agents. From classic job types such as docommand and scripts, to more specific job types such as Java, database, file transfers, InfoSphere DataStage, Cognos, and OSLC to name a few. Dynamic agents can be defined in pools and dynamic pools.

Run workload dynamically

It communicates with the server the status of its resources. In this way the product is able to dynamically run your workload to the best available resources by:

- Automatically discovering scheduling environment resources.
- Automatically following resource changes

- Requesting additional resources when needed
- Matching job requirements to available resources
- Controlling and optimizing use of resources

The characteristics listed above provides high availability and load balancing potentialities to your environment and well suite virtualized environments.

When a job is submitted, either as part of a job stream in the plan or through ad hoc submission, IBM Workload Scheduler checks the job requirements, the available resources and the related characteristics and submits the job to the resource that best meets the requirements to run it.

Run both existing job types and job types with advanced options

It can run:

- Existing job types. For example docommand and scripts.
- Job types with advanced options, both those supplied with the product and the additional types implemented through the custom plug-ins. For example, those supplied with the product are DB2®, file transfer, and web services. Those implemented through the custom plug-ins are the ones you developed using the Workload Automation, Lutist Development Kit. To run these job types you must also install the Java™ run time.

Manage dynamic workload broker logical resource

It can remotely run, from the agent, the dynamic workload broker **resource** command on the server. To manage the **resource** command you must also install the Java™ run time.

Pool

Pools are groups in which you add specific dynamic agent workstations according to the requirements of the job. Jobs are assigned dynamically to the best available dynamic agent applying a load balancing policy by choosing the agent with the lesser amount of jobs to run.

Dynamic pool

Dynamic pools are groups in which you specify the requirements and IBM® Workload Scheduler selects the dynamic agents that meet the requirements. Jobs are assigned dynamically to the best available dynamic agent applying an optimization policy to identify the best agent.

Standard agent

A workstation that launches jobs only under the direction of its domain manager. It is not fault-tolerant.

Extended agent

A logical workstation definition that helps you launch and control jobs on other systems and applications, such as PeopleSoft, SAP, and z/OS®.

z-centric agent

Runs jobs scheduled from IBM Z Workload Scheduler but is installed in the IBM Workload Scheduler environment. It has no fault-tolerance and communicates directly with the IBM Z Workload Scheduler controller

through the RESTful HTTP interface. In IBM Z Workload Scheduler it has the same functionality as a computer automatic workstation even though it runs in the distributed environment.

Network Agent

A logical workstation definition for creating dependencies between jobs and job streams in separate IBM Workload Scheduler networks.

Topology

About this task

A key to choosing how to set up IBM Workload Scheduler domains for an enterprise is the concept of localized processing. The idea is to separate or localize the enterprises's scheduling needs based on a common set of characteristics.

Common characteristics are things such as geographical locations, business functions, and application groupings. Grouping related processing can limit the amount of interdependency information that needs to be communicated between domains.

The benefits of localizing processing in domains are:

- Decreased network traffic. Keeping processing localized to domains eliminates the need for frequent interdomain communications.
- Provides a convenient way to tighten security and simplify administration. Security and administration can be defined at, and limited to, the domain level. Instead of network-wide or workstation-specific administration, you can have domain administration.
- Network and workstation fault tolerance can be optimized. In a multiple domain IBM Workload Scheduler network, you can define backups for each domain manager, so that problems in one domain do not disrupt operations in other domains.

Networking

About this task

The following questions will help in making decisions about how to set up your enterprise's IBM Workload Scheduler network. Some questions involve aspects of your network, and others involve the applications controlled by IBM Workload Scheduler. You may need to consult with other people in your organization to resolve some issues.

- How large is your IBM Workload Scheduler network? How many computers does it hold? How many applications and jobs does it run?

The size of your network will help you decide whether to use a single domain or the multiple domain architecture. If you have a small number of computers, or a small number of applications to control with IBM Workload Scheduler, there may not be a need for multiple domains.

- How many geographic locations will be covered in your IBM Workload Scheduler network? How reliable and efficient is the communication between locations?

This is one of the primary reasons for choosing a multiple domain architecture. One domain for each geographical location is a common configuration. If you choose single domain architecture, you will be more reliant on the network to maintain continuous processing.

- Do you need centralized or decentralized management of IBM Workload Scheduler?

A IBM Workload Scheduler network, with either a single domain or multiple domains, gives you the ability to manage IBM Workload Scheduler from a single node, the master domain manager. If you want to manage multiple locations separately, you can consider the installation of a separate IBM Workload Scheduler network at each location. Note that some degree of decentralized management is possible in a stand-alone IBM Workload Scheduler network by mounting or sharing file systems.

- Do you have multiple physical or logical entities at a single site? Are there different buildings, and several floors in each building? Are there different departments or business functions? Are there different applications?

These may be reasons for choosing a multi-domain configuration. For example, a domain for each building, department, business function, or each application (manufacturing, financial, engineering, etc.).

- Do you run applications, like SAP R/3, that will operate with IBM Workload Scheduler?

If they are discrete and separate from other applications, you may choose to put them in a separate IBM Workload Scheduler domain.

- Would you like your IBM Workload Scheduler domains to mirror your Windows™ domains?

This is not required, but may be useful.

- Do you want to isolate or differentiate a set of systems based on performance or other criteria?

This may provide another reason to define multiple IBM Workload Scheduler domains to localize systems based on performance or platform type.

- How much network traffic do you have now?

If your network traffic is manageable, the need for multiple domains is less important.

- Do your job dependencies cross system boundaries, geographical boundaries, or application boundaries? For example, does the start of Job1 on CPU3 depend on the completion of Job2 running on CPU4?

The degree of interdependence between jobs is an important consideration when laying out your IBM Workload Scheduler network. If you use multiple domains, you should try to keep interdependent objects in the same domain. This will decrease network traffic and take better advantage of the domain architecture.

- What level of fault-tolerance do you require?

An obvious disadvantage of the single domain configuration is the reliance on a single domain manager. In a multi-domain network, the loss of a single domain manager affects only the agents in its domain.

IBM Workload Scheduler components

About this task

IBM Workload Scheduler uses several manager processes to efficiently segregate and manage networking, dependency resolution, and job launching. These processes communicate among themselves through the use of message queues. Message queues are also used by the Console Manager (conman) to integrate operator commands into the batch process.

On any computer running IBM Workload Scheduler there are a series of active management processes. They are started as a system service, or by the `StartUp` command. The following are the main processes of a fault-tolerant agent:

Netman

The network management process that establishes network connections between remote `mailman` processes and local `writer` processes.

Mailman

The mail management process that sends and receives inter-CPU messages.

Batchman

The production control process. Working from Symphony™, the production control file, it runs jobs streams, resolves dependencies, and directs `jobman` to launch jobs.

Writer

The network `writer` process that passes incoming messages to the local `mailman` process.

Jobman

The job management process that launches and tracks jobs under the direction of `batchman`.

The following are the main processes of a dynamic agent:

agent

The network management process that establishes network connections to the master domain manager or the dynamic domain manager.

JobManager

The process that manages all of the dynamic agent operations, such as, sending events, and directing `TaskLauncher` to launch jobs.

TaskLauncher

The job management process that launches and tracks jobs under the direction of `JobManager`.

JobManagerGW

Manages the communication between dynamic agents and the master or dynamic domain manager through a firewall.

Monman

The event management process that turns on the event monitoring engine on the workstation.

In addition, IBM Workload Scheduler uses two command line interfaces:

Composer

The command-line program used to define, manage, and store scheduling objects in the IBM Workload Scheduler database. The composer command-line program can be installed and used on any computer connected through TCP/IP to the system where the master domain manager is installed. It does not require the installation of a IBM Workload Scheduler workstation as a prerequisite. It communicates through HTTP/HTTPS with the master domain manager where the relational database management system (RDBMS) is installed. The HTTP/HTTPS communication setup and the authentication check are managed by the WebSphere® Application Server - Express® infrastructure. The composer uses edit files to update the scheduling database.

Conman

The console manager. It is the user interface for plan running activities by means of the command line interface. `Conman` writes information that is received by either the local `netman` or `mailman` processes.

IBM Workload Scheduler scheduling objects

About this task

Scheduling with IBM Workload Scheduler includes the capability to do the following:

- Schedule jobs across a network.
- Group jobs into job streams according, for example, to function or application.
- Set limits on the number of jobs that can run concurrently.
- Create job streams based on day of the week, on specified dates and times, or by customized calendars.
- Ensure correct processing order by identifying dependencies such as successful completion of previous jobs, availability of resources, or existence of required files.
- Set automatic recovery procedures for unsuccessful jobs.
- Forward incomplete jobs to the next production day.

Starting from version 8.3, the IBM Workload Scheduler scheduling objects are stored in a relational database. This results in a significant improvement, in comparison with previous versions, of how objects are defined and managed in the database. Each object can now be managed independently without having to use lists of scheduling objects like calendars, parameters, prompts and resources. The command syntax used to define and manage these objects has also become direct and powerful.

IBM Workload Scheduler administrators and operators work with these objects for their scheduling activity:

Workstation

Also referred to as *CPU*. Usually an individual computer on which jobs and job streams are run. Workstations are defined in the IBM Workload Scheduler database as a unique object. A workstation definition is required for every computer that executes jobs or job streams in the IBM Workload Scheduler network.

Workstation class

A group of workstations. Any number of workstations can be placed in a class. Job streams and jobs can be assigned to execute on a workstation class. This makes replication of a job or job stream across many workstations easy.

Domain

All workstations in a distributed IBM Workload Scheduler network are organized in one or more domains, each of which consists of one or more agents and a domain manager acting as the management hub. Most communication to and from the agents in the domain is routed through the domain manager. You can organize all agents in your network in a single domain, or in multiple domains. A single domain network consists of a master domain manager and any number of agents. A multiple domain network consists of a master domain manager, any number of lower tier domain managers, and any number of agents in each domain.

Folder

A smart workflow folder is a container of jobs, job streams, and other folders. Use smart workflow folders to organize your jobs and job streams according to your lines of business or other custom categories. A folder can contain one or more jobs or job streams, while each job stream can be associated to one folder. If no folder is defined for a job stream, the root folder (/) is used by default.

Job

A script or command, run on the user's behalf, run and controlled by IBM Workload Scheduler.

Job stream

A list of jobs that run as a unit (such as a weekly backup application), along with run cycles, times, priorities, and other dependencies that determine the exact order in which the jobs run.

Calendar

A list of scheduling dates. Each calendar can be assigned to multiple job streams. Assigning a calendar to a job stream causes that job stream to run on the dates specified in the calendar. A calendar can be used as an inclusive or as an exclusive run cycle.

Run cycle

A cycle that specifies the days that a job stream is scheduled to run. Run cycles are defined as part of job streams and may include calendars that were previously defined. There are three types of run cycles: a Simple run cycle, a Weekly run cycle, or a Calendar run cycle (commonly called a calendar). Each type of run cycle can be inclusive or exclusive. That is, each run cycle can define the days when a job stream is included in the production cycle, or when the job stream is excluded from the production cycle.

Run cycle group

A run cycle group is a list of run cycles that are combined together to produce a set of run dates. You can optionally define a run cycle group for your job stream instead of, or in addition to, a number of single run cycles. Unlike run cycles, run cycle groups are not tied to a specific job stream and therefore can be reused in several different job streams to apply the same scheduling rules each time.

Run cycles are organized into subsets within a run cycle groups. The subsets are in a logical OR relationship with each other. They enable you to apply exclusive run cycles to a subset of the inclusive run cycles. The result of the run cycle group is a run date or a set of run dates.

Prompt

An object that can be used as a dependency for jobs and job streams. A prompt must be answered affirmatively for the dependent job or job stream to launch. There are two types of prompts: predefined and ad hoc. An ad hoc prompt is defined within the properties of a job or job stream and is unique to that job or job stream. A predefined prompt is defined in the IBM Workload Scheduler database and can be used by any job or job stream.

Resource

An object representing either physical or logical resources on your system. Once defined in the IBM Workload Scheduler database, resources can be used as dependencies for jobs and job streams. For example, you can define a resource named *tapes* with a unit value of two. Then, define jobs that require two available tape drives as a dependency. Jobs with this dependency cannot run concurrently because each time a job is run the *tapes* resource is in use.

Variable and variable table

A variable can be used to substitute values in scheduling objects contained in jobs and job streams; that is, in JCL, log on, prompts dependencies, file dependencies, and recovery prompts. The values are replaced in the job scripts at run time. Variables are global (that is, they can be used on any agent of the domain) and are defined in the database in groups called `variable tables`.

Parameter

A parameter can be used to substitute values in jobs and job streams just like global variables. The difference is that a parameter is defined on the specific workstation where the related job is to run and has no global effect, but only on that specific workstation. Parameters cannot be used when scripting extended agent jobs.

User

On Windows™ workstations, the user name specified in the `Logon` field of a job definition must have a matching user definition. The definitions provide the user passwords required by IBM Workload Scheduler to launch jobs.

Event rule

A scheduling event rule defines a set of actions that are to run upon the occurrence of specific event conditions. The definition of an event rule correlates events and triggers actions. When you define an event rule, you specify one or more events, a correlation rule, and the one or more actions that are triggered by those events. Moreover, you can specify validity dates, a daily time interval of activity, and a common time zone for all the time restrictions that are set.

Workload application

A workload application is an IBM Workload Scheduler database object that acts as a container for one or more job streams. You can use workload applications to standardize a workload automation solution so that the

solution can be reused in one or more IBM Workload Scheduler environments thereby automating business processes.

You prepare a workload application template in a source IBM Workload Scheduler environment and then export it so that it can be deployed in a target environment. The export process extracts from the source environment all of the elements necessary to reproduce the solution in another environment. It produces a compressed file containing a number of files required to import the workload application into the target environment.

You can control how jobs and job streams are processed with the following attributes:

Dependencies

Conditions that must be satisfied before a job or job stream can run. You can set the following types of dependencies:

- A predecessor job or job stream must have completed successfully.
- One or more specific resources must be available.
- Access to specific files must be granted.
- An affirmative response to a prompt.

Time constraints

Conditions based on time, such as:

- The time at which a job or job stream should start.
- The time after which a job or job stream cannot start.
- The repetition rate at which a job or job stream is to be run within a specified time slot.

Job priority

A priority system according to which jobs and job streams are queued for execution.

Job fence

A filter defined for workstations. Only jobs and job streams whose priority exceeds the job fence value can run on a workstation.

Limit

Sets a limit to the number of jobs that can be launched concurrently on a workstation.

The production process

IBM Workload Scheduler production is based on a plan that runs in a *production period*. The production period is defined by the user when creating or extending the production plan. It can span from a few hours to some days (by default it is a 24 hours period). Before the start of each production period, IBM Workload Scheduler executes a program that creates the production plan starting from the modeling data stored in the database and from an intermediate plan called *preproduction plan*. Then another program includes the uncompleted schedules from the previous production period into the current plan and another different program logs all the statistics of the previous production into an archive.

All of the required information for that production period is placed into a production control file named Symphony™. During the production period, the production control file is continually being updated to reflect the work that needs to be done, the work in progress, and the work that has been completed. A copy of the Symphony™ file is sent to all subordinate domain managers and to all the fault-tolerant agents in the same domain. The subordinate domain managers distribute their copy to all the fault-tolerant agents in their domain and to all the domain managers that are subordinate to them, and so on down the line. This causes fault-tolerant agents throughout the network to continue processing even if the network connection to their domain manager is down. From the graphical interfaces or the command line interface, the operator can view and make changes in the current production by making changes in the Symphony™ file.

IBM Workload Scheduler processes monitor the production control file and make calls to the operating system to launch jobs as required. The operating system runs the job, and in return informs IBM Workload Scheduler if the job completed successfully or not. This information is entered into the production control file to indicate the status of the job.

Scheduling

About this task

Scheduling objects are workstations, workstation classes, domains, jobs, job streams, resources, prompts, calendars, variables and variable tables, parameters, event rules, and run cycle groups. Scheduling objects are managed with the Composer program and are stored in the IBM Workload Scheduler database. To create or modify an object, you can use either the IBM Workload Scheduler command line interface or one of the graphical interfaces.

Scheduling includes the following tasks:

- Defining and maintaining workstations
- Defining scheduling objects
- Defining job streams
- Starting and stopping production processing
- Viewing and modifying jobs and job streams.

Creating job streams

The primary processing task of IBM Workload Scheduler is to run job streams. A job stream is an outline of batch processing consisting of a list of jobs. Job streams can be defined using either the command line interface or one of the graphical interfaces. Using either graphical interface you can easily create and modify job streams. You can use their job stream editors to work with the jobs and the `follows` dependencies between the jobs, as well as the job stream run cycles. You can also easily specify time restrictions, resource dependencies, file dependencies, and prompt dependencies at the job stream level.

Job streams can be defined as *draft*. A draft job stream is not considered when resolving dependencies and is not added to the production plan. It becomes actual only after the *draft* keyword is removed from its definition, and the `JnextPlan` command is run to add it to the preproduction plan and so to the production plan.

Setting job recovery

About this task

When defining a job, consider that in some instances it might not complete successfully. The administrator can define a recovery option and recovery actions when defining the job. The following recovery options are available:

- Do not continue with the next job. This stops the execution of the job stream and puts it in the *stuck* state. This is the default action.
- Continue with the next job.
- Run the job again.

Optionally, a recovery prompt can be associated to the job. A recovery prompt is a local prompt to display when the job completes unsuccessfully. Processing does not continue until the prompt is answered affirmatively.

Another option is to define a recovery job that can be run in the place of the original job if it completes unsuccessfully. The recovery job must have been defined previously. Processing stops if the recovery job does not complete successfully.

Defining and managing mission-critical jobs

About this task

Job schedulers can use the IBM Workload Scheduler command line or the Dynamic Workload Console to flag jobs as mission-critical and specify their deadlines. A critical job and all its predecessors make up what is called a *critical network*. At planning time, IBM Workload Scheduler calculates the start time of the critical job and of each of its predecessors starting from the critical job deadline and estimated duration. While the plan runs, this information is dynamically kept up-to-date based on how the plan is progressing. If a predecessor, or the critical job itself, is becoming late, IBM Workload Scheduler automatically prioritizes its submission and promotes it to get more system resources and thus meet its deadline.

You can also customize the **riskConfidence** `optman` parameter to define when a critical job should be flagged as `High Risk`. The job status is calculated comparing the confidence factor of completing before deadline and the percentage specified in this parameter. If the probability of completing before the deadline is below **riskConfidence**, then the critical job is considered at high risk. For more information about the **riskConfidence** `optman` parameter, see the section about global options in *Administration Guide*.

Within a critical network, IBM Workload Scheduler dynamically identifies the path of predecessors that is potentially most at risk; this is called the *critical path*. IBM Workload Scheduler calculates the level of risk that each critical job has of missing its deadline; a high risk indicates that the estimated end of the critical job is after its deadline while a potential risk indicates that some predecessors of the critical job have a warning condition, for example are late or in error.

The Dynamic Workload Console provides specialized views for tracking the progress of critical jobs and their predecessors. Job schedulers and operators can access the views from the Dashboard or by creating Monitor Critical Jobs tasks.

The initial view lists all critical jobs for the engine, showing the status: normal, potential risk, or high risk. From this view, an operator can navigate to see:

- The hot list of jobs that put the critical deadline at risk.
- The critical path.
- Details of all critical predecessors.
- Details of completed critical predecessors.
- Job logs of jobs that have already run.

Using the views, operators can monitor the progress of the critical network, find out about current and potential problems, release dependencies, and rerun jobs.

For example:

1. To flag a critical job and follow it up, the Job scheduler opens the Workload Designer on the Dynamic Workload Console, marks the specific job as critical, and sets the deadline for 5 a.m.

When `jnextPlan` is run, the critical start dates for this job, and all the jobs that are identified as its predecessors, are calculated.

2. To track a specific critical job, the operator proceeds as follows:

- a. The operator checks the dashboards and sees that there are critical jobs scheduled on one of the engines.
- b. He clicks the link to get a list of the critical jobs.

The specific job shows a `Potential Risk` status.

- c. He selects the job and clicks **Hot List** to see the predecessor job or jobs that are putting the critical job at risk.

One of the predecessor jobs is listed as being in error.

- d. He selects the job and clicks **Job log**.

The log shows that the job failed because of incorrect credentials for a related database.

- e. After discovering that the database password was changed that day, he changes the job definition in the symphony file and reruns the job.
- f. When he comes back to the dashboard, he notices that there are no longer any jobs in potential risk. Also, the critical jobs list that was opened when clicking on the potential risk link no longer shows the critical job after the job is rerun.
- g. The job is now running after being automatically promoted, getting higher priority for submission and system resources.
- h. No further problems need fixing and the critical job finally completes at 4.45 a.m.

Scheduling workload dynamically

About this task

You can choose to set IBM Workload Scheduler to dynamically associate your submitted workload (or part of it) to the best available resources at run time.

The IBM Workload Scheduler installation process includes the option to install the dynamic scheduling capability. If you select this option, you get the following functionality:

- Automatically discover scheduling environment resources
- Match job requirements to available resources
- Control and optimize use of resources
- Automatically follow resource changes
- Request additional resources when needed

You can submit IBM Workload Scheduler jobs, including jobs defined to run on extended agents, as well as J2EE applications (if you selected the option to schedule J2EE at installation time). To schedule workload dynamically, you:

1. Use the Dynamic Workload Console to define the agents you want to use for running workload as logical resources or groups of resources.
2. Update your IBM Workload Scheduler job definitions to make as destination CPU the dynamic workload broker workstation (this workstation works as a bridge between the scheduler engine and the pool of resources)
3. For every IBM Workload Scheduler job, add a JSDL (Job Submission Description Language) job definition where you match the job with required resources, candidate hosts, and scheduling and optimization preferences. Use the Dynamic Workload Console to do this easily.

When a job is thus submitted, either as part of a job stream in the plan or through ad hoc submission, IBM Workload Scheduler checks the job requirements, the available resources and the related characteristics and submits the job to the resource that best meets the requirements.

Running production

Production consists of taking the definitions of the scheduling objects from the database, together with their time constraints and their dependencies, and building and running the production control file.

Running the plan

About this task

The production plan contains information about which jobs to run, on which fault-tolerant agent, and what dependencies must be satisfied before each job is launched. IBM Workload Scheduler creates the production plan starting from the modeling data stored in the database and from an intermediate plan called the preproduction plan. The preproduction plan is automatically created and managed by the product. To avoid problems, the database is locked during the generation of the plan and is unlocked when the generation completes or if an error condition occurs. The preproduction plan is used to identify in advance the job stream instances and the external follows job stream dependencies involved in a specified time period.

You use the `JnextPlan` command on the master domain manager to generate the production plan and distribute it across the IBM Workload Scheduler network.

To generate and start a new production plan, IBM Workload Scheduler performs the following steps:

1. Updates the preproduction plan with the objects defined in the database that were added or updated since the last time the plan was created or extended.
2. Retrieves from the preproduction plan the information about the job streams to run in the specified time period and saves it in an intermediate production plan.
3. Includes in the new production plan the uncompleted job streams from the previous production plan.
4. Creates the new production plan and stores it in a file named Symphony™. The production plan data is also replicated in the database.
5. Distributes a copy of the Symphony™ file to the workstations involved in the new product plan processing.
6. Logs all the statistics of the previous production plan into an archive.
7. Updates the job stream state in the preproduction plan.

The copy of the newly-generated Symphony™ file is used starting from the top domain's fault-tolerant agents and domain managers of the child domains and down the tree to all subordinate domains.

Each fault-tolerant agent that receives the production plan can continue processing even if the network connection to its domain manager goes down.

At each destination fault-tolerant agent, IBM Workload Scheduler performs the following actions to manage job processing:

1. Accesses the copy of the Symphony™ file and reads the instructions about which job to run.
2. Makes calls to the operating system to launch jobs as required.
3. Updates its copy of the Symphony™ file with the job processing results and sends notification back up the tree to the master domain manager and to all full status fault-tolerant agents. The original copy of the Symphony™ file is stored on the master domain manager and the copies stored on the backup master domain managers, if defined, are updated accordingly.

This means that during job processing, each fault-tolerant agent has its own copy of the Symphony™ file updated with the information about the jobs it is running (or that are running in its domain and child domains if the fault-tolerant agent is full-status or a domain manager), and the master domain manager (and backup master domain manager if defined) has the copy of the Symphony™ file that contains all updates coming from all fault-tolerant agents. In this way the Symphony™ file on the master domain manager is kept up-to-date with the jobs still to run, the jobs running, and the jobs already completed.

After the production plan is generated for the first time, it can be extended to the next production period with the `JnextPlan` command. The Symphony™ file is refreshed with the latest changes and redistributed throughout the network.

Running job streams

Depending on their run cycle definition, job streams are taken from the IBM Workload Scheduler database and automatically inserted into the current production plan.

While the job stream is in the plan, and has not completed, you can still modify any of its components. That is, you can modify the job stream properties, the properties of its jobs, their sequence, the workstation or resources they use, and so on, to satisfy last-minute requirements.

You can also hold, release, or cancel a job stream, as well as change the maximum number of jobs within the job stream that can run concurrently. You can change the priority previously assigned to the job stream and release the job stream from all its dependencies.

Last minute changes to the current production plan include the possibility to submit jobs and job streams that are already defined in the IBM Workload Scheduler database but were not included in the plan. You can also submit jobs that are being defined ad hoc. These jobs are submitted to the current plan but are not stored in the database.

Starting from version 8.3, you can create and manage multiple instances of the same job stream over a number of days or at different times within the same day. This new feature introduced the possibility to have in the same plan more than one instance of the same job stream with the same name. Each job stream instance is identified by the job stream name, the name of the workstation where it is scheduled to run, and by the start time defined in the preproduction plan.

Monitoring

About this task

Monitoring is done by listing plan objects. Using lists, you can see the status of all or of subsets of the following objects in the current plan:

- Job stream instances
- Job instances
- Domains
- Workstations
- Resources
- File dependencies
- Prompt dependencies

You can use these lists also to manage some of these objects. For example, you can reallocate resources, link or unlink workstations, kill jobs, or switch domain manager.

Additionally, you can monitor the daily plan with Tivoli Business Service Manager, an object-oriented systems management application that provides monitoring and event management of resources, applications, and subsystems, that is integrated with IBM Workload Scheduler.

Network managers can use IBM Workload Scheduler/NetView, a NetView® application, to monitor and diagnose IBM Workload Scheduler networks from a NetView® management node. It includes a set of submaps and symbols to view IBM Workload Scheduler networks topographically, and determine the status of job scheduling activity and critical IBM Workload Scheduler processes on each workstation. Menu actions are provided to start and stop IBM Workload Scheduler processing and to run *conman* on any workstation in the network.

Controlling with IBM® Tivoli® Monitoring

About this task

IBM® Tivoli® Monitoring is a product that applies pre-configured best practices to the automated monitoring of essential system resources. It helps you to detect bottlenecks and other potential problems and provides you with the means for automatic recovery from critical situations. In this way it eliminates the need for system administrators to manually scan through extensive performance data.

IBM Workload Scheduler integrates with IBM® Tivoli® Monitoring through the installation of a customizable agent, the IBM Z Monitoring agent that collects and distributes data to the Tivoli® Enterprise Monitoring Server.

By using data providers that are the interfaces of the IBM Z Monitoring agent, you can collect data from data sources, client programs, URLs, scripts, relation tables, or SNMP agents and transfer the collected data to the custom IBM Z Monitoring agent that monitors the collected data.

The IBM Z Monitoring agent uses two different data providers to integrate with IBM Workload Scheduler:

- A script data provider that collects IBM Workload Scheduler data by running a script at regular intervals and parsing the script output to look for errors, anomalies, and threshold conditions. The following resources are monitored:
 - Host availability
 - Application status
 - IBM Workload Scheduler used space
 - IBM Workload Scheduler available disk space
 - Page-outs
 - Swap space available
 - IBM Workload Scheduler message file size
- An agent log file data source that monitors data in a sequential text file. IBM Workload Scheduler writes events for all monitored objects (for example, jobs and job streams), after each new plan production, to an event file that is accessed by the Z Monitoring agent for IBM Workload Scheduler. The collected data is retrieved and displayed by the Tivoli® Enterprise Portal.

Reporting

About this task

As part of the preproduction and post-production processes, reports are generated which show summary or detail information about the previous or next production day. These reports can also be generated ad-hoc. The following reports are available:

- Job details listing
- Prompt listing
- Calendar listing
- Parameter listing
- Resource listing
- Job History listing
- Job histogram
- Planned production schedule

- Planned production summary
- Planned production detail
- Actual production summary
- Actual production detail
- Cross reference report

In addition, during production, a standard list file (STDLIST) is created for each job instance launched by IBM Workload Scheduler. Standard list files contain header and trailer banners, echoed commands, and errors and warnings. These files can be used to troubleshoot problems in job execution.

Auditing

About this task

An auditing option helps track changes to the database and the plan.

For the database, all user modifications, except for the delta of the modifications, are logged. If an object is opened and saved, the action is logged even if no modification is made.

For the plan, all user modifications to the plan are logged. Actions are logged whether or not they are successful.

Audit files are logged to a flat text file on individual machines in the IBM Workload Scheduler network. This minimizes the risk of audit failure due to network issues and allows a straightforward approach to writing the log. The log formats are basically the same for both the plan and the database. The logs consist of a header portion which is the same for all records, an "action ID", and a section of data which varies according to the action type. All data is stored in clear text and formatted to be readable and editable from a text editor such as vi or notepad.

Using event-driven workload automation

About this task

Use this optional feature to set up and run rules that perform predefined actions in response to particular events occurring on your agents. Your organization can benefit from using this feature by adding on-demand workload automation to plan-based job scheduling, gaining savings in time and resources.

Event-driven workload automation is based on the concept of event rule. In IBM Workload Scheduler an event rule is a scheduling object and is made up of events, event-correlating conditions, and actions. When you define an event rule, you specify one or more events, a correlation rule, and one or more actions that are triggered by those events. Moreover, you can specify validity dates, a daily time interval of activity, and a common time zone for all the time restrictions that are set.

You can set up event rules to:

- Trigger the execution of batch jobs and job streams based on the occurrence or combination of real time events
- Reply to prompts

- Notify users when anomalous conditions occur in the IBM Workload Scheduler scheduling environment or batch scheduling activity
- Invoke an external product when a particular event condition occurs

IBM Workload Scheduler includes a set of predefined event and action plug-ins, but also provides a software development kit with samples and templates for your application programmers to develop their own plug-ins.

Options and security

The IBM Workload Scheduler options files determine how IBM Workload Scheduler runs on your system. Several performance, tuning, security, logging, and other configuration options are available.

Setting the IBM Workload Scheduler options

About this task

You can set two types of properties to configure your IBM Workload Scheduler run time environment, properties that are set on the master domain manager and affect processing on all workstations in the IBM Workload Scheduler network, and properties that are set locally on a workstation and affect processing on that workstation only. The former are managed using the IBM Workload Scheduler command line program named `optman`, and the latter you define locally on the workstation by customizing the configuration files `useropts`, `localopts`, and `jobmanrc`.

Global options are used to:

- Define if the security files of all the workstations of the network can be created and managed only from the master domain manager or if the root user or administrator of each workstation can create and manage their own.
- Select whether to enable or disable database auditing.
- Control which objects in the plan the user is permitted to list when running a query.
- Select whether to enable plan auditing.
- Select whether to enable strong encryption.
- Enter the number of days for which you want to save job statistics.
- Set the minimum and maximum lengths of the preproduction plan in days.
- Determine if uncompleted job streams are carried forward from the old to the new production control file.
- Remove or clean up job stream instances from the plan that ended in error.
- Define the start time of the IBM Workload Scheduler processing day.

Local options are used to:

- Specify the name of the local workstation
- Prevent the launching of jobs run by root in UNIX™
- Prevent unknown clients from connecting to the system
- Specify a number of performance options
- Specify a number of logging preferences
- Set SSL security options.

Setting security

About this task

Security is accomplished with the use of a security file that contains one or more user definitions. Each user definition identifies a set of users, the objects they are permitted to access, and the types of actions they can perform.

A template file is installed with the product. Edit the template to create the user definitions and compile and install it with a utility program to create a new operational security file. After it is installed, you make further modifications by creating an editable copy with another utility.

An individual file can be maintained on each workstation, or a single security file can be created on the master domain manager and copied to each domain manager, fault-tolerant agent, and standard agent.

Secure authentication and encryption

About this task

Security is enhanced for connections between protected and non-protected domains by applying the authentication and encryption mechanism based on the Secure Sockets Layer (SSL) protocol. SSL uses digital certificates to authenticate the identity of a workstation.

The IBM Workload Scheduler administrator must plan how authentication will be used within the network:

- Use one certificate for the entire IBM Workload Scheduler network.
- Use a separate certificate for each domain.
- Use a separate certificate for each workstation.

SSL support is automatically installed with IBM Workload Scheduler.

Work across firewalls

About this task

For previous versions of IBM Workload Scheduler, running the commands to start or stop a workstation or to get the standard list required opening a direct TCP/IP connection between the originator and the destination nodes. In a firewall environment, this forces users to break the firewall to open a direct communication path between the master and each fault-tolerant agent in the network.

IBM Workload Scheduler provides a configurable attribute, named *behindfirewall*, in the workstation's definition in the database. You can set this attribute to ON to indicate that a firewall exists between that particular workstation and its domain manager, and that the link between the domain manager and the workstation (which can be another domain manager) is the only allowed link between the domains.

Also, for all the workstations having this attribute set to ON, the commands to start or stop the workstation or to get the standard list will be transmitted through the domain hierarchy instead of opening a direct connection between the master (or domain manager) and the workstation.

Centralized security mechanism

About this task

A new global option makes it possible to change the security model in the IBM Workload Scheduler network. If you use this option, then the security files for the fault-tolerant agents in the network can be created or modified only on the master domain manager. The IBM Workload Scheduler administrator is responsible for creating, updating, and distributing the security files for all the agents where user access is required. Setting this global option triggers a security mechanism to identify and trust the IBM Workload Scheduler network corresponding to that master domain manager.

If you prefer the traditional security model, you can still use it by not activating the global variable.

Using time zones

About this task

IBM Workload Scheduler supports different time zones. Enabling time zones provides you with the ability to manage your workload across a multiple time zone environment. Both the 3-character and the variable length notations are supported with the current version of IBM Workload Scheduler. The variable length notation format is area/city, for example Europe/Paris as equivalent to ECT (European Central Time). The 3-character notation is supported for capability with earlier versions of the product.

Once configured, time zones can be specified for start and deadline times within jobs and job streams.

Using the access methods of IBM Workload Scheduler

Access methods are used to extend the job scheduling capability of IBM Workload Scheduler to other systems and applications. They run on:

- **Extended agents**
- **Dynamic agents** and **IBM Z Workload Scheduler agents**

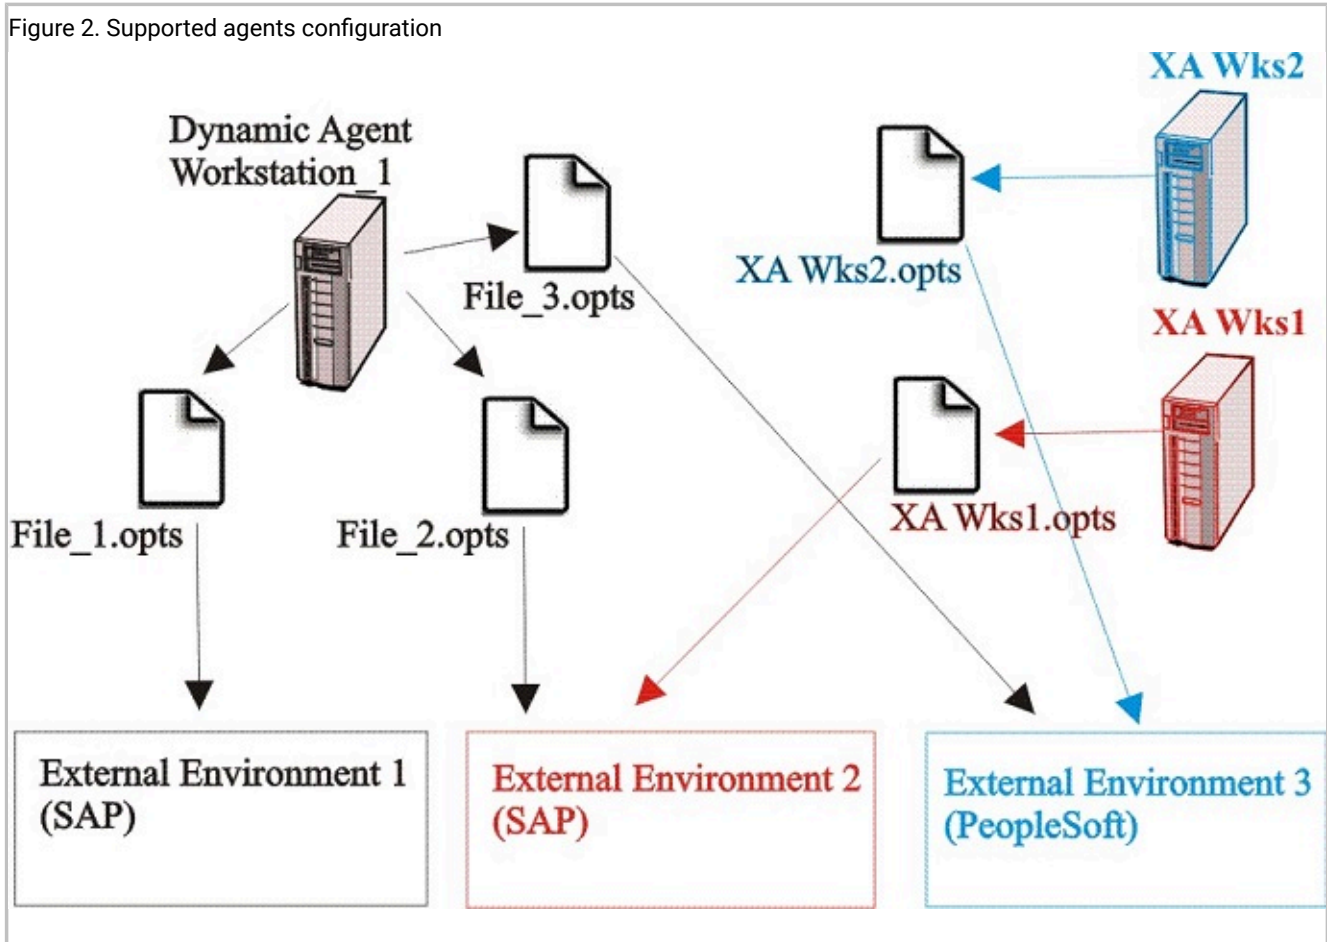
For more details, see [Supported agents on page 119](#).

An access method needs to interact with the external system either through its command line or the Dynamic Workload Console. IBM Workload Scheduler includes the following access methods:

- PeopleSoft (psagent)
- SAP R/3 (r3batch)

[Figure 2: Supported agents configuration on page 119](#) shows the typical configuration for the extended agent and the dynamic agent.

Figure 2. Supported agents configuration



Supported agents

The IBM Workload Scheduler access methods can run on extended agents and on dynamic agents (including IBM Z Workload Scheduler agents).

You can extend IBM Workload Scheduler capabilities for scheduling on one or more external applications, such as PeopleSoft and SAP R/3. The agents needed to extend scheduling capabilities on these external applications using IBM Workload Scheduler are:

Extended agents

Extended agents extend the static scheduling capability. They are logical workstations related to an access method hosted by a physical IBM Workload Scheduler workstation (not another extended agent). More than one extended agent workstation can be hosted by the same IBM Workload Scheduler workstation and rely on the same access method. The extended agent is defined in a standard IBM Workload Scheduler workstation definition, which gives the extended agent a name and identifies the access method. The access method is a program that is run by the hosting workstation whenever IBM Workload Scheduler submits a job to an external system.

Jobs are defined for an extended agent in the same manner as for other IBM Workload Scheduler workstations, except for the job attributes that depend on the external system or application.

To launch and monitor a job on an extended agent, the host runs the access method, passing it job details as command line options. The access method communicates with the external system to launch the job and returns the status of the job. To launch a job in an external environment, IBM Workload Scheduler runs the extended agent access method providing it with the extended agent workstation name and information about the job. The method looks at the corresponding file named `XANAME_accessmethod.opts` (where `XANAME` is the name of the extended agent workstation) to determine which external environment instance it will connect to. The access method can then launch jobs on that instance and monitor them through completion, writing job progress and status information in the standard list file of the job.

A physical workstation can host a maximum of 255 extended agents.

Extended agents can also be used to run jobs in an end-to-end environment, where their scheduling and monitoring tasks are performed by a IBM Z Workload Scheduler controller.

Dynamic agents and IBM Z Workload Scheduler agents (z-centric)

To run access methods on external applications using dynamic agents, you define a job of type **Access Method**. The access method communicates with the external system to launch the job and returns the status of the job. The method looks at the corresponding file named `optionsfile_accessmethod.opts` (where `optionsfile` is the configuration file that depends on the selected access method). The dynamic agent and the IBM Z Workload Scheduler agent can have more than one `.opts` associated file to determine which external environment instance to connect to. The access method can launch jobs on that instance and monitor them through completion, writing job progress and status information in the standard list file of the job.

Chapter 4. IBM Z Workload Scheduler

IBM Z Workload Scheduler expands the scope for automating your data processing (DP) operations. It plans and automatically schedules the production workload. From a single point of control, it drives and controls the workload processing at both local and remote sites. By using IBM Z Workload Scheduler to increase automation, you use your DP resources more efficiently, have more control over your DP assets, and manage your production workload processing better.

How your production workload is managed

How does IBM Z Workload Scheduler give you all this? This section describes functions that make your information systems (IS) operations run more efficiently. But first, here is a brief introduction to the structure of the product and some concepts.

Structure

About this task

IBM Z Workload Scheduler consists of a base product, the *agent* and a number of features. Every z/OS® system in your complex requires the base product. One z/OS® system in your complex is designated the *controlling* system and runs the *engine* feature. Only one engine feature is required, even when you want to start standby engines on other z/OS® systems in a sysplex.

IBM Z Workload Scheduler with IBM Workload Scheduler addresses your production workload in the distributed environment. You can schedule, control, and monitor jobs in IBM Workload Scheduler from IBM Z Workload Scheduler. For example, in the current plan, you can specify jobs to run on workstations in IBM Workload Scheduler.

The workload on other operating environments can also be controlled with the open interfaces provided with IBM Z Workload Scheduler. Sample programs using TCP/IP or an NJE/RSCS (network job entry/remote spooling communication subsystem) combination show you how you can control the workload on environments that at present have no scheduling feature.

Additionally, national language features let you see the dialogs and messages, in the language of your choice. These languages are currently available:

- English
- German
- Japanese
- Korean
- Spanish

Panel and message text can also be modified to include enterprise-specific instructions or help.

Concepts

About this task

Terminology definition. Throughout this chapter the following synonyms are used alternately to express the same concepts:

Concept	Workload scheduling term	IBM Z Workload Scheduler term
A list of jobs and of the information used to run them. A scheduling object defined in the product database.	job stream	application
The specific job stream (application) when it is entered in the IBM Z Workload Scheduler plan. It is defined also by a specific execution date and time.	job stream instance	occurrence
A task, a command, or a number of tasks.	job	operation
A job in the plan.	job instance	operation

In managing production workloads, IBM Z Workload Scheduler builds on several important concepts.

Plans. IBM Z Workload Scheduler constructs operating *plans* based on user-supplied descriptions of the DP operations department and its production workload. These plans provide the basis for your service level agreements and give you a picture of the status of the production workload at any point in time.

You can simulate the effects of changes to your production workload, calendar, and installation by generating trial plans.

Applications.

Also known as *job stream*, an application is a description of a unit of production work. It can include the following:

- A list of the operations, also known as *jobs*: they are the tasks associated with that unit of work, such as:
 - Data entry
 - Job preparation
 - Job submission or started-task initiation
 - Communication with the NetView® program
 - File transfer to other operating environments
 - Printing of output
 - Postprocessing activities, such as quality control or dispatch
 - Other tasks related to the unit of work that you want to schedule, control, and track
- A description of dependencies between jobs within a job stream and between jobs in other job streams
- Information about resource requirements, such as exclusive use of a data set
- Special operator instructions that are associated with a job

- How and where each job should be processed
- Run policies for that unit of work; that is, when it should be scheduled or alternatively the name of a group definition that records the run policy

IBM Z Workload Scheduler schedules work based on the information you provide in your application (job stream) descriptions.

Workstations.

When scheduling and processing work, IBM Z Workload Scheduler considers the processing requirements of each job. Some typical processing considerations are:

- Which human or machine resources are required for processing the work, for example, operators, processors, or printers?
- When are these resources available?
- How are these jobs to be tracked?
- Can this work be processed somewhere else if the resources become unavailable?

IBM Z Workload Scheduler supports a range of work process types, called *workstations*, that map the processing requirements of any task in your production workload. Each workstation supports one type of activity. This gives you the flexibility to schedule, monitor, and control any type of DP activity, including the following:

- Job setup, both manual and automatic
- Job submission
- Started-task actions
- Communication with the NetView® program
- Print jobs
- Manual preprocessing or postprocessing activity

You can plan for maintenance windows in your hardware and software environments. IBM Z Workload Scheduler helps you perform a controlled and incident-free shut down of the environment, preventing last-minute cancellation of active tasks. You can choose to reroute the workload automatically during any outage, planned or unplanned.

IBM Z Workload Scheduler tracks jobs as they are processed at workstations and dynamically updates the plan with real-time information on the status of jobs. You can view or modify this status information online using the workstation ready lists in the dialog.

Virtual Workstations.

Using virtual workstations improves workload balancing and the monitoring of system availability. This feature automatically directs the submission of workload to different destinations removing the need to associate a workstation to a specific

destination. You can define a list of destinations for the submission of workload and the scheduler distributes the workload to automatically-selected active destinations, according to a round-robin scheduling approach.

You can activate this feature by specifying the new virtual option at workstation definition level. This option is allowed for computer workstations with the automatic reporting attribute, and is supported by all the interfaces available to define, modify, and monitor workstations.

Using virtual workstations the scheduler distributes the workload across your trackers evenly, thus avoiding bottlenecks when submitting or running jobs. In fact, the scheduler splits the workload among the available destinations, so that the Job Entry System (JES) and Workload Manager (WLM) do not find overloaded input queues when selecting jobs for their action.

Dependencies.

In general, every DP-related activity must occur in a specific order. Activities performed out of order might create invalid output and possibly even corrupt your corporate data. This might cause costly reruns, missed deadlines, and unsatisfied customers.

You can define *dependencies* for operations (jobs) when a specific processing order is required. When IBM Z Workload Scheduler manages the dependent relationships for you, the jobs are always started in the correct order every time they are scheduled. A dependency is called *internal* when it is between two jobs in the same job stream, and *external* when it is between two jobs in different job streams. A dependency takes place between a *predecessor* operation and a *successor* operation, whereby the successor can start after its predecessor has completed.

The resolution of a dependency is resolved based on the input arrival times - theoretical start times that help define a specific application occurrence - of the applications of which the predecessors and successors are part. A dependency is resolved when a best matching predecessor is found according to the criteria defined for that dependency. The criteria can be that the best matching predecessor can be found within the closest preceding occurrence with respect to the successor, or in one that runs in the same day or within a specific interval of days or hours. Also, a dependency can be defined as mandatory to various degrees.

In addition, you can specify *conditional dependencies*, where you use the return code and status of an operation to determine the start of another operation. Standard logical operators are supported to define the check on status or return code values, to implement the definition of dependencies with a *conditional logic*. If the predecessor operation is associated to a job with different steps, you can specify a conditional *step-level dependency* on individual step return codes.

IBM Z Workload Scheduler lets you serialize work based on the status of any DP resource. A typical example is a job that uses a data set as input, but must not start until the data set is successfully created and loaded with valid data. You can use *resource serialization* support to send availability information about a DP resource to IBM Z Workload Scheduler.

Special resources.

Special resources are typically defined to represent physical or logical objects used by jobs. A special resource can be used to serialize access to a data set or to limit the number of file transfers on a particular network link. The resource does not have to represent a physical object in your configuration, although it often does.

IBM Z Workload Scheduler keeps a record of the state of each resource and its current allocation status. You can choose to hold resources if a job allocating the resources ends abnormally. You can also use the IBM Z Workload Scheduler interface with the Resource Object Data Manager (RODM) to schedule jobs according to *real* resource availability. You can subscribe to RODM updates in both local and remote domains.

IBM Z Workload Scheduler lets you *subscribe* to data set activity on z/OS® systems. The data set triggering function of IBM Z Workload Scheduler automatically updates special resource availability when a data set is closed. You can use this notification to coordinate planned activities or to add unplanned work to the schedule.

Calendars.

IBM Z Workload Scheduler uses information about when the job departments work, so that job streams are not scheduled to run on days when processing resources are not available (for example, Sundays and holidays). This information is stored in a *calendar*. IBM Z Workload Scheduler supports *multiple calendars* for enterprises where different departments have different work days and non-working days. Different groups within a business operate according to different calendars.

The multiple calendar function is critical if your enterprise has installations in more than one geographical location (for example, with different local or national holidays).

Business processing cycles. IBM Z Workload Scheduler uses business processing cycles, or *periods*, to calculate when your job streams run, for example, weekly, or every 10th working day. Periods are based on the business cycles of your customers. IBM Z Workload Scheduler supports a range of periods for processing the different job streams in your production workload.

When you define a job stream, you specify when it is planned to run, using a *run cycle*, which can be:

- A *rule* with a format such as

```
ONLY the SECOND TUESDAY of every MONTH
```

```
EVERY FRIDAY in the user-defined period SEMESTER1
```

where the words in upper case are selected from lists of ordinal numbers, names of days, and common calendar intervals or period names.

- A combination of period and *offset*. For example, an offset of 10 in a monthly period specifies the 10th day of each month.

Run cycle groups. One of the elements that make up the definition of an application is the run cycle, where you specify the temporal details of when the application should run in terms of time, days, weeks, months, or periods (and several other details). An application can have several definitions of run cycles, which become part of the definition of that application. There are several types of run cycles, such as regular, exclusive, rule-based, or period.

Going one step further, you can define run cycle groups. These are database objects in their own right, and are not part of the definition of an application, but the same run cycle group can be used by more than one application. A run cycle group is a list of run cycles that, combined together, produce a set of run dates.

You can structure a run cycle group into subsets. Within a subset you can match an exclusive run cycle against a positive one to generate negative occurrences, which identify the days when an application is normally scheduled to run but is bypassed.

You can make use of the logical AND between two run cycles in a group. This enables you to easily define rules that schedule work on complex run dates.

Using Plans in IBM Z Workload Scheduler

IBM Z Workload Scheduler plans your production workload schedule. It produces both high-level and detailed plans. These plans both drive the production workload and show you the status of the production workload on your system at any specified time. You can produce trial plans to forecast future workloads.

Long-term planning

About this task

The *long-term plan* is a high-level schedule of your anticipated production workload. It lists, by day, the instances of job streams to be run during the period of the plan. Each instance of a job stream is called an *occurrence*. The long-term plan shows when occurrences are to run, as well as the dependencies that exist between the job streams. You can view these dependencies graphically on your workstation as a network, to check that work has been defined correctly. The plan can help you in forecasting and planning for heavy processing days. The long-term-planning function can also produce histograms showing planned resource use for individual workstations during the plan period.

You can use the long-term plan as the basis for documenting your service level agreements. It lets you relate service level agreements directly to your production workload schedules so that your customers can see when and how their work is to be processed.

The long-term plan provides a window to the future. You can decide how far into the future, from one day to four years. You can also produce long-term plan simulation reports for *any* future date. IBM Z Workload Scheduler can automatically extend the long-term plan at regular intervals. You can print the long-term plan as a report, or you can view, alter, and extend it online using the dialogs.

Detailed planning

About this task

The *current plan* is the center of IBM Z Workload Scheduler processing. It drives the production workload automatically and provides a way to check its status. The current plan is produced by the run of batch jobs that extract from the long-term plan the occurrences that fall within the specified period of time from the job details. The current plan selects a window from the long-term plan and makes the jobs ready to be run. They are started depending on the decided restrictions (for example, dependencies, resources availability, or time-dependent jobs).

The current plan is a rolling plan that can cover several days. A common method is to cover 1 to 2 days with regular extensions each shift. Production workload processing activities are listed by minute.

You can either print the current plan as a report, or view, alter, and extend it online, by using the dialogs.

Automatically controlling the production workload

About this task

IBM Z Workload Scheduler automatically drives the production workload by monitoring the flow of work and by directing the processing of jobs to follow the business priorities established in the plan.

Through its interface to the NetView® program or its management-by-exception ISPF dialog, IBM Z Workload Scheduler can alert

the production control specialist to problems in the production workload processing. Furthermore, the NetView® program can automatically trigger IBM Z Workload Scheduler to perform corrective actions in response to these problems.

IBM Z Workload Scheduler automatically:

- Starts and stops started tasks
- Edits job statements: z/OS® JCL or equivalent job statements for other operating environments before submission
- Submits jobs in the specified sequence to the target operating environment every time
- Tracks each scheduled job in the plan
- Determines the success or failure of the jobs
- Displays status information and instructions to guide workstation operators
- Provides automatic recovery of jobs when they end in error, regardless of the operating environment
- Generates processing dates for your job stream run cycles using rules, such as:
 - Every second Tuesday of the month
 - Only the last Saturday in June, July, and August
 - Every third workday in the user-defined PAYROLL period
- Starts jobs with regard to real resource availability
- Performs data set cleanup in error and rerun situations for the z/OS® workload
- Tailors the JCL for step restarts of z/OS® jobs and started tasks
- Dynamically schedules additional processing in response to unplannable activities

- Provides automatic notification when an updated data set is closed; this can be used to trigger subsequent processing
- Generates alerts when abnormal situations are detected in the workload

IBM Z Workload Scheduler also provides manual control facilities, which are described in [Manual control and intervention on page 133](#).

Automatic workload submission

About this task

IBM Z Workload Scheduler automatically drives work through the system, taking into account work that requires manual or program-recorded completion. Program-recorded completion refers to situations where the status of a scheduler-controlled job is set to “complete” by a user-written program. It also promotes the optimum use of resources, improves system availability, and automates complex and repetitive operator tasks. IBM Z Workload Scheduler automatically controls the submission of work according to:

- Dependencies between jobs
- Workload priorities
- Specified time for the submission of particular work
- Availability of resources

By saving a copy of the JCL for each separate run, or occurrence, of a particular job in its plans, IBM Z Workload Scheduler prevents the unintentional reuse of temporary JCL changes, such as overrides.

Job tailoring. IBM Z Workload Scheduler provides automatic job tailoring functions to automatically edit jobs. This can reduce your dependency on time-consuming and error-prone manual editing of jobs. IBM Z Workload Scheduler job tailoring provides:

- Automatic variable substitution
- Dynamic inclusion and exclusion of inline job statements
- Dynamic inclusion of job statements from other libraries or from an exit

For jobs submitted on a z/OS® system, these job statements are z/OS® JCL. Scheduler JCL tailoring directives can be included in jobs that are submitted on other operating systems, such as AIX®/6000.

Variables can be substituted in specific columns, and you can define verification criteria to ensure that invalid strings are not substituted. Special directives supporting a variety of date formats used by job stream programs let you dynamically define the required format and change the multiple times for the same job. You can define arithmetic expressions to calculate values such as the current date plus four work days. And you can set a temporary variable to a specific value or to an expression composed of other temporary variables.

System Automation commands tailoring. IBM Z Workload Scheduler provides a function that edits system automation commands automatically. This helps you to save time and reduce the possibility of editing errors. IBM Z Workload Scheduler command tailoring provides automatic variable substitution.

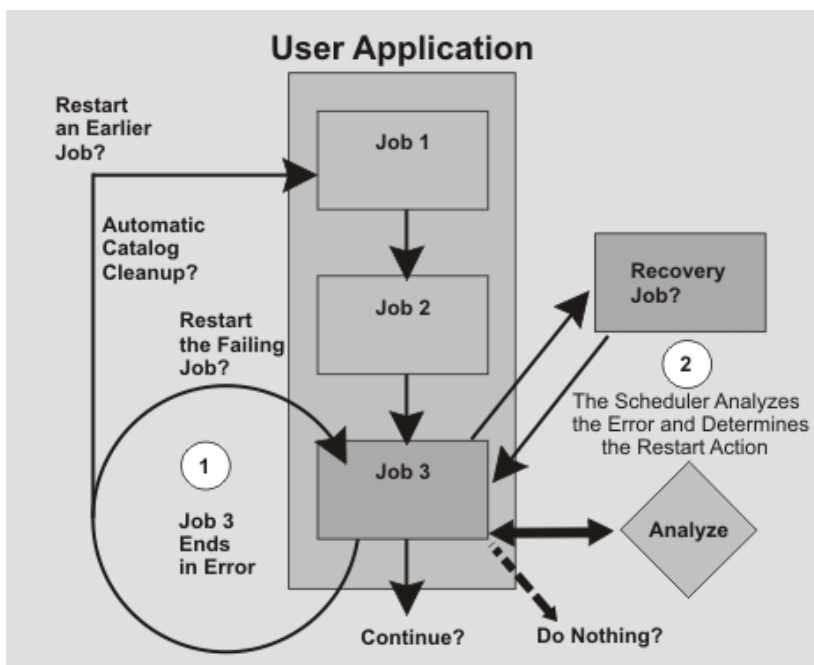
Automatic recovery and restart

IBM Z Workload Scheduler provides automatic restart facilities for your production work. You can specify the restart actions to take if work initiated by IBM Z Workload Scheduler ends in error (see [Figure 3: Automatic recovery and restart on page 129](#).) You can use these functions to predefine automatic error recovery and restart actions for jobs and started tasks. The scheduler's integration with the NetView® program allows it to automatically pass alerts to the NetView® program in error situations. Using the z/OS® cross-system coupling facility (XCF) enables IBM Z Workload Scheduler to maintain production workload processing when system failures occur.

Recovery of jobs and started tasks. Automatic recovery actions for failed jobs are specified in user-defined control statements. Parameters in these statements determine the recovery actions to be taken when a job or started task ends in error.

Example

Figure 3. Automatic recovery and restart



Restart and cleanup. You can use restart and cleanup to catalog, uncatalog, or delete data sets when a job ends in error, or when you need to rerun a job. Data set cleanup handles JCL in the form of in-stream JCL, in-stream procedures, and cataloged procedures on both local and remote systems. This function can be initiated automatically by IBM Z Workload Scheduler or manually by using the panels. IBM Z Workload Scheduler resets the catalog to the status that it was before the job ran for both generation data set groups (GDGs) and for DD allocated data sets contained in JCL. In addition, restart and cleanup supports the use of Removable Media Manager in your environment.

Restart at both the step- and job-level is also provided in the IBM Z Workload Scheduler panels. It manages resolution of generation data group (GDG) names, JCL containing nested INCLUDEs or PROC, and IF-THEN-ELSE statements. IBM Z Workload Scheduler also automatically identifies problems that can prevent successful restart, providing a logic of the "best restart step."

You can browse the job log or request a step-level restart for any z/OS® job or started task even when there are no catalog modifications. The job-log browse functions are also available for the workload on other operating platforms, which is especially useful for those environments that do not support an SDSF-like facility. If you use a SYSOUT archiver, for example RMDS, you can interface with it from IBM Z Workload Scheduler and so prevent duplication of job log information.

These facilities are available to you without the need to make changes to your current JCL.

IBM Z Workload Scheduler gives you an enterprise-wide data set cleanup capability on remote agent systems.

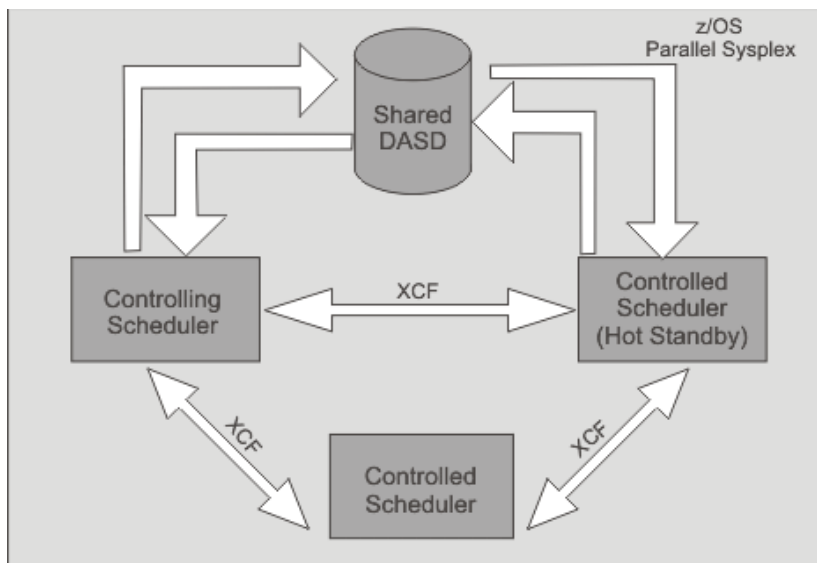
Production workload restart. IBM Z Workload Scheduler provides a production workload restart, which can automatically maintain the processing of your work if a system or connection fails. Scheduler-controlled production work for the unsuccessful system is rerouted to another system. Because IBM Z Workload Scheduler can restart and manage the production workload, the integrity of your processing schedule is maintained, and service continues for your customers.

IBM Z Workload Scheduler uses the VTAM® Model Application Program Definition feature and the z/OS® defined symbols to ease the configuration and job in a sysplex environment, giving you a single system view of the sysplex.

Starting, stopping, and managing your engines and agents do not require you to know on which sysplex z/OS® image they are actually running on.

Example

Figure 4. Production workload restart and hot standby



Hot standby. IBM Z Workload Scheduler provides a single point of control for your z/OS® production workload. If this controlling system fails, IBM Z Workload Scheduler can automatically transfer the controlling functions to a backup system within a Parallel Sysplex®, see [Figure 4: Production workload restart and hot standby on page 130](#). Through XCF, IBM Z Workload Scheduler can automatically maintain production workload processing during system or connection failures.

z/OS® automatic restart manager support

About this task

All the scheduler components are enabled to be restarted by the Automatic Restart Manager (ARM) of the z/OS® operating system, in the case of program failure.

Workload Manager (WLM) support

About this task

With Workload Manager (WLM), you can make the best use of resources accessed by your scheduled jobs. In addition, your jobs maintain the highest possible throughput with WLM and IBM Z Workload Scheduler. When used with WLM, the scheduler can achieve the best possible system response times.

Automatic status checking

About this task

To track the work flow, IBM Z Workload Scheduler interfaces directly with the operating system, collecting and analyzing status information about the production work that is currently active in the system. IBM Z Workload Scheduler can record status information from both local and remote processors. When status information is reported from remote sites in different time zones, IBM Z Workload Scheduler makes allowances for the time differences.

Status reporting from heterogeneous environments

About this task

The processing on other operating environments can also be tracked by IBM Z Workload Scheduler. You can use supplied programs to communicate with the engine from any environment that can establish communications with a z/OS® system.

Status reporting from user programs

About this task

You can pass status information about production workload processing to IBM Z Workload Scheduler from your own user programs through a standard supplied routine.

Additional job-completion checking

About this task

If required, IBM Z Workload Scheduler provides further status checking by scanning SYSOUT and other print data sets from your processing when the success or failure of the processing cannot be determined by completion codes. For example, IBM Z Workload Scheduler can check the text of system messages or messages originating from your user programs. Using information contained in job completion checker (JCC) tables, IBM Z Workload Scheduler determines what actions to take when it finds certain text strings. These actions can include:

- Reporting errors
- Requeuing SYSOUT
- Writing incident records to an incident data set

Managing unplanned work

About this task

IBM Z Workload Scheduler can be automatically triggered to update the current plan with information about work that cannot be planned in advance. This allows IBM Z Workload Scheduler to control unexpected work. Because IBM Z Workload Scheduler checks the processing status of this work, automatic recovery facilities are also available.

Integration with IBM Tivoli Monitoring

IBM Z Monitoring monitors and manages system and network applications on a variety of platforms and can keep track of the availability and performance of your enterprise in its entirety. You can use the reports provided by IBM Z Monitoring to track trends and troubleshoot problems.

The integration between IBM Z Workload Scheduler and IBM Z Monitoring requires the installation of a customizable agent, the IBM Z Monitoring Agent for IBM Z Workload Scheduler. This agent is installed on the systems that you want to monitor, and it collects and distributes data to the Tivoli® Enterprise Monitoring Server.

The IBM Z Monitoring Agent for IBM Z Workload Scheduler (called Z Monitoring agent from now on) sends events and alerts to IBM Z Monitoring (ITM). IBM Z Workload Scheduler establishes an IP connection with the Socket Data Source of the Z Monitoring agent which acts like a listener task. Every time one of the monitored operations changes its status, or an alert is issued, the related event is sent to the Z Monitoring agent.

The data sent by IBM Z Workload Scheduler is interpreted by the Z Monitoring agent. The agent stores the event information in its cache and this information is later consolidated by the Tivoli® Enterprise Monitoring Server. The collected data is then retrieved and displayed by a component named Tivoli® Enterprise Portal.

The Tivoli® Enterprise Portal is an interactive interface used for viewing and monitoring your enterprise network. The Tivoli® Enterprise Portal client connects to the Tivoli® Enterprise Portal server that enables retrieval, manipulation, and analysis of data collected by all existing IBM Z Monitoring agents in your enterprise.

You can use the Tivoli® Enterprise Portal to display and manage events related to IBM Z Workload Scheduler operations such as operation start and end time. In the same way, you can configure the automatic selection of monitored operations and subtasks, providing a single point of management and control for the resources of your enterprise.

Interfacing with other programs

About this task

IBM Z Workload Scheduler provides a program interface (PIF). Using this interface, you can automate most actions that you can perform online through the dialogs. This interface can be called from CLISTS, user programs, and using TSO commands.

The application programming interface (API) lets your programs communicate with IBM Z Workload Scheduler from any compliant platform. You can use Common Programming Interface for Communications (CPI-C), advanced program-to-program communication (APPC), or your own logical unit (LU) 6.2 verbs to converse with IBM Z Workload Scheduler through the API. You can use this interface to query and update the current plan. The programs can be running on any platform that is connected locally, or remotely through a network, with the z/OS® system where the engine runs.

Manual control and intervention

About this task

IBM Z Workload Scheduler lets you check the status of work and intervene manually when priorities change or when you want to run unplanned work. You can query the status of the production workload and then modify the schedule if needed.

Status inquiries

About this task

With the ISPF dialogs or the Dynamic Workload Console, you can make queries online and receive timely information on the status of the production workload.

Time information that is displayed by the dialogs is in the local time of the dialog user. Using the dialogs, you can request detailed or summary information on individual job streams, jobs, and workstations, as well as summary information concerning workload production as a whole. You can also display dependencies graphically as a network at both job stream and job level. Status inquiries:

- Provide you with overall status information that you can use when considering a change in workstation capacity or when arranging an extra shift or overtime work.
- Help you supervise the work flow through the installation; for example, by displaying the status of work at each workstation.
- Help you decide whether intervention is required to speed up the processing of specific job streams. You can find out which job streams are the most critical. You can also check the status of any job stream, as well as the plans and actual times for each job.
- Help you to check information before making modifications to the plan. For example, you can check the status of a job stream and its dependencies before deleting it or changing its input arrival time or deadline. See [Modifying the current plan on page 134](#) for more information.
- Provide you with information on the status of processing at a particular workstation. Perhaps work that should have arrived at the workstation has not arrived. Status inquiries can help you locate the work and find out what has happened to it.

Modifying the current plan

About this task

IBM Z Workload Scheduler makes status updates to the plan automatically, using its tracking functions. However, it lets you change the plan manually to reflect unplanned changes to the workload or to the operations environment, which often occur during a shift. For example, you might want to change the priority of a job stream, add unplanned work, or reroute work from one workstation to another. You might also want to correct operational errors manually. Modifying the current plan might be the best way to handle these situations.

You can modify the current plan online. For example, you can:

- Include unexpected jobs or last-minute changes to the plan. IBM Z Workload Scheduler then automatically creates the dependencies for this work.
- Manually modify the status of jobs.
- Delete occurrences of job streams.
- Graphically display job dependencies before you modify them.
- Modify the data in job streams, including the JCL.
- Respond to error situations by:
 - Rerouting jobs
 - Rerunning jobs or occurrences
 - Completing jobs or occurrences
 - Changing jobs or occurrences

- Change the status of workstations by:
 - Rerouting work from one workstation to another
 - Modifying workstation reporting attributes
 - Updating the availability of resources
 - Changing the way resources are handled
- Replan or extend the current plan

In addition to using the dialogs, you can modify the current plan from your own job streams using the program interface or the application programming interface. You can also trigger IBM Z Workload Scheduler to dynamically modify the plan using TSO commands or a batch program. This adds unexpected work automatically to the plan.

Management of critical jobs

About this task

IBM Z Workload Scheduler uses the capability of the Workload Manager component of z/OS® to ensure that critical jobs are completed on time. If a critical job is late, IBM Z Workload Scheduler favors it using the Workload Manager interface.

Management of critical path

About this task

In addition to the handling of critical jobs based on Workload Manager, IBM Z Workload Scheduler provides the dynamic handling of the critical path calculated by the daily planning batch jobs process.

The critical path is the path, within a network of jobs, with the least slack time.

The slack time, in a critical job predecessor path, is the amount of time that processing of the predecessor jobs can be delayed without exceeding the deadline of a critical job. It is the spare time calculated using the deadline, input arrival, and duration settings of predecessor jobs.

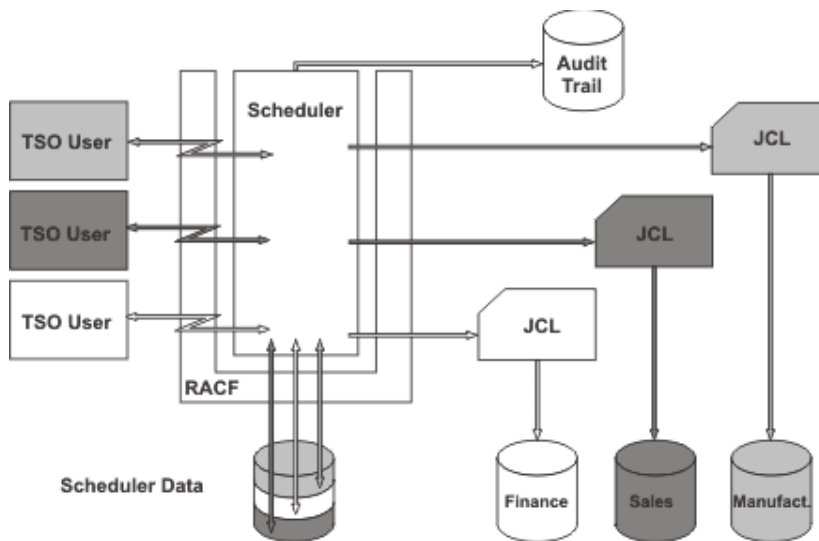
The capabilities include:

- Monitoring of critical job predecessors that are late, long running, or ended with an error. This process uses the same internal logic that the scheduler applies to monitor alert conditions.
- Monitoring of the paths that are consuming their slack time, becoming more critical than the paths calculated at plan generation.
- Enhanced critical jobs monitoring, using ISPF dialog flows.
- Back-end support for new views available using the Dynamic Workload Console.

Security

Today, DP operations increasingly require a high level of data security, particularly as the scope of DP operations expands and more people within the enterprise become involved. IBM Z Workload Scheduler provides complete security and data integrity within the range of its functions. It provides a shared central service to different user departments even when the users are in different companies and countries. IBM Z Workload Scheduler provides a high level of security to protect scheduler data and resources from unauthorized access. With IBM Z Workload Scheduler, you can easily organize, isolate, and protect user data to safeguard the integrity of your user applications (see [Figure 5: Security on page 136](#)). IBM Z Workload Scheduler can plan and control the work of many user groups, and maintain complete control of access to data and services.

Figure 5. Security



Audit trail

With the audit trail, you can define how you want IBM Z Workload Scheduler to log accesses (both reads and updates) to scheduler resources. Because it provides a history of changes to the databases, the audit trail can be extremely useful for staff that work with debugging and problem determination.

A sample program is provided for reading audit-trail records. The program reads the logs for a period that you specify and produces a report detailing changes that have been made to scheduler resources.

System authorization facility

About this task

IBM Z Workload Scheduler uses the system authorization facility (SAF), a function of z/OS®, to pass authorization verification requests to your security system, for example RACF®. This means that you can protect your scheduler data objects with any security system that uses the SAF interface.

Protection of data and resources

About this task

Each user request to access a function or to access data is validated by SAF. This is some of the information that can be protected:

- Calendars and periods
- Job stream names or job stream owner, by name
- Workstation, by name
- Job stream-specific data in the plan
- Operator instructions
- JCL

To support distributed, multi-user handling, IBM Z Workload Scheduler lets you control the level of security you want to implement, right down to the level of individual records. You can define generic or specific RACF® resource names to extend the level of security checking.

If you have RACF® Version 2 Release 1 installed, you can use the IBM Z Workload Scheduler reserved resource class to manage your IBM Z Workload Scheduler security environment. This means that you do not have to define your own resource class by modifying RACF® and restarting your system.

Data integrity during submission

About this task

IBM Z Workload Scheduler can ensure the correct security environment for each job it submits, regardless of whether the job is run on a local or a remote system. IBM Z Workload Scheduler lets you create tailored security profiles for individual jobs or groups of jobs.

Compliance with the Federal Information Processing Standards (FIPS)

You can configure IBM Z Workload Scheduler to comply with Federal Information Processing Standards (FIPS) standard over SSL secured connections.

You can configure IBM Z Workload Scheduler to activate FIPS 140-2 compliance when scheduling end-to-end with z-centric agents or with fault-tolerant agents, and in communications between the controller and the following interfaces: tracker, server, datastore, remote ISPF dialog, program interface (PIF). To do so, you must configure the parameter named ENABLEFIPS in the TCPOPTS and/or HTTPPTS initialization statements of the IBM Z Workload Scheduler controller.

Configurations of IBM Z Workload Scheduler

About this task

IBM Z Workload Scheduler supports many configuration options using a variety of communication methods:

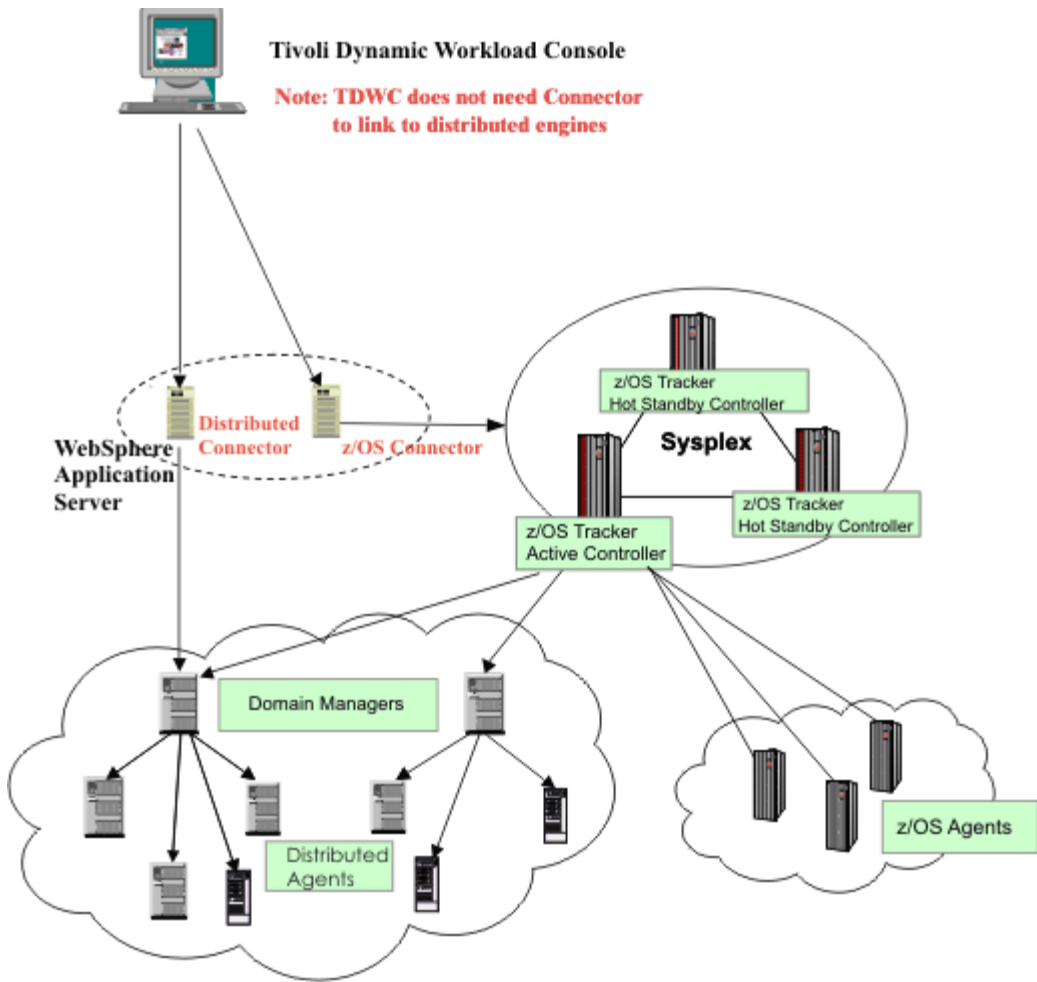
- [The controlling system on page 137](#)
- [Controlled z/OS systems on page 138](#)
- [Remote panels and program interface applications on page 139](#)
- [Scheduling jobs that are in IBM Workload Scheduler on page 139](#)

The controlling system

The controlling system requires both the agent and the engine. One controlling system can manage the production workload across all your operating environments.

The engine is the focal point of control and information. It contains the controlling functions, the dialogs, and the scheduler's own batch programs. Only one engine is required to control the entire installation, including local and remote systems (see [Figure 6: IBM Z Workload Scheduler configurations on page 138](#)).

Figure 6. IBM Z Workload Scheduler configurations



Controlled z/OS® systems

About this task

An agent is required for every controlled z/OS® system in a configuration. This includes, for example, local controlled systems within shared DASD or sysplex configurations.

The agent runs as a z/OS® subsystem and interfaces with the operating system (through JES2 or JES3, and SMF), using the subsystem interface and the operating system exits. The agent monitors and logs the status of work, and passes the status information to the engine via shared DASD, XCF, or ACF/VTAM®.

You can use z/OS® and cross-system coupling facility (XCF) to connect your local z/OS® systems. Instead of being passed to the controlling system using shared DASD, work status information is passed directly through XCF connections. XCF lets you use all the production-workload-restart facilities and its hot standby function. See [Automatic recovery and restart on page 129](#).

Remote systems

About this task

The agent on a remote z/OS® system passes status information about the production work in progress to the engine on the controlling system. All communication between IBM Z Workload Scheduler subsystems on the controlling and remote systems is done through ACF/VTAM.

IBM Z Workload Scheduler lets you link remote systems using ACF/VTAM networks. Remote systems are frequently used locally “*on-premises*” to reduce the complexity of the data processing (DP) installation.

Remote panels and program interface applications

About this task

ISPF panels and program interface (PIF) applications can run in a different z/OS® system from the one where the engine is running. Dialogs and PIF applications send requests to and receive data from an IBM Z Workload Scheduler server that is running on the same z/OS® system where the target engine is running, using advanced program-to-program communications (APPC). The server communicates with the engine to perform the requested actions.

The server is a separate address space, started and stopped either automatically by the engine or by the user through the z/OS® start command. There can be more than one server for an engine.

If the dialogs or the PIF applications run on the same z/OS® system where the target engine is running, the server might not be involved.

Scheduling jobs that are in IBM Workload Scheduler

About this task

IBM Z Workload Scheduler also allows you to access job streams (schedules in IBM Workload Scheduler) and add them to the current plan in IBM Z Workload Scheduler. In addition, you can build dependencies among IBM Z Workload Scheduler job streams and IBM Workload Scheduler jobs. From IBM Z Workload Scheduler, you can monitor and control the distributed agent.

Using fault-tolerant workstations

In the IBM Z Workload Scheduler current plan, you can specify jobs to run on fault-tolerant agents in IBM Workload Scheduler. IBM Z Workload Scheduler passes the job information to the IBM Workload Scheduler Symphony™ file, which in turn passes the jobs in the current plan to IBM Workload Scheduler to distribute and process. In turn, IBM Workload Scheduler reports the status of running and completed jobs back to the current plan for monitoring in IBM Z Workload Scheduler.

Using z-centric workstations

z-centric workstations are agents that are installed in a IBM Workload Scheduler network and that can be connected to IBM Z Workload Scheduler by HTTP or HTTPS. They provide the means to schedule from IBM Z Workload Scheduler jobs that need to run on distributed platforms (UNIX™, Linux™, Windows™). They are equivalent to computer automatic workstations in IBM Z Workload Scheduler and require less configuration and a smaller supporting infrastructure than fault-tolerant workstations.

Chapter 5. Dynamic Workload Console

About this task

The Dynamic Workload Console is a Web-based user interface for:

- IBM Workload Scheduler
- IBM Z Workload Scheduler

It is the strategic user interface for the IBM Workload Automation suite of products and includes support for the latest functions and enhancements available with the scheduling engines.

The Dynamic Workload Console is a light, powerful and user-friendly single point of operational control for the entire scheduling network. It allows for single sign-on and authentication to one or many schedulers, is highly scalable, and provides real-time monitoring, management and reporting of enterprise workloads. It also greatly simplifies report creation and customization.

With Dynamic Workload Console you can:

- Manage your workload to design objects in the database, handle plans, submit jobs or job streams, and monitor objects in the plan.
- Design and control the topology of your scheduling environment, that is workstations and domains.
- Define and run reports to gather historical data or details about your plans. You can also generate and run customized SQL reports.
- Define and manage logical resources or groups of logical resources for use with dynamic scheduling.

You can access the Dynamic Workload Console from any computer in your environment using a web browser through both secure HTTPS or HTTP protocol.

The first and main actions you perform when you connect to the Dynamic Workload Console are:

Creating a connection to a scheduling engine (IBM Workload Scheduler or IBM Z Workload Scheduler)

You type the details (such as IP address, user name, and password) to access a scheduling engine, and, optionally, a database to operate with objects defined in plans or stored in the database. You can also define new scheduling objects in the database.

From the Dynamic Workload Console you can access the current plan, a trial plan, a forecast plan, or an archived plan for the distributed environment or the current plan for the z/OS® environment.

You might want to access the database to perform actions against objects stored in it or generate reports showing historical or statistical data.

In addition, working both on the database and on plans, you can create and run event rules to define and trigger actions that you want to run in response to events occurring on IBM Workload Scheduler nodes.

Specifying filtering criteria to manage scheduling objects in the plan

You specify some filtering criteria to query a list of scheduling objects whose attributes satisfy the criteria you specified. Starting from this list, you can navigate and modify the content of the plan, switching between objects, opening more lists and accessing other plans or other IBM Workload Scheduler or IBM Z Workload Scheduler environments.

The console provides also the following graphical views tools to manage your workload:

Graphical view (for modeling)

A graphical extension to the Workload Designer that shows graphical representations of job stream definitions in the database. It provides an intuitive way to create and maintain them.

Plan view (for monitoring)

A high-level representation of a plan of any type, showing a filtered set of job streams and their mutual dependencies.

Job stream view (for monitoring, troubleshooting, and impact analysis)

A expansible graphical representation of job streams and jobs in the plan. It provides a straightforward, multilevel analysis of how job and job stream completion affects plan progress and provides a direct way to work with jobs, job streams and their dependencies.

Preproduction plan view (for workload planning)

Contains job stream instances and job stream dependencies to be run during a specified time interval covered by the plan.

From each view, you can take actions on objects, view their properties, and easily switch between the views. Graphics can be exported to PNG and SVG files.

Chapter 6. End-to-end scheduling

About this task

By using end-to-end scheduling, you can schedule and control jobs on mainframe, Windows™, and UNIX™ environments, for truly distributed scheduling. In the end-to-end configuration, IBM Z Workload Scheduler is used as the planner for the job scheduling environment. IBM Workload Scheduler domain managers, standard, fault-tolerant, and z-centric agents are used to schedule on the distributed platforms. The agents replace the use of tracker agents.

IBM Z Workload Scheduler also allows you to access job streams (schedules in IBM Workload Scheduler) and add them to the current plan in IBM Z Workload Scheduler. In addition, you can build dependencies among IBM Z Workload Scheduler job streams and IBM Workload Scheduler jobs. From IBM Z Workload Scheduler, you can monitor and control the distributed agents.

The benefits that can be gained from using end-to-end scheduling are the following:

- Connecting either fault-tolerant or z-centric IBM Workload Scheduler agents to IBM Z Workload Scheduler.
- Scheduling on additional operating systems.
- Synchronization of work in mainframe and distributed environments.
- The ability for IBM Z Workload Scheduler to use multi-tier architecture with domain managers.

You can manage distributed scheduling by activating either of the following features:

- [End-to-end scheduling with fault tolerance capabilities on page 143](#)
- [End-to-end scheduling with z-centric capabilities on page 145](#)

End-to-end scheduling with fault tolerance capabilities

Learn about End-to-end scheduling with fault tolerance capabilities.

End-to-end scheduling with fault tolerance capabilities directly connects IBM Workload Scheduler standard agents, fault-tolerant agents, and domain managers (with their underlying agents and domains) to IBM Z Workload Scheduler. IBM Z Workload Scheduler is seen by the distributed network as the master domain manager.

IBM Z Workload Scheduler creates the production plan also for the distributed network and sends it to the domain managers and to the directly-connected agents. The domain managers send a copy of the plan to each of their agents and subordinate domain managers for execution.

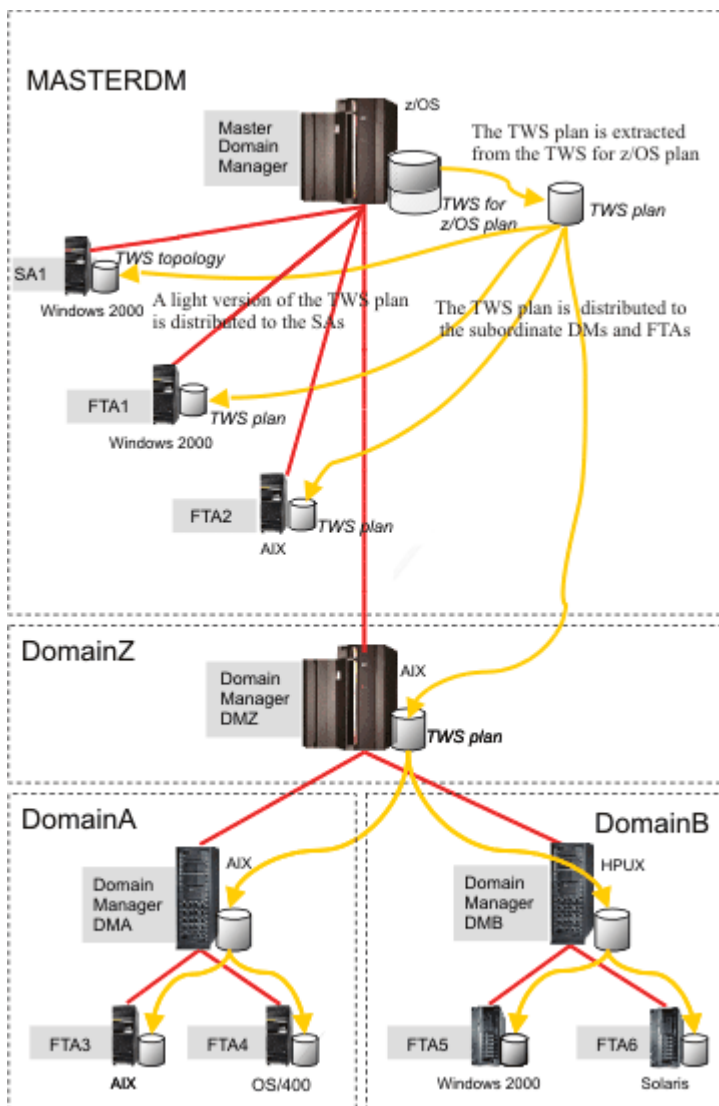
The IBM Workload Scheduler domain managers function as the broker systems for the distributed network by resolving all dependencies for their subordinate managers and agents. They send their updates (in the form of events) to IBM Z Workload Scheduler so that it can update the plan accordingly. IBM Z Workload Scheduler handles its own jobs and notifies the domain managers of all the status changes of the IBM Z Workload Scheduler jobs that involve the IBM Workload Scheduler plan. In this configuration, the domain managers and all the distributed agents recognize IBM Z Workload Scheduler as the master domain manager and notify it of all the changes occurring in their own plans. At the same time, the agents are not permitted

to interfere with the IBM Z Workload Scheduler jobs, because they are viewed as running on the master that is the only node that is in charge of them.

In the IBM Z Workload Scheduler current plan, you can specify jobs to run on workstations in the IBM Workload Scheduler network. IBM Z Workload Scheduler passes the job information to the Symphony™ file in the IBM Z Workload Scheduler server, which in turn passes the Symphony™ file to the IBM Workload Scheduler domain managers (DMZ) to distribute and process. In turn, IBM Workload Scheduler reports the status of running and completed jobs back to the current plan for monitoring in the IBM Z Workload Scheduler engine.

Figure 7: End-to-end with fault tolerance capabilities configuration on page 144 shows an IBM Workload Scheduler network managed by an IBM Z Workload Scheduler and the flow of data.

Figure 7. End-to-end with fault tolerance capabilities configuration



End-to-end scheduling with z-centric capabilities

Learn about End-to-end scheduling with z-centric capabilities.


End-to-end scheduling with z-centric capabilities directly connects IBM Workload Scheduler z-centric agents to IBM Z Workload Scheduler, that is the master domain manager for the distributed network.

Powerful mainframe capabilities, such as standard variable substitution, automatic recovery statements and alternate workstation, are supported to manage distributed workload.

Communication between the z-centric agents and IBM Z Workload Scheduler controller is direct, through the HTTP or HTTPS protocol.

Figure 8: End-to-end with z-centric capabilities configuration on page 145 shows a network with this configuration.

Figure 8. End-to-end with z-centric capabilities configuration



Distributed agents

About this task

A distributed agent is a computer running IBM Workload Scheduler on which you can schedule jobs from IBM Z Workload Scheduler. Examples of distributed agents are the following: standard agents, extended agents, fault-tolerant agents, and domain managers.

The following is a description of the types of distributed agents:

Domain Manager

The management hub in a domain. All communications to and from the agents in a domain are routed through the domain manager.

Backup Domain Manager

A fault-tolerant agent or domain manager capable of assuming the responsibilities of its domain manager for automatic workload recovery.

Dynamic workload broker

It is installed on the domain manager and performs the job management and resource management activities. It always presides over the dynamic scheduling.

Fault-tolerant Agent (FTA)

A workstation capable of resolving local dependencies and launching its jobs in the absence of a domain manager.

Standard Agent

A workstation that launches jobs only under the direction of its domain manager.

Extended Agent

A logical workstation definition that helps you launch and control jobs on other systems and applications, such as PeopleSoft, SAP, and z/OS® JES2 and JES3.

z-centric Agent

A workstation that runs jobs scheduled from IBM Z Workload Scheduler. The controller directly handles the communication with this type of agent.

Pool

Pools are groups in which you add specific dynamic agent workstations according to the requirements of the job. Jobs are assigned dynamically to the best available dynamic agent applying a load balancing policy by choosing the agent with the lesser amount of jobs to run.

Dynamic pool

Dynamic pools are groups in which you specify the requirements and jobs are assigned dynamically to the best available dynamic agent meeting the requirements and applying an optimization policy to identify the best agent.

Distributed agents replace tracker agents in IBM Z Workload Scheduler. The distributed agents help you schedule on non-z/OS systems with a more reliable and scalable agent.

In the IBM Z Workload Scheduler plan, the logical representation of a distributed agent is called a fault-tolerant workstation or a z-centric workstation.

Chapter 7. Hybrid scheduling environment

Hybrid Workload Automation allows you to use a common User Interface to manage both on premises and on Cloud engines. Manage Workload Automation on Cloud or On premises, or choose a mixture of both. You can monitor and run actions in plan on the following objects: job, critical job, workstation, resource, prompt, job stream, file, and domain.

A solution for you

The same workload efficiently managed in an on premises environment can be scheduled to run in a cloud environment. Workload Automation on Cloud reduces IT infrastructure costs and is a solution for using resources only when you have a business need to run your workload. You can transfer your workload from your on premises environment to your Workload Automation on Cloud by using a simple import procedure. The workload is instantly reproduced and is ready to be employed on your agents. Handling your workload in a cloud environment is a more cost effective way of using resources in the cloud only when you need them without relying on an on premises infrastructure.

How to work with hybrid scheduling environment

To work with the on Cloud engine from the on premises environment, define the engine connection in the Dynamic Workload Console by using port 443. This setting allows you to work with the REST APIs needed for managing Hybrid Workload Automation.

To transfer your on premises workload to the Workload Automation on Cloud environment, complete the following steps:

- Create a workload application template.
- Export the workload application template.
- Import the workload application template into the Workload Automation on Cloud environment.

Create a workload application template containing the job streams that you want to transfer and define some general information about the template, such as a name and description. Select the job streams, each containing one or more jobs, that you want to add to your template. Select the export action and save the template in a compressed file format. The compressed file contains all the files and information required to enable the workload to run in the new cloud environment with some minimal customization to adapt the workload to the cloud environment.

From the Workload Automation on Cloud dashboard, add the template to your on Cloud environment by selecting Import a custom application access point and sign in with your IBM ID. Upload the compressed workload application template file and select the agent on which to import it. You need only adjust two objects to adapt the workload to cloud environment: the prompt definition and the variable table. You must rename both of these objects by prefixing their names with your Workload Automation subscription environment ID. When the import process completes, the workload application is ready to be used in the on Cloud environment.

Notices

This document provides information about copyright, trademarks, terms and conditions for product documentation.

© Copyright IBM Corporation 1993, 2016 / © Copyright HCL Technologies Limited 2016, 2026

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing

IBM Corporation

North Castle Drive, MD-NC119

Armonk, NY 10504-1785

US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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 and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© (HCL Technologies Limited) (2026).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2016

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM® or other companies. A current list of IBM® trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe™, the Adobe™ logo, PostScript™, and the PostScript™ logo are either registered trademarks or trademarks of Adobe™ Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library™ is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open™, LTO™, the LTO™ Logo, Ultrium™, and the Ultrium™ logo are trademarks of HP, IBM® Corp. and Quantum in the U.S. and other countries.

Intel™, Intel™ logo, Intel Inside™, Intel Inside™ logo, Intel Centrino™, Intel Centrino™ logo, Celeron™, Intel Xeon™, Intel SpeedStep™, Itanium™, and Pentium™ are trademarks or registered trademarks of Intel™ Corporation or its subsidiaries in the United States and other countries.

Linux™ is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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



Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine™ is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

ITIL™ is a Registered Trade Mark of AXELOS Limited.

UNIX™ is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Index

Numerics

10.2.7 13, 14, 14, 15, 15, 16, 17
1027 13, 14, 14, 15, 15, 16, 17
2.1.3 24

A

access methods 118
actions 27, 59
advanced program-to-program communication (APPC) 133
AES algorithm 40
AES encryption 40
agent_certificate 66
agents
 authenticating with JWT 67
agents supported in IBM Workload Scheduler 119
agentscope 41
ai data advisor 61
ai data advisor 61
AI Data Advisor 72
aida 61, 61
AIDA 72
alerts, passing to NetView 129
alias 27
altrpri command 58
Anomaly detection) 72
API (application programming interface) 133
API Key authentication 66
api v2 49
APPC (advanced program-to-program communication) 133
application
 definition of 122
application programming interface (API) 133
approval 61
approval flow 61
Artificial Intelligence 72
audit-trail facility 136
authenticating with API Key 66
authenticating with JWT 67
authority checking 136
automatic
 job submission 128
 status checking 131
 status reporting 131
automatic job and started-task recovery 129, 131
automatic job stream submission 36

B

backup domain manager 99
backup dynamic domain manager 99
backup master 99
backup system 130
batchman 103
brand 48
business processing cycle 125

C

calendar 105
 definition 125
Certificate Authority 48
certificate management 41
certificate update 27
certificates 27, 48
Certificates 34
certman 41

Certman 48
certman delete 41
command-line authentication 66
commands 48
Common Programming Interface for Communications (CPI-C) 133
composer
 authenticating 66
Composer 104
configurations 137
conman
 authenticating 66
Conman 104
controlled systems 138
controlling system
 description 137
 recovery of 130
CPI-C (Common Programming Interface for Communications) 133
CREATEPERSONALAPIKEY
 security file access keyword 66
CREATESERVICEAPIKEY
 security file access keyword 66
cross-system coupling facility (XCF) 129, 131, 138
current plan 126
custom certificates 48
custom logo 48
customize 59
CyberArk
 backward compatibility 39
 previous syntax compatibility 39
 upgrading from versions 10.2.0 and earlier 39
 upgrading from versions 10.2.1 and later 39
CyberArk integration 39
CyberArk upgrade scenarios 39

D

database 47, 48
DELETEALLAPIKEYS
 security file access keyword 66
dependencies 43
 defining 124
description 16, 28
description field 16
description job stream 16
design 42, 42, 43
designer 42, 42, 47, 48
documentation 28
domain manager 99, 99
domain managers
 no-root installation 51
dynamic agent 45
dynamic agents 27
 authenticating with JWT 67
dynamic domain manager 99
dynamic domain managers
 no-root installation 51
dynamic workload console 59

E

edit monitoring 11
edwa 28
ENABLEFIPS
 FIPS data encryption 137
encryption at rest 72
encryption in motion 73

enhance integrations 61
enhancements 33, 33, 34, 36, 39, 47, 47, 48
event driven 28
event rule 17
events 28
extended agent 100
extended agents 118

F

fault-tolerant agent 99
fault-tolerant agents
 no-root installation 51
fence command 58
file 16
file dependency 113
File Transfer 11, 25
filtering 37
filtering job stream 16
filters 27
folder 16

G

global options 116
graphical designer 42, 42, 43, 47

H

human-in-the-loop 61

I

IBM Workload Scheduler 118, 121, 139
 supported agents 119
IBM Workload Scheduler agents 119
IBM Workload Scheduler/NetView 113
integration 11, 25, 61
integrations 61
 moved to Automation Hub 75
ISPF (Interactive System Productivity Facility) dialog 127

J

jira 61
Jira 61
job completion checker (JCC) 131
job dependencies 124
job monitoring 43
job recovery
 automatic 129
 manual 134
job stream 14, 14, 15, 15, 16
Job Stream Submission job 36
job stream view 29
Job stream view 44
job streams 139
job streams within job streams 36
job submission
 automatic 128
 manual 134
job tailoring 128
jobman 103
jobtable dir encryption 72
js 14, 14, 15, 15, 16
JSONata 61
JWT
 downloading to dynamic agents 67
 installing on dynamic agents 67

L

layout locking 29
limit cpu command 58
limit sched command 58

LISTALLAPIKEYS
security file access keyword 66
local options 116
login 59
logo 48
long-term plan 126

M

Machine Learning 72
mailman 103
manual status control 134
master domain 97
master domain manager 99
master domain managers
no-root installation 51
memory management 45
merge 13, 13
message boxes encryption 72
MFT 11, 25
modeling 42, 42, 43
monitor 60
monitor workload 60
monitoring 37, 38, 43, 59, 60
Monitoring 33, 33, 34
multi-selection 29
multi-tier architecture 143
multiple job stream 14, 15, 15
multiple job stream submission 14, 14, 15, 15
multiple workspaces 24

N

national language features 121
nested job streams 36
netman 103
NetView
alerts 129
network Agent 101
new tool 48
no-root installation
domain managers 51
dynamic domain managers 51
fault-tolerant agents 51
master domain managers 51
server components 51

O

occurrences 126, 126
OCLI 48
OCLI model 48
OCLI model commands 48
operation dependencies 124
operator instructions 28, 38
oql 16, 27, 37
OQL 33, 33
oql filters 37
Orchestration CLI 48
Orchestration CLI model commands 48
orchestration monitor 16, 27, 27, 37, 38, 43,
44, 45, 59
Orchestration Monitor 34
Orchestration Query Language 33, 33

P

parameter 106
password management 39, 40
password resolver 36, 39
password solver 39
password storage 39
password vault 39
performance 45
periods 125
PIF (program interface) 132

PIF applications
applications 139
plan
current 126
definition of 126
detailed 126
long-term 126
modification of 134
trial 122
types 122
planning
trial plans 122
plug-in 11, 25, 43
plug-ins
moved to Automation Hub 75
postgres 47, 48
postgresql 47, 48
production control file 98
production period 107
production workload restart 129, 131
program interface (PIF) 132
prompt 16, 106
prompt dependency 113

R

RACF (Resource Access Control Facility) 136
recovery 129, 131
recovery job 109
recovery prompt 109
remote 25
remote certificate update 27
remote dialogs
dialogs 139
remote-to-remote 25
remove hold for definition 44
remove no operation for definition 44
Requests
Enhancements 56, 69, 77
Requests for Enhancements 56, 69, 77
resource 106
Resource Access Control Facility (RACF) 136
resources 43
rest api 49
rest api v2 49
restapi 49
restapi v2 49
restart 129, 131
restart management 129, 129, 131
RESTful 61
RFEs 56, 69, 77
run cycle 105
run cycle group 42
run cycle groups 125
run cycle preview 14

S

SAF (system authorization facility) 136
schedule 126
search folder 16
secure communication 48
security 136
security file access keyword 66
self-contained job 49
self-service catalog 61
send event 17
server components
no-root installation 51
service 61
service now 17
servicenow 61
ServiceNow 61
services 61

set hold for definition 44
set no operation for definition 44
simulation with trial plans 122
size monitoring 11
special resources
definition of 124
SSL communication for FTAs 73
standalone job definition 49
standard agent 100
standard list file 115
started 27
status checking, automatic 131
status control
manual 133, 133
status inquiries 133
status reporting
automatic 131
from heterogeneous environments 131
from user programs 131
step-level restart 129
submission 14, 14, 15, 15
submitting job streams 36
SVG export 29
symphony 98, 108
Symphony encryption 72
syntax 37
SYSOUT, checking of 131
system authorization facility (SAF) 136
system automation commands tailoring 128
system failures 129
Systems Application Architecture
Common Programming Interface for
Communications 133

T

TECEvent Forwarder
moved to Automation Hub 76
Tivoli Business Systems Manager 113
tracker agents 143
tree view 16
treeview 16
trial plans 122
triggers 42
Triple DES
deprecated 40
replaced with AES 40
TWSObjectsMonitor 28
TWSObjectsMonitor events 28

U

ui 24
unplannable work 132
update agent certificates 27
user 106
user authority checking 136

V

variable 106
variable table 106
variables 28
virtual workstation
definition of 123

W

wappman
authenticating 66
what's new 24
work submission, automatic 128
workload 60
workload automation 29
workload designer 42, 42
Workload Manager (WLM) 131

- workload restart 129, 131
- workspace 13
- workspace merge 13
- workspace merging 13
- workspaces 24
- workstation 42
 - changing the status of 134
 - definition 123
- workstation class 104
- workstations 42
- writer 103

X

- XCF (cross-system coupling facility) 129, 131, 138

Z

- z-centric agent 100