IBM

IBM® Workload Scheduler
Glossary

Note

Before using this information and the product it supports, read the information in Notices on page xl.

This edition applies to version 10, release 2, modification level 1 of IBM® Workload Scheduler (program number 5698-T09) and to all subsequent releases and modifications until otherwise indicated in new editions.

Chapter 1. Glossary

Use the glossary to find terms and definitions for the IBM Workload Automation products.

The following cross-references are used:

- See refers you from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- See also refers you to a related or contrasting term.

To view glossaries for other IBM products, go to www.ibm.com/software/globalization/terminology.

A on page 4 B on page 6 C on page 7 D on page 11 E on page 14 F on page 16 G on page 18 H on page 19 I on page 19 J on page 21 L on page 23 M on page 24 N on page 25 O on page 26 P on page 27 Q on page 30 R on page 30 S on page 32 T on page 35U on page 36 V on page 37 W on page 37 X on page 39Z on page 39

A

access method

An executable file used by extended agents to connect to and control jobs on other operating systems (for example, z/OS®) and applications (for example, Oracle Applications, PeopleSoft, and SAP R/3). The access method is specified in the workstation definition for the extended agent. See also "extended agent on page 16".

active application description

An application description that is complete and ready for use in planning or scheduling. See also pending application description on page 27.

actual duration

At a workstation, the actual time in hours, minutes, and seconds it takes to process an operation from start to finish.

actual start time

The time that an IBM Workload Scheduler job instance or job stream instance actually starts. See also:

- "earliest start time on page 14"
- "latest start time on page 23"
- "planned start time on page 28"
- "scheduled time on page 32"

adjusted quantity

The current quantity of a special resource, taking the deviation into account.

advanced ISPF panel

A panel style that provides a quick, at-a-glance scrollable view of the AD and CP operations, with color-coded fields that represent application and operation status. You can select administrative tasks to perform from an Action menu on the main menu. You can list and browse a single AD or list and browse a single operation in the plan.

AD

application description

ad hoc job

A job that is inserted into the current production plan. These jobs are unique to the plan, and are not saved in the database. See also:

- "database on page 11"
- "plan on page 28"

ad hoc prompt dependency

A prompt dependency that is defined within the properties of a job or job stream and is unique to that job or job stream. See also "prompt dependency on page 30".

agent

An installed component that enables jobs to be run on a computer or a computer partition, provided that the computer or computer partition is also defined as a workstation in the IBM Workload Scheduler database. Agents can be standard, fault-tolerant, extended, or network. Specially configured agents are also used as backups for domain managers and the master domain manager. See also:

- "backup domain manager on page 6"
- "backup master domain manager on page 7"
- "fault-tolerant agent on page 16"
- "network agent on page 25"
- "standard agent on page 33"
- "extended agent on page 16"

all-days cyclic period

A cyclic period where all days are counted when calculating the interval.

application

A measurable and controllable unit of work that completes a specific user task, such as the running of payroll or financial statements. The smallest entity that an application can be broken down into is an operation. Generally, several related operations make up an application.

application description

A database description of an application.

application group

Type of application description which holds run cycle and calendar information for standard applications or job descriptions which have been defined as a member of the group.

application ID

The name of an application. (For example, PAYROLL or DAILYJOBS.)

application version

See version.

audit

A process that logs modifications to the database and plan.

automatic events

Events recognized by or triggered by an executing program. Automatic events are usually generated by tracking programs but can also be created by a user-defined program.

automatic hold/release

Function used to control jobs that are submitted outside IBM® Z Workload Scheduler. It allows you to define whether such jobs should be automatically released at the appropriate time if placed in HOLD status when submitted.

automatic job and started-task recovery

In IBM® Z Workload Scheduler, a function that lets you specify, in advance, alternative recovery strategies for operations that end in error.

automatic-reporting workstation

A workstation (for example, a processor or printer) that reports events (the starting and stopping of operations) in real time to IBM® Z Workload Scheduler.

availability

The degree to which a system (and in IBM® Z Workload Scheduler, an application) or resource is ready when needed to process data.

B

backup domain manager

An agent in a distributed IBM Workload Scheduler network that can assume the responsibilities of its domain manager. It is installed as a full status, fault-tolerant agent. See also:

- "fault-tolerant agent on page 16"
- "full status on page 17"
- "domain manager on page 13"

backup dynamic domain manager

An agent, in a distributed and z/OS IBM Workload Scheduler network, that can assume the responsibilities of its dynamic domain manager.

backup master domain manager

An agent, in a distributed IBM Workload Scheduler network, that can assume the responsibilities of the master domain manager. It is installed as a full status, fault-tolerant agent. See also:

- "fault-tolerant agent on page 16"
- "full status on page 17"
- "master domain manager on page 24"

batch loader

In IBM® Z Workload Scheduler, a batch program that you can use to create and update information in the application-description and operator-instruction databases.

batchman

A production control process that interacts directly with a copy of the `Symphony` file distributed to workstations at the beginning of the production period and updates it, resolving dependencies. It is the only process that can update the `Symphony` file. See also:

- "processes on page 29"
- "production period on page 29"
- "symphony file on page 34"

C

calendar

A list of scheduling dates. Calendars are defined in the database and are mostly assigned to run cycles. Calendars can be used either to identify the dates when job streams or jobs can be run (when used with inclusive run cycles), or when they cannot be run (when used with exclusive run cycles). A calendar can also be designated for use as a non-working days calendar in a job stream. See also:

- "exclusive run cycle on page 15"
- "inclusive run cycle on page 20"
- "non-working days calendar on page 26"

capacity

The actual number of parallel servers and workstation resources available during a specified open interval.

capacity ceiling

The maximum number of operations that a workstation can handle simultaneously.

carry forward

If a job stream is not completed before the end of the current production period it can be carried forward to the next and then to subsequent periods, until the latest start time is reached or the job completes. See also "latest start time on page 23".

catalog

A directory of files and libraries, with reference to their locations. A catalog may contain other information such as the types of devices in which the files are stored, passwords, blocking factors.

closed workstation

A workstation that is unavailable to process work for a specific time, day, or period.

command-line client

A component you use to run selected IBM Workload Scheduler master domain manager commands from any workstation where it is installed. It communicates by TCP/IP with the command-line server, which is part of the master domain manager. The command-line client does not need to be installed on the master domain manager, and is a selectable option for installation on other nodes in the network. For details of the supported commands see the *IBM Workload Scheduler: Planning and Installation*. See also "master domain manager on page 24".

command-line server

See "command line client on page 8".

completion code

In IBM® Z Workload Scheduler, a system code that indicates how the processing of an operation ended at a workstation. See "error code on page 14".

composer

A command-line program for managing the definitions of scheduling objects in the database. See also "database on page 11".

computer workstation

A workstation that performs z/OS processing of jobs and started-task operations, and that usually reports status to IBM® Z Workload Scheduler automatically. A processor used as a workstation. It can refer to single processors or multiprocessor complexes serving a single job queue (for example, JES2 or JES3 systems).

condition

A set of one or more condition dependencies aggregated by a rule that can be:

AND

All condition dependencies are true

OR

At least *n* condition dependencies are true.

condition dependency

A specific check of the status or return code of a conditional predecessor within a condition defined on the conditional successor.

condition dependency status

The evaluation result of the specific conditional predecessor check:

False

Return code or status check was done and the result was false

True

Return code or status check was done and the result was true

Undefined

Not yet evaluated.

condition status

The evaluation result of the rule applied to the condition dependencies statuses:

False

According to the rule:

ALL

At least one condition dependency is false

OR

At least *n* condition dependencies (Simples) are false

True

According to the rule:

ALL

All condition dependencies are true

OR

At least *n* condition dependencies are true

Undefined

Not yet possible to evaluate.

conditional dependency

A relationship between an operation or an operation step and another operation set up by using a condition.

conditional predecessor

A predecessor job involved in a condition.

conditional successor

A successor job involved in a condition.

conman

A command-line program for monitoring and managing the production environment. See also "processes on page 29".

connector

An installed component that provides the interface between the engine and the Dynamic Workload Console. See also:

- "engine on page 14"
- "Dynamic Workload Console on page 14"

controller

In IBM® Z Workload Scheduler, the component that runs on the controlling system, and that contains the tasks that manage the plans and databases.

controlling system

The system that the controller runs on.

control on servers

If a workstation is defined with control on servers, IBM® Z Workload Scheduler will not start more operations at the workstation than there are available servers.

CP

Current plan.

CPU

See "workstation on page 38".

cpu time

The processor time used by a job. See also "duration on page 13".

cross dependency

A dependency of a local job on a remote job running in a different scheduling environment. It is achieved by using a shadow job, which runs in the same environment as the local job and maps the remote job processing.

See also:

- remote engine workstation on page 30
- shadow job on page 33

current plan (CP)

A detailed plan of system activity that covers a period of at least 1 minute, and not more than 21 days. A current plan typically covers 1 or 2 days.

cyclic interval

The number of days in a cyclic period.

cyclic period

A period that represents a constant number of days. There are two types of cyclic periods:

- Work-days-only cyclic period, where only the work days are counted when calculating the number of days in the period.
- All-days cyclic period, where all days are counted.

D

daily planning

The process of creating a current plan.

database

In IBM Workload Scheduler distributed environments it consists of a set of tables defined in a relational database, such as DB2® or Oracle, containing the definitions for all scheduling objects (jobs, job streams, resources, workstations, domains, parameters, prompts, variables, and files). It also contains data such as job and job stream statistics, user data, and object creation and modification time stamps.

In IBM Z Workload Scheduler, it is defined as a collection of data that is fundamental to a system. The product uses six databases: calendar, period, workstation description, JCL variable table, application description, and operator instruction.

See also "plan on page 28".

data store

In IBM® Z Workload Scheduler, the component managing the job runtime information at the tracked system. It is dedicated to the storing and possible retrieval of sysout data sets belonging to submitted jobs, to optimize the sysout availability.

deadline

The time by which a job or job stream is set to complete. When a job or job stream passes the deadline, notifications are sent to users and integrated applications, but the job or job stream is not prevented from running if all time restrictions and dependencies are satisfied. Jobs or job streams that have not yet started or that are still running after the deadline time has expired are considered "late" in the plan. See also "plan on page 28".

deadline date

The latest date by which an occurrence must be complete.

deadline time

The latest time by which an occurrence must be complete.

deadline WTO message

You can specify that IBM® Z Workload Scheduler issue an operator message (EQQW776I) when a started operation has not been marked as completed before the deadline time. In addition to the standard message, the user-defined text that describes the operation is issued as part of the WTO.

default calendar

A calendar that you have defined for IBM® Z Workload Scheduler to use when you do not specify a calendar in an application description. A calendar that IBM® Z Workload Scheduler uses if you have neither specified a calendar in an application description, nor defined your own default calendar.

dependency

A prerequisite that must be satisfied before a job or job stream can start. See also:

- "external dependency on page 16"
- "file dependency on page 16"
- "follows dependency on page 17"
- "prompt dependency on page 30"
- "resource dependency on page 31"

In IBM® Z Workload Scheduler, a relationship between two operations in which the status or the return code of the first operation determines the starting of the second operation.

descriptive text

User-written text describing the operation. This text is also issued as part of the write-to-operator message if the operation has been started, exceeds its deadline, and has the deadline write-to-operator (WTO) option specified.

deviation

A temporary variation in the quantity of a special resource.

distributed agent

The software running on a computer that is part of a IBM Workload Scheduler domain, 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. Contrast with fault-tolerant workstation and tracker agent.

distributed network

A connected group of workstations that use the IBM Workload Scheduler distributed engine to perform workload scheduling. See also:

- "engine on page 14"
- "workstation on page 38"

distributed workstation

A workstation on which jobs and job streams are run using the distributed engine. See also:

- "engine on page 14"
- "workstation on page 38"

domain

A named group of workstations in a distributed IBM Workload Scheduler network, consisting of one or more agents and a domain manager acting as the management hub. All domains have a parent domain except for the master domain. See also:

- "domain manager on page 13"
- "master domain manager on page 24"

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 agents in the domain is routed through the domain manager. See also "workstation on page 38"

duration

The elapsed time that a job is expected to take to complete (estimated duration) and actually takes (actual duration). See also:

- "cpu time on page 10"
- "time restriction on page 36"

dynamic agent

An agent, in an IBM Workload Scheduler distributed or end-to-end network, that runs:

- Dynamically, workload on the workstation that best meets the requirements needed to run it
- Existing job types (docommand and script)
- Job types with advanced options that are either supplied with the product or implemented by using custom plug-ins.

You can group dynamic agents in pools and dynamic pools. See also pool on page 28 and dynamic pool on page 14.

dynamic domain manager

An installed component, in a IBM Workload Scheduler distributed and z/OS network, that is the management hub in a domain running both static and dynamic workload. All communications to and from the dynamic agents in the domain are routed through the dynamic domain manager. See also dynamic agent on page 13 and backup dynamic domain manager on page 7.

dynamic pool

A logical workstation grouping a set of dynamic agents, which is dynamically defined based on the resource requirements you specify. See also dynamic agent on page 13.

Dynamic Workload Console

A web-based graphical user interface used to create, modify, and maintain job scheduling objects, to manage the production environment, and to produce reports. See also "views on page 37".

E

earliest start time

The time before which a job or job stream cannot start. The job or job stream can start after the earliest start time provided that all other time restrictions and dependencies are satisfied. It is set using the *at* Dynamic Workload Console option or in the command-line scheduling language using the *at* keyword . See also:

- "actual start time on page 4"
- "latest start time on page 23"
- "planned start time on page 28"
- "scheduled time on page 32"

end-to-end network

A network obtained by connecting one or more IBM Workload Scheduler fault-tolerant agents in a distributed network to an IBM Workload Scheduler for z/OS® node in a z/OS® network using TCP/IP, to perform workload scheduling. In this configuration, the IBM Workload Scheduler for z/OS® node becomes the master domain manager of the fault-tolerant agents to schedule and manage jobs in the distributed network. See also:

- "engine on page 14"
- "workstation on page 38"

engine

The core software for the scheduling environment. The engine can be either a z/OS® engine (installed as part of IBM Workload Scheduler for z/OS®) or a distributed engine (installed as part of IBM Workload Scheduler).

error code

A code set by IBM® Z Workload Scheduler to describe how the processing of an operation ended at a computer workstation.

estimated duration

The estimated length of time an operation will use a workstation. This is initially based on a value that is provided when the operation is defined, but can be adjusted automatically by the feedback mechanism to reflect actual durations. The minimum value of planned duration is 1 second, and the maximum value is 99 hours 59 minutes 00 seconds. If you specify 99 hours 59 minutes 01 seconds, you do not receive an alert message if the actual duration is greater than the planned duration.

ETT

See event-triggered tracking.

event

An action that changes the status of an operation and changes the current plan.

event manager

In IBM® Z Workload Scheduler, the function that processes all tracking events and determines which of these are related to IBM® Z Workload Scheduler.

event reader

In IBM® Z Workload Scheduler, a task that reads event records from an event data set.

event tracking

A function of IBM® Z Workload Scheduler that follows events in the operations department in real time and records status changes in the current plan.

event-triggered tracking (ETT)

A component of IBM® Z Workload Scheduler that waits for specific events to occur, and then adds a predefined application to the current plan. ETT recognizes two types of events: the reader event, which occurs when a job enters the JES reader, and the resource event, which occurs when the availability status of a special resource is set to “yes”.

event writer

In IBM® Z Workload Scheduler, a task that writes event records in an event data set.

exclusive resource

A resource that can be used by only one operation at a time.

exclusive run cycle

A run cycle that specifies the days and times that a job stream *cannot* be run. Exclusive run cycles take precedence over inclusive run cycles. See also "run cycle on page 32".

expected arrival time

The time when an operation is expected to arrive at a workstation. It can be calculated by daily planning or specified in the long-term plan.

explorer view

A graphical view in the Dynamic Workload Console used to modify and maintain job streams in the database and the plan. See also:

- "database on page 11"
- "plan on page 28"
- "views on page 37"

extended agent

An agent used to integrate IBM Workload Scheduler job control features with other operating systems (for example, z/OS®) and applications (for example, Oracle Applications, PeopleSoft, and SAP R/3). Extended agents must be hosted by a master domain manager, domain manager, or an agent (not another extended agent), and use access methods to communicate with external systems. See also "access method on page 4".

extended status code

Together with the normal status codes, IBM® Z Workload Scheduler maintains extended status codes that provide additional information about the status of operations. The extended status code is not always present.

external dependency

A dependency defined in one job or job stream that refers to another job stream or to a job in another job stream. In IBM® Z Workload Scheduler, A relationship between two occurrences, in which an operation in the first occurrence (the predecessor) must successfully finish before an operation in the second occurrence (the successor) can begin processing.

external job

A job referred to in an external dependency. See also "external dependency on page 16".

F

fault-tolerant agent

An installed agent component in a distributed IBM Workload Scheduler network capable of resolving local dependencies and launching its jobs in the absence of a domain manager.

fault-tolerant workstation

In IBM® Z Workload Scheduler, a computer workstation configured to schedule jobs on a distributed computer. A fault-tolerant workstation is the logical representation in the IBM® Z Workload Scheduler plan.

feedback limit

A numeric value that defines the limits within which actual data that is collected during tracking is fed back and used by IBM® Z Workload Scheduler.

fence

Regulates whether a job can be run on a workstation. The job fence is a priority level that the priority of a job must exceed before it can run.

file dependency

A dependency where a job or job stream cannot start until it finds a specific file is present in a specific path on a specific workstation. Sometimes called an *opens file* dependency. See also "dependency on page 12".

filter criteria

Input values that are used to limit the mass update of applications to only those specified. This term is used in the ISPF panels.

final job stream

The last job stream that is run in a production period. It contains scripts that generate the next production plan. See also:

- "production period on page 29"
- "production plan on page 29"

first critical operation

An operation of an occurrence that has the earliest latest-start-time. The first critical operation of an occurrence determines the critical path.

first operation

An operation in an occurrence that has no internal predecessor. The start node in a network.

fixed resources

A set of resource names used to check the authority of users to access the panels.

follows dependency

A dependency where a job or job stream cannot start until other jobs or job streams have completed successfully. See also "dependency on page 12".

forecast plan

A projection over a selected timeframe based on the job streams and dependencies defined in the database. See also

- "database on page 11"
- "plan on page 28"

form number

A user-defined code that identifies the type of paper to be used for an operation on a printer workstation. IBM® Z Workload Scheduler can use the form number to identify the different print operations belonging to one job.

free day

In IBM® Z Workload Scheduler, any day that is not a work day.

freedays calendar

See "non-working days calendar on page 26".

FTA

See "fault-tolerant agent on page 16"

full status

An attribute of an agent that enables it to be updated with the status of jobs and job streams running on all other workstations in its domain and in subordinate domains, but not on peer or parent domains. A backup domain manager or master domain manager must be full status. See also:

- "backup domain manager on page 6"
- "domain on page 13"
- "master domain manager on page 24"

G

general workstation

A workstation where activities other than printing and processing are carried out. A general workstation reporting to IBM® Z Workload Scheduler is usually manual, but it can also be automatic. Manual activities can include data entry and job setup.

generic alert

An alert that is broadcast by IBM® Z Workload Scheduler, and collected by NetView, when an operation ends in error. You can specify this as an option when defining application descriptions.

global options

Configuration options defined on the master domain manager using **optman**. These options apply to all workstations in the IBM Workload Scheduler network. See also:

- "local options on page 23"
- "optman on page 27"
- "user options on page 36"

global search character

In IBM® Z Workload Scheduler, a percent sign (%), which represents any single character, or an asterisk (*), which represents any character string of any length.

global variable table

The JCL variable table that IBM® Z Workload Scheduler checks for a variable substitution value if no value is found in the specific JCL variable table that is associated with the operation.

graph view

A graphical view in the Dynamic Workload Console used to modify and maintain job streams in the database and the plan. See also:

- "database on page 11"
- "plan on page 28"
- "views on page 37"

group definition

The application group to which the application description or job description is a member.

H

highest return code

A numeric value in the range 0–4095. If this return code is exceeded during job processing, the job will be reported as ended-in-error.

holidays calendar

The default non-working days calendar for all job streams. You must create it and name it "holidays", otherwise the default non-working days are considered to be all Saturdays and Sundays. See also:

- "calendar on page 7"
- "non-working days calendar on page 26"

host

A workstation required by extended agents. It can be any IBM Workload Scheduler workstation except another extended agent.

I

IBM Workload Scheduler distributed Agent for z/OS (ddriven)

An agent used to submit jobs (JCL) from IBM Workload Scheduler on the JES2 subsystem of z/OS.

IBM Z Workload Scheduler Agent (z-centric)

An agent that can be configured, in an IBM® Z Workload Scheduler end-to-end environment to connect to:

- IBM® Z Workload Scheduler controller to run:
  - Existing job types (scripts)
  - Job types with advanced options that are either supplied with the product or implemented by using custom plug-ins.
- Dynamic Domain Manager to run:
  - Existing job types (scripts)
  - Job types with advanced options that are either supplied with the product or implemented by using custom plug-ins.
  - Jobs on agents that best meet the job requirements.

iCalendar

A standard (RFC 2445) for calendar data exchange. Specific iCalendars can be supplied in place of IBM Workload Scheduler calendars to determine the dates on which jobs or job streams should run. See also "calendar on page 7".

impact view

A graphical view in the Dynamic Workload Console used to modify and maintain job stream instance dependencies in the plan. See also:

- "plan on page 28"
- "views on page 37"

incident log

An optional function available under the job completion checker.

inclusive run cycle

A run cycle that specifies the days and times that a job stream *is scheduled* to be run. Exclusive run cycles take precedence over inclusive run cycles. See also "run cycle on page 32".

initiator/terminator

The job scheduler function that selects jobs and job steps to be executed, allocates input/output devices for them, places them under task control, and at completion of the job, supplies control information for writing job output on a system output unit.

input arrival time (IAT)

The user-defined date and time when an operation or an application is planned to be ready for processing.

interactive jobs

A job that runs interactively on a Windows® desktop.

intermediate start

The date and time an operation started after processing was interrupted.

internal date

Internally, IBM® Z Workload Scheduler uses a two-digit year format when handling dates. In order to handle dates before and after 31 December 1999 correctly, IBM® Z Workload Scheduler uses an origin year of 72 for the internal century window. This means that internally the year 1972 is represented as 00 and 2071 is represented as 99.

internal dependency

A relationship between two operations within an occurrence, in which the first operation (the predecessor) must successfully finish before the second operation (the successor) can begin.

internal status

The current status of jobs and job streams in the IBM Workload Scheduler engine. The internal status is unique to IBM Workload Scheduler. See also "status on page 34".

internetwork dependencies

A dependency between jobs or job streams in separate IBM Workload Scheduler networks. See also "network agent on page 25".

internetwork job or job stream

A job or job stream in a remote IBM Workload Scheduler network that is referenced by an internetwork dependency defined for a job or job stream in the local network. See also "network agent on page 25".

J

JCC

See *job completion checker*.

JCL tailoring

IBM® Z Workload Scheduler provides automatic JCL tailoring facilities, which enable jobs to be automatically edited using information that is provided at job setup or submit.

Jnextday

The previously used term for: "JnextPlan on page 21".

Jnextplan

A job that creates or extends the production plan. See also "production plan on page 29".

job

A unit of work that is processed at a workstation. The job definition consists of a unique job name in the database along with other information necessary to run the job. See also "job definition on page 21". In IBM® Z Workload Scheduler, an operation performed at a computer workstation. A job usually includes all necessary computer programs, linkages, files, and instructions to the operating system.

job class

Any one of a number of job categories that can be defined. By classifying jobs and directing initiators to initiate specific classes of jobs, it is possible to control a mixture of jobs that can be run concurrently.

job-completion checker (JCC)

An optional function of IBM® Z Workload Scheduler that allows extended checking of the results from CPU operations.

job conditional dependency

A specific check of the status or return code of a conditional predecessor job within a condition.

job definition

A definition of a unit of work that resides in the database of the distributed IBM Workload Scheduler engine and can be added to a job stream. Job definitions can be created before creating a job stream, or can be created as part of the creation or modification of a job stream. See also "job stream on page 22".

job description

A single processor (job or started-task) operation and its dependencies.

job ID

The JES job ID of the job associated with the operation.

job instance

A job scheduled for a specific run date in the plan. See also "job on page 21".

job limit

See "limit on page 23"

jobman

A job management process that controls the launching of jobs under the direction of **batchman** and reports job status back to **mailman**. The **jobman** process is responsible for tracking job states and for setting the environment as defined in `jobmanrc` and `.jobmanrc` when requesting job launches. See also:

- "batchman on page 7"
- "jobmon on page 22"
- "mailman on page 24"

jobmon

A job management and monitoring process in the Windows® version of IBM Workload Scheduler. A separate **jobmon** process is spawned to launch and monitor each job. It reports job status back to **jobman**. See also:

- "jobman on page 22"
- "processes on page 29"

job name

The name of the job associated with an operation. The job name is assigned in the JOB statement of a job. It identifies the job to the system.

job preparation

Job preparation involves modifying jobs in preparation for processing. This can be performed manually, by an operator, or automatically by the JCL tailoring functions of IBM® Z Workload Scheduler.

job setup

The preparation of a set of JCL statements for a job at a job setup workstation. Job setup can be performed manually by an operator, or automatically by IBM® Z Workload Scheduler.

job setup workstation

A general workstation defined with the job setup option. A job setup workstation lets you modify your job or STC JCL before execution.

job status

See "status on page 34".

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.

job stream instance

A job stream that is scheduled for a specific run date in the plan. See also "job stream on page 22".

job submission

In IBM® Z Workload Scheduler, a process that presents jobs to z/OS for running on a workstation once the scheduling criteria for the operation is met.

job tracking

In IBM® Z Workload Scheduler, a process that communicates with operating systems that control computer workstations.

L

last operation

An operation in an occurrence that has no internal successor. The terminating node in a network.

latest out time

See *latest start time*.

latest start time

The time before which the job or job stream must start. The job or job stream can start before the latest start time provided that all other dependencies are satisfied. It is set in the command-line scheduling language using the *until* keyword. See also:

- "actual start time on page 4"
- "earliest start time on page 14"
- "planned start time on page 28"
- "scheduled time on page 32"

The latest out time for an operation is identical to the latest start time.

limit

A means of allocating a specific number of job slots into which IBM Workload Scheduler is allowed to launch jobs. A job limit can be set for each job stream, and for each workstation. For example, setting the workstation job limit to 25 permits IBM Workload Scheduler to have no more than 25 jobs running concurrently on the workstation.

list

A means of filtering plan and database objects and presenting them in a table.

local options

Configuration options defined on each workstation in the `localopts` file. Each workstation in the IBM Workload Scheduler network must have a `localopts` file. The settings in this file are changed using a text editor, and apply only to that workstation. See also:

- "global options on page 18"
- "user options on page 36"

localopts file

A file in which the local options are defined. Each workstation in the IBM Workload Scheduler network must have a localopts file.

local processor

In a complex of processors under JES3, a processor that executes users' jobs and that can assume global functions if the global processor fails. In IBM® Z Workload Scheduler, a processor in the same installation that communicates with the controlling IBM® Z Workload Scheduler processor through shared DASD or XCF communication links.

logman

A command that produces job statistics from the previous production plan log file, and updates the preproduction plan.

long-term plan (LTP)

A high-level plan of system activity that covers a period of at least 1 day, and not more than 4 years. It serves as the basis for a service level agreement with your users, and as input to daily planning.

M

makesec

A command-line utility that compiles the security file. See also "security file on page 32".

mailman

A mail management process. It routes messages to local and remote workstations. Additional **mailman** processes named **ServerID**s are created on domain managers to divide the load on mailman and improve the efficiency of message handling. When the domain manager starts up, it creates a separate **mailman** process instance for each **ServerID** specified in the workstation definitions of the agents it manages. Each workstation then contacts its own **ServerID** on the domain manager instead of contacting the main **mailman** process. See also "processes on page 29".

manual reporting

A type of workstation reporting in which events, once they have taken place, are manually reported to IBM Workload Scheduler for z/OS. This type of reporting requires that some action be taken by a workstation operator. Manual reporting is usually performed from a list of ready operations.

mass updating

A function of the Application Description dialog in which a large update to the application database can be requested.

master domain manager

An installed component that performs the role of management hub of the top-level domain in the IBM Workload Scheduler network. It maintains the database of all scheduling objects in the domain and the central

configuration files. The master domain manager generates the plan and creates and distributes the Symphony file. In addition, logs and reports for the network are maintained on the master domain manager. See also:

- "backup master domain manager on page 7"
- "database on page 11"
- "domain on page 13"
- "plan on page 28"

MDM

See "master domain manager on page 24".

metronome

An application that takes a snapshot of the IBM Workload Scheduler configuration and generates an HTML report. It is used in problem determination to provide information to IBM Software Support.

modify current plan (MCP)

In IBM® Z Workload Scheduler, a dialog function used to dynamically change the contents of the current plan to respond to changes in the operation environment. Examples of special events that would cause alteration of the current plan are: a rerun, a deadline change, or the arrival of an unplanned application.

most critical application occurrences

Those unfinished applications whose latest start time is less than or equal to the current time.

mozart

The previously used term for the "database on page 11".

N

netman

A network management process that is started by the **Startup** script in UNIX®, or as a service in Windows®. **Netman** behaves like a network listener program which receives **conman start**, **stop**, **link** or **unlink** requests from the network. The **netman** process examines each request received and either implements the request itself or spawns a local IBM Workload Scheduler process to do so. See also "processes on page 29".

network agent

A logical extended agent used to create dependencies between jobs and job streams on separate IBM Workload Scheduler networks. See also "internetwork dependencies on page 20".

NOERROR table

A storage control block containing the currently active NOERROR definitions, specifying a list of error codes that, for job-tracking purposes, are treated as normal completion codes. The controller builds the NOERROR table at startup. You can check the NOERROR table content by using the modify command /F subsys,LSTNOERR.

noncyclic period

A period that does not represent a constant number of days or work days. Examples: quarter, academic semester.

nonreporting

A reporting attribute of a workstation, which means that information is not fed back to IBM® Z Workload Scheduler.

non-working days calendar

A calendar assigned to a job stream to represent the non-working days when job streams and jobs are not to be run. It can also be used to designate Saturdays or Sundays, or both, as workdays. Formerly called *freedays* calendar. See also:

- "calendar on page 7"
- "holidays calendar on page 19"

O

occurrence

An instance of an application in the long-term plan or current plan. An application occurrence is one attempt to process that application. Occurrences are distinguished from one another by run date, input arrival time, and application ID. For example, an application that runs four times a day is said to have four occurrences per day.

occurrence group

Consists of one or more application occurrences added to the long-term plan or current plan, where such occurrences are defined as belonging to a particular application group specified in the group definition field of the application description or job description.

offset

Values, in the ranges 1 to 999 and −1 to −999, that indicate which days of a calendar period an application runs on. This is sometimes called displacement.

offset-based run cycle

A run cycle that uses a combination of user-defined periods and offsets. For example, an offset of 3 in a period of 15 days is the third day from the beginning of the period. It is more practical to use offset-based run cycles when the cycle is based on cyclic periods. This term is only used as such in IBM Workload Scheduler for z/OS®, but the concept applies also to the distributed product. See also:

- "rule-based run cycle on page 31"
- "run cycle on page 32"

OI

See *operator instruction*.

open interval

The time interval during which a workstation is active and can process work.

opens file dependency

See "file dependency on page 16".

operation

A unit of work that is part of an application and that is processed at a workstation.

operation deadline

The latest time when the operation must be complete.

operation latest out

For an operation that has predecessors, the latest out date and time are the latest start time for the first critical operation in the application occurrence. If the first critical operation has not started by this date and time, then the operation is flagged as late, because it will be impossible for it to start on time based on the sum of the planned durations of all the operations on its critical path.

operation number

The number of the operation. This uniquely identifies each operation in an application.

operator instruction (OI)

An instruction that an operator can view when the operator must manually intervene in operations.

optman

A command-line program that maintains the global options in the product database.

origin date

The date that a period (cyclic or noncyclic) starts on.

owner ID

An identifier that represents the application owner.

P

parallel operations

Operations that are not dependent on one another and that can, therefore, run at the same time.

parameter

An entity that enables job instance-specific values to be substituted in job and job stream scripts, either from values in the database or at run time. Parameters cannot be used when scripting extended agent jobs.

pending application description

An application description that is incomplete and not ready for use in planning or scheduling. See active application description.

pending occurrence

The dummy occurrence created by the daily planning process to honor a dependency that has been resolved in the long-term plan but cannot be resolved in the current plan because the input arrival time of the predecessor is not within the current plan end time.

pending predecessor

A predecessor dependency to an occurrence which is defined in the long-term plan but not yet included in the current plan. See also *pending occurrence*.

period

A time period defined in the IBM® Z Workload Scheduler calendar.

plan

The means of scheduling jobs. Objects in the database become instances in the plan. See also:

- "database on page 11"
- "final job stream on page 17"
- "forecast plan on page 17"
- "JnextPlan on page 21"
- "planman on page 28"
- "preproduction plan on page 29"
- "production plan on page 29"
- "trial plan on page 36"

See also *current plan*.

planman

An application you use to create, extend, and reset plans of all types. See also "plan on page 28".

planned start time

The time that IBM Workload Scheduler estimates a job instance will start. This estimate is based on start times of previous instances of the job. See also:

- "actual start time on page 4"
- "earliest start time on page 14"
- "latest start time on page 23"
- "scheduled time on page 32"

pool

A logical workstation on which you group a set of dynamic agents with similar hardware or software characteristics. See also dynamic agent on page 13.

predecessor

A job or job stream that must complete successfully before successor jobs or job streams can be started. See also "successor on page 34".

predefined prompt dependency

A prompt dependency that is defined in the database and can be associated to any job or job stream. See also "prompt dependency on page 30".

print workstation

A workstation that prints output and usually reports status to IBM® Z Workload Scheduler automatically.

priority

A way of determining the order in which jobs and job streams start. In IBM Workload Scheduler, priorities for each job and job stream range from *0* to *101*. A job or job stream with a priority of *0* does not run. In IBM® Z Workload Scheduler, the priority of an operation is a value from 1 to 9 (where 1=low, 8=high, and 9=urgent). It is one of the factors that determines how IBM® Z Workload Scheduler schedules applications.

preproduction plan

A high-level plan of system activity containing job streams and dependencies. It is created automatically when the production plan is created for the first time. It is extended if the production plan is extended. It is similar to the *long-term plan* used in IBM Workload Scheduler for z/OS®. See also "plan on page 28".

processes

Network processes that control the production environment and network traffic. See also:

- "batchman on page 7"
- "jobman on page 22"
- "jobmon on page 22"
- "mailman on page 24"
- "netman on page 25"
- "writer on page 39"

production period

The time frame covered by the production plan. See also "production plan on page 29".

production plan

Contains all job scheduling activity planned for a period. The plan is created or extended by the **Jnextplan** job or by **planman**. It is stored in the `Symphony` file, and consists of all the jobs, job streams, and dependency objects that are scheduled to run for that period, including any jobs or job streams carried forward from the previous plan. See also:

- "carry forward on page 8"
- "JnextPlan on page 21"
- "plan on page 28"

program interface (PIF)

In IBM® Z Workload Scheduler, an interface that lets user-written programs issue various requests to IBM® Z Workload Scheduler.

prompt dependency

A dependency where an operator must respond affirmatively to a prompt so that the dependent job or job stream can run. See also:

- "ad hoc prompt dependency on page 5"
- "predefined prompt dependency on page 29"

Q

query current plan (QCP) dialog

A dialog that displays information taken directly from the current plan. The information includes information on operations, workstations, and application occurrences.

R

ready list

A display list of all the operations ready to be processed at a workstation. Ready lists are the means by which workstation operators manually report on the progress of work.

remote engine workstation

A workstation that represents locally a remote IBM Workload Scheduler engine. It is used to run only shadow jobs.

See also:

- cross dependency on page 10
- shadow job on page 33

remote job

A job that runs in a remote scheduling environment and is mapped by a shadow job to become a dependency for a job that runs in a local environment. See also shadow job on page 33

remote job tracking

The function of tracking jobs on remote processors connected by network links to a controller. This function enables a central site to control the submitting, scheduling, and tracking of jobs at remote sites.

remote processor

A processor connected to the IBM® Z Workload Scheduler host processor via a network. An IBM® Z Workload Scheduler event writer and an event transmitter are installed on the remote processor and transmit events to the IBM Workload Scheduler for z/OS host processor through the network link.

replan current period

In IBM® Z Workload Scheduler, a function that recalculates planned start times for all occurrences to reflect the actual situation.

reporting attribute

A code that specifies how a workstation will report events to IBM® Z Workload Scheduler.

reroutable

IBM® Z Workload Scheduler can reroute operations if the workstation that they are scheduled to run on is inactive. An example of this can be if communication links to the system where the workstation is located fail. This option applies to operations only when they have status R (ready) or W (waiting).

rerun

In IBM® Z Workload Scheduler, a function that lets an application or part of an application that ended in error be run again.

resource

Either physical or logical system resources. Resources are used as dependencies for jobs and job streams. See also "resource dependency on page 31".

resource dependency

A dependency where a job or job stream cannot start until the required quantity of the defined resource is available. See also "resource on page 31".

restartable

If an operation is defined as restartable, IBM® Z Workload Scheduler can automatically restart that operation if the workstation that it is using becomes inactive. This option applies only to the operation while it has status S (started). The operation will be reset to status R (ready).

restart and cleanup

A recovery function that ensures the restart of a job and the related cleanup actions, for example, deleting or uncataloging data sets created in a job run.

return code

An error code that is issued by IBM Workload Scheduler for z/OS for automatic-reporting workstations.

row command

An ISPF dialog command used to manipulate data in a table.

rule

A named definition of a run cycle that determines when an application will run.

rule-based run cycle

A run cycle that uses rules based on lists of ordinal numbers, types of days, and common calendar intervals (or period names in IBM Workload Scheduler for z/OS®). For example, the last Thursday of every month. Rule-based run cycles are based on conventional periods, such as calendar months, weeks of the year, and days of

the week. In IBM Workload Scheduler for z/OS®, run cycles can also be based on periods that you define, such as a semester. This term is only used as such in IBM Workload Scheduler for z/OS®, but the concept applies also to the distributed product. See also:

- "offset-based run cycle on page 26"
- "run cycle on page 32"

run cycle

A schedule that specifies when a job stream runs. See also:

- "calendar on page 7"
- "exclusive run cycle on page 15"
- "iCalendar on page 19"
- "inclusive run cycle on page 20"
- "rule-based run cycle on page 31"
- "simple run cycle on page 33"
- "weekly run cycle on page 37"

S

schedule

See "job stream on page 22". In IBM® Z Workload Scheduler, the current or long-term plan. To determine the input arrival date and time of an occurrence or operation.

scheduled time

The time when a job or job stream is scheduled to run. See also:

- "actual start time on page 4"
- "earliest start time on page 14"
- "latest start time on page 23"
- "planned start time on page 28"

security file

The file where access rights of users to objects in the database and the plan are defined. It is created by **makesec**. See also "makesec on page 24".

server

In IBM® Z Workload Scheduler, the optional component that runs on the controlling system and handles requests from remote ISPF dialogs, remote PIF applications, and the Graphical User Interface for Application Description.

service functions

Functions of IBM® Z Workload Scheduler that let the user deal with exceptional conditions, such as investigating problems, preparing APAR tapes, and testing during implementation.

shadow job

A job that runs locally that is used to map another job instance running on a remote engine workstation.

See also:

- job instance on page 21
- remote engine workstation on page 30

shared resource

A special resource or workstation resource that can be used simultaneously by more than one operation.

simple run cycle

A specific set of user-defined days a job stream is run. A simple run cycle is defined for a specific job stream and cannot be used by other job streams. See also "run cycle on page 32".

slack

Refers to spare time. This extra time can be calculated for the critical path by taking Deadline less the Input Arrival less the sum of Operation Durations.

smoothing factor

A value that controls the extent to which actual durations and deadlines are fed back into the application description database.

special resource

A resource that is not associated with a particular workstation, such as a data set.

splittable

Refers to a workstation where operations can be interrupted while being processed.

standard

User-specified open intervals for a typical day at a workstation.

standard agent

An installed agent component in a distributed IBM Workload Scheduler network that runs jobs, but requires a domain manager to resolve local dependencies and launch the jobs.

started-task computer workstation

You can specify that a computer workstation will support started tasks by giving the workstation the STC option. Operations defined to this workstation will be treated as started tasks, not as jobs.

started-task operations

Operations that start or stop started tasks. These operations are run at a computer workstation with the STC option specified.

status

The current job or job stream status within the Dynamic Workload Console. The Dynamic Workload Console status is common to IBM Workload Scheduler and IBM Z Workload Scheduler. See also "internal status on page 20".

status code

Codes that represent the current state of an operation. The status code is often associated with an extended status code.

step condition dependency

A specific check of the return code of a single step belonging to the conditional predecessor job within a condition.

step dependency

A specific check of the return code of a single step.

submit/release data set

A data set shared between the IBM® Z Workload Scheduler host and a local IBM® Z Workload Scheduler processor that is used to send job-stream data and job-release commands from the host to the local processor.

subresources

A set of resource names and rules for the construction of resource names. IBM® Z Workload Scheduler uses these names when checking a user's authority to access individual IBM® Z Workload Scheduler data records.

subsystem

A secondary or subordinate system, usually capable of operating independently of, or asynchronously with, a controlling system.

successor

A job that cannot start until all of the predecessor jobs or job streams on which it is dependent are completed successfully. See also: "predecessor on page 29".

Symphony file

A file containing the scheduling information needed by the production control process (**batchman**) to run the plan. The file is built and loaded when the production plan is created or extended on the master domain manager. During the production phase, it is continually updated to indicate the current status of production processing: work completed, work in progress, and work to be done. To manage production processing, the contents of the `Symphony` file (plan) can be displayed and altered using **conman** or the Dynamic Workload Console. See also:

- "batchman on page 7"
- "conman on page 10"
- "plan on page 28"

In IBM® Z Workload Scheduler where end-to-end scheduling is configured, the file is produced during daily planning and is sent to distributed agents to run their part of the plan.

SYSOUT

A system output stream, also an indicator used in data definition statements to signify that a data set is to be written on a system output unit.

SYSOUT class

An indicator used in data definition statements to signify that a data set is to be written on a system output unit. It applies only to print workstations.

T

table view

A graphical view in the Dynamic Workload Console used to display database and plan object data in tabular format. See also:

- "database on page 11"
- "plan on page 28"
- "views on page 37"

tail plan

Created during the daily planning process, includes only tail work; that is, work that started during or before the current planning period and that extends beyond its end.

task

In the Dynamic Workload Console, a filter, by scheduling object type, which returns a list of objects with attributes matching those specified in the task definition.

temporary operator instructions

Operator instructions that have a specific time limit during which they are valid. They will be displayed to the workstation operator only during that time period.

time dependent

IBM® Z Workload Scheduler attempts to start operations as soon as possible, when all dependencies have been resolved and processing resources are available. However, you can specify that an operation is time-dependent, so IBM® Z Workload Scheduler will not start it until a specific time.

timeline view

A graphical view in the Dynamic Workload Console used to modify and maintain job stream instance time restrictions. See also:

- "time restriction on page 36"
- "views on page 37"

time restriction

Determines the times before which, after which, or both, that a job or job stream cannot be run. Specifying both defines a time frame within which a job or job stream runs. Jobs can also have a repetition rate. For example, IBM Workload Scheduler can launch the same job every 30 minutes between the hours of 8:30 a.m. and 1:30 p.m.

time zone support

A feature of IBM® Z Workload Scheduler that lets applications be planned and run with respect to the local time of the processor that runs the application. Some networks might have processors in different time zones. The controlling processor will make allowances for differences in time during planning activities to ensure that interacting activities are correctly coordinated.

tracker

In IBM® Z Workload Scheduler, a component that runs on every system in your complex. It acts as the communication link between the z/OS system that it runs on and the controller.

tracking event log

A log of job-tracking events and updates to the current schedule.

transport time

The time allotted for transporting materials from the workstation where the preceding operation took place to the workstation where the current operation is to occur. The transport time is used only for planning purposes. Operations will be started irrespective of the transport time specified.

trial plan

A projection of the current production plan for a different period, using the same start date. It is used to determine the effect of different plan decisions. See also "plan on page 28".

turnover

A subfunction of IBM® Z Workload Scheduler that is activated when IBM® Z Workload Scheduler creates an updated version of the current plan.

U

user options

Configuration options defined for each user on a workstation, in a `useropts` file for each user on a workstation. The settings in this file apply only to that user on that workstation. See also:

- "local options on page 23"
- "global options on page 18"

utility commands

A set of utilities invoked from the operating system's command line for managing IBM Workload Scheduler.

V

validity period

The time interval defined by an origin date and an end date within which a run cycle or an application description is valid.

version

Applications with the same ID but different validity dates.

views

Elements of the graphical user interface of the Dynamic Workload Console used for viewing and modifying scheduling objects. See also:

- "explorer view on page 15"
- "graph view on page 18"
- "impact view on page 19"
- "table view on page 35"
- "timeline view on page 36"

W

waiting list

A list of jobs that have been submitted but still have uncompleted predecessors. Operations will be included in the waiting list if the JCL is not submitted by the controller and the tracker has been started with HOLDJOB(YES).

weekly run cycle

A run cycle that specifies the days of the week that a job stream is run. For example, a job stream can be specified to run every Monday, Wednesday, and Friday using a weekly run cycle. A weekly run cycle is defined for a specific job stream and cannot be used by multiple job streams. See also "run cycle on page 32".

work day

A day on which applications can normally be scheduled to start.

work-days-only cyclic period

A cyclic period where only work days are counted when calculating the interval.

work-day end time

The time when one work day ends in IBM® Z Workload Scheduler and the next day begins. By default, this time is midnight.

For example, if the work-day end time is 02:00, work for Friday can continue until 02:00 on Saturday morning, even if Saturday is a free day. If Saturday and Sunday are free days, no new work will be started until 02:00 on Monday.

workload service assurance

An optional feature that you use to flag jobs as mission-critical for your business and to ensure that they are processed in a timely manner. The progress of jobs and job streams on the critical path can be monitored and events raised if the planned finish time is at risk.

workstation

A definition of an individual computer or computer partition on which jobs and job streams are run. Types of workstation vary depending on the type of engine. See also:

- "distributed workstation on page 13"
- "z/OS® workstation on page 39"

In IBM® Z Workload Scheduler, a unit, place, or group that performs a specific data processing function. A logical place where work occurs in an operations department. IBM® Z Workload Scheduler requires that you define the following characteristics for each workstation: the type of work it does, the quantity of work it can handle at any particular time, and the times it is active. The activity that occurs at each workstation is called an operation.

workstation class

A workstation class is a group of workstations with similar job-scheduling characteristics. Any number of workstations can be placed in a class. Job streams and jobs can be assigned to run on a workstation class. This makes replication of a job or job stream across many workstations easy. See also "workstation on page 38"

workstation description database

In IBM® Z Workload Scheduler, a database containing descriptions of the workstations in the operations department.

workstation resource

A physical resource, such as a tape drive, that must be allocated among jobs. When you define a workstation, you can specify the quantity of each of two resources (R1 and R2) that are available to operations. When defining operations to that workstation, you can specify the number of these resources that must be available for the operation to start on that workstation.

workstation type

Each workstation can be one of three types: computer, printer, or general.

writer

A process started by **netman**. The **writer** process passes incoming messages to the local **mailman** process. The **writer** processes (there may be more than one on a domain manager workstation) are started by **link** requests and are stopped by **unlink** requests or when the communicating **mailman** process ends. See also "processes on page 29".

write-to-operator workstation

A general workstation that lets you use scheduling facilities to issue a write-to-operator (WTO) message at a specific operator console defined by the workstation destination.

WTO message

Write-to-operator message.

WTO operation

An operation that consists of an operator instruction that IBM® Z Workload Scheduler passes to NetView. These operations are run at a general workstation with the WTO option specified.

X

x-agent

See "extended agent on page 16".

Z

z/OS® network

A connected group of workstations that use the IBM Z Workload Scheduler engine to perform workload scheduling. See also:

- "engine on page 14"
- "workstation on page 38"

z/OS® workstation

A representation of system configuration elements in the IBM Z Workload Scheduler network. For the z/OS® engine, workstations can be:

- Computer
- General
- Printer

See also "workstation on page 38".

Notices

This document provides information about copyright, trademarks, terms and conditions for product documentation.

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.

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.