



SEMI E42-0998

RECIPE MANAGEMENT STANDARD: CONCEPTS, BEHAVIOR, AND MESSAGE SERVICES

1 Introduction

This standard defines the concepts required for management of recipes, the operations or behavior provided by the Recipe Management Standard (RMS), and the messages through which services are provided through an interface between the provider and the user of these services.

1.1 *Purpose* — The purpose of this standard is two-fold:

- To enable applications software to be developed that can assume the existence of standard concepts, behaviors, and message services that collectively form Recipe Management and that take advantage of them.
- To enable software to be developed to offer the Recipe Management capabilities.

1.2 *Scope* — This is a provisional standard that defines concepts, behavior, and services to support the integration of automated recipe management within a semiconductor factory. These services are applicable to a variety of relationships, including both traditional **host/equipment** and cluster tool controller/attached module communications and control.

The standard provides a set of communications services which allows such systems to transfer and manage recipes to ensure the correct processing of material within semiconductor manufacturing *equipment* and systems. RMS also requires compliance to SEMI E39 (Object Services Standard (OSS): Concepts, Behavior, and Services) for completeness.

This document describes several different hierarchical relationships: supervisory agents and their supervised agents, recipes and their subrecipes, and recipe classes and their subclasses. Such hierarchical relationships provide a natural organizational and classification structure that is reflected in many different kinds of systems, such as telephone switching systems and directory trees. It is the intent of this standard to support logical hierarchical relationships rather than to require that strict hierarchical relationships be implemented in systems architecture.

RMS places no restriction on where the set of defined services is implemented.

1.3 Referenced Documents

1.3.1 *Semiconductor Equipment and Materials International (SEMI)*¹

1.3.1.1 *SEMI Equipment Automation/Hardware Volume*

SEMI E10 — Standard for Definition and Measurement of Equipment Reliability, Availability, and Maintainability (RAM)

1.3.1.2 *SEMI Equipment Automation/Software 2 Volume*

SEMI E30 — Generic Model for Communications and Control of SEMI Equipment (GEM)

SEMI E39 — Object Services Standard: Concepts, Behavior, and Services

SEMI E53 — Event Reporting

1.3.2 Other References

James Rumbaugh, Michael Blaha, William Premerlani, Frederick Eddy, William Lorensen, *Object-Oriented Modeling and Design*, Englewood Cliffs, New Jersey: Prentice-Hall, 1991.

D. Harel, *Statecharts: A Visual Formalism for Complex Systems*, Science of Computer Programming 8, 1987.

1.4 *Definitions* — Basic definitions for objects, services, and form are provided in SEMI E39, Sections 4.1 through 4.4. This section provides additional definitions.

Definitions in Section 1.4.3 are specific to RMS.

1.4.1 Services

service provider — An application (a component of an *agent*) responsible for providing services to the service user.

¹ Semiconductor Equipment and Materials International (SEMI), 805 East Middlefield Road, Mountain View, CA 94043, 650.964.5111, FAX 650.967.5375



service user — (service consumer) An application that uses the services provided.

1.4.2 Form

binary — A string of bit values (zeroes and ones), with a format that is either left unspecified or specified by bit position, with the most significant bit first. The total length of the string is a multiple of eight. Messaging protocol may impose restrictions on length.

1.4.3 Recipe Management — This section introduces basic terminology used in RMS. Entries are in alphabetical order. Additional definitions and specifications are provided in later sections.

agent — An intelligent system within a factory that provides one or more service resources and uses the services of other agents. This is a generalization that includes host, equipment, cell, cluster, cluster module, station controller, and work station. *Agents* are associated with a physical system or a collection of physical systems, such as computer platforms.

authorized user — A user who can be identified to an *agent* as having the level of authority required for a particular activity, such as *certifying* a recipe for that equipment.

collection event — A detectable occurrence of interest to a service user.

component agent — A subordinate *agent* that provides services to a *supervisory agent*.

download — An operation that transfers a recipe (down) to an execution storage area.

edit — An operation which creates a new recipe body or changes the body of an existing recipe.

editor — A service which allows a *user* to edit a recipe. *Editors* are not specified in RMS.

*equipment*² — An *agent* with associated hardware that provides, at a minimum, recipe **execution** services.

event — A detectable occurrence significant to an object.

execution (recipe execution) — The process of reading the recipe contents and implementing its instructions, process parameters, or other information required for its own processing.

executing agent — An *agent* that provides *recipe execution* capabilities.

execution area — The storage location of the recipe(s) currently *selected* (ready) for *execution*.

host — A *supervisory agent* that represents the factory to its subordinates.

logical recipe — A recipe with a particular set of *attributes* and a particular *body*, considered independently from its physical location. A *logical recipe* may have multiple instances or copies.

name — A text-based *attribute* of an object that may be used as all or part of its *identifier*.

namespace — In general, a domain within which object identifiers are unique. In RMS, the term *namespace* is used as a synonym for *recipe namespace*, unless otherwise stated.

operator — The user who interacts locally with *agent* through the *agent's* interface.

recipe — The pre-planned and reusable portion of the set of instructions, *settings*, and parameters under control of an *agent* that determines the processing environment seen by the manufactured object and that may be subject to change between runs or processing cycles.

recipe class — A formal grouping of recipes with a common language syntax and functionality.

recipe executor — The component of an *executing agent* that executes recipes.

recipe namespace — A logical management domain with the responsibility for the storage and management of recipes, the ensurance of the uniqueness of recipe **identifiers** within that domain, and the provision of services pertaining to recipes stored within that domain.

recipe parameter — A control value that affects the *agent's* process.

select — The act of preparing a recipe for execution.

setting — A static value accessible to the *user*, through one or more methods, that is used by equipment to control its process. *Settings* include, but are not limited to, setpoint values. *Settings* typically may be specified within a recipe.

storage area — An area where objects and data are stored.

² The term "equipment" is restricted in RMS to "intelligent equipment."



subordinate agent — An *agent* that is a component of, or managed by, another *agent*.

supervisory agent — An *agent* with supervisory responsibilities for one or more subordinate *agents*.

timestamp — The notation of the date and time of the occurrence of an event.

upload — An operation that transfers a recipe (up) from an execution storage area.

user — A person interacting with an *agent* directly through the *agent's* human interface or indirectly through the *agent's* supervisor.

validate — The action of checking recipe contents to ensure that parameter type and range are valid for the equipment configuration prior to execution. [Note that *validation* and *verification* are used in different ways.]

variable parameter — A formally defined variable (*setting*) defined in the body of a recipe permitting the actual value to be supplied externally.

verify — The operation of reading a recipe's contents to ensure that it is syntactically correct and identifying elements that must be made public.

version — Part of a recipe's *identifier* that is used to show its heritage.

1.5 Conventions — The following conventions are used in this document:

- To highlight terms specific to RMS (excluding terms defined in Sections 1.4.1 through 1.4.3 that are common to multiple standards), a defined term appears in **boldface** wherever it first appears and wherever it is defined. This alerts the reader to those terms with specific meanings. Except for terms that are very common, such as recipe, host, and equipment, defined terms otherwise are in *italics* wherever they appear.
- Terms related to objects and object services are in conformance with SEMI E39 (Object Services Standard: Concepts, Behavior, and Services). A brief discussion of objects is provided in the Appendix of that document.
- Attribute names are underlined.
- Attributes called "names" in RMS are generally intended to be used for the ObjID attribute of a *standardized object*.

- To prevent the definition of numerous message parameters named "XxxList," this document adopts the convention of referring to the list as "(List of Xxx)". In this case, the definition of the parameter Xxx will be given, not of the list. The term "list" indicates a collection (or set) of zero or more items of the same data type.

For attributes that are lists, this convention is not followed, as the entire attribute, as a list, must be assigned a specific name.

1.5.1 Text String Restrictions — Text strings used in attribute names, attribute values, or message parameters, are subject to the restrictions defined by OSS: Text in ASCII is restricted to the characters between 20₁₆ and 7D₁₆, excluding the question mark "?" (3F₁₆), the asterisk "*" (2A₁₆), and the tilde "~" (7E₁₆).

Text strings used as, or within, the object identifier ObjID are additionally restricted to exclude the "greater than" symbol ">" (2E₁₆) and the colon character ":" (3A₁₆) to conform with OSS requirements.

1.5.2 Harel State Model — This document uses the Harel State Chart notation to describe the dynamic behavior of the objects defined. An overview of this notation is presented in an Appendix of SEMI E30. The formal definition of this notation is presented in Science of Computer Programming 8, "Statecharts: A Visual Formalism for Complex Systems," by D. Harel, 1987.

This document also adopts the extension of Harel notation to show the deletion of an object as used by Rumbaugh, et al. (see Section 1.5.2.1).

1.5.3 Objects — *Standardized objects* defined by RMS conform to the requirements of SEMI E39 (Object Services Standard: Concepts, Behavior, and Services). RMS adopts the convention of showing *standardized objects* as drawn with a heavy line in object models. *Non-standardized objects* are used to illustrate concepts and relationships but are not formally defined and cannot be accessed through Object Services. A list of *standardized objects* defined by RMS is included in Related Information 1.

1.5.3.1 OMT Object Information Model — The object models are presented using the Object Modeling Technique, developed by Rumbaugh et al., in *Object-Oriented Modeling and Design*, Prentice-Hall, Englewood Cliffs, NJ, 1991. Overviews of this notation are pro-



vided in an appendix of SEMI E39 (Object Services Standard: Concepts, Behavior, and Services).

1.5.3.2 Object Attribute Representation — The object information models for *standardized* objects will be supported by an **attribute definition** table with the following column headings:

<i>Attribute Name</i>	<i>Definition</i>	<i>Access</i>	<i>Rqmt</i>	<i>Form</i>
The formal text name of the attribute.	Description of the information contained.	RO or RW	Y or N	(see below)

The Access column uses RO (Read Only) or RW (Read and Write) to indicate the access that users of the service have to the attribute.

A 'Y' or 'N' in the Requirement (Rqmt) column indicates if this attribute must be supported in order to meet fundamental compliance for the service.

The Form column is used to indicate the format of the attribute. (See Section 1.4 for definitions.)

1.5.4 Service Message Representation

Service Resource Definition

A **service definition** table defines the specific set of messages for a given *service resource*, as shown in the following table:

<i>Message Service Name</i>	<i>Type</i>	<i>Description</i>
Message name	N or R	The intent of the service.

Type can be either N = Notification or R = Request.

Notification type messages are initiated by the service provider, and the provider does not expect to get a response from the user.

Request messages are initiated by a service user. Request messages ask for data or an activity from the provider. Request messages expect a specific response message (no presumption on the message content).

Service Parameter Dictionary

A **service parameter dictionary** table defines the parameters for one or more services, as shown in the following table:

<i>Parameter</i>	<i>Form</i>	<i>Description</i>
Parameter X	Data type	A parameter called X is B in A.

A row is provided in the table for each parameter of the service. The first column contains the name of the parameter. This is followed by columns describing the form and contents of the corresponding primitive.

The Form column is used to indicate the type of data contained in a parameter. (See Section 1.4 for definitions.)

The Description column in the Service Parameter Dictionary table describes the meaning of the parameter, the values it can take on, and any interrelationships with other parameters.

To prevent the definition of numerous parameters named "XxxList", this document adopts the convention of referring to the list as "(List of)Xxx". In this case, the definition of the variable Xxx will be given, not of the list. The term "list" indicates a collection (or set) of zero or more items of the same data type. Where a list is used in both the request and the response, the list order in the request is retained in the response. A list must contain at least one element, unless zero elements are specifically allowed.

Service Message Definition

A **service message definition** table defines the parameters used in a service, as shown in the following table:

<i>Parameter</i>	<i>Req/Ind</i>	<i>Rsp/Conf</i>	<i>Description</i>
Parameter X	(see below)	(see below)	A description of the service.

The columns labeled Req/Ind and Rsp/Conf link the parameters to the direction of the message. The message sent by the initiator is called the "Request". When receiver terms this message the "Indication" or the "Request", the receiver may then send a "Response", which the original sender terms the "Confirmation".



The following codes appear in the Req/Ind and Rsp/Conf columns and are used in the definition of the parameters (e.g., how each parameter is used in each direction):

"M" — **Mandatory parameter** — must be given a valid value.

"C" — **Conditional parameter** — may be defined in some circumstances and undefined in others. Whether a value is given may be completely optional or may depend on the value of the other parameter.

"U" — **User-defined parameter**.

"-" — The parameter is not used.

"=" — (for Response only) Indicates that the value of this parameter in the response must match that in the primary (if defined).

1.6 Requirements — Requirements for recipe management are varied. For example, there is a fundamental need to retain information about each recipe, to provide that information upon request, and to transfer it with the recipe so that it remains available. There is a need to classify recipes, to reuse them, and to share them with multiple installations of *equipment*. The need for sharing recipes introduces a new requirement for safeguarding the relationship between a recipe's **identifier** and its contents, to be able to protect a recipe from unauthorized changes, and to ensure that the *identifier* for a recipe that is used by multiple *equipment* is in fact the same recipe.

The requirements for managing recipes in a factory are given in tabular form in Table 1.1, which presents them according to their functional areas and the specific issues involved. Table R1-1, in Related Information, shows the specific concepts defined in RMS to address each of these requirements.



Table 1.1 Requirements

<i>Functional Areas</i>	<i>Issue</i>	<i>Requirements</i>
Management of Recipes	Identification	Uniquely store, identify, and select recipes in a system.
		Easy and clear identification of a recipe.
	Multiple Access	Share recipes among <i>equipment</i> .
		Enable synchronized change of shared recipes among <i>equipment</i> .
	History and Traceability	Capture the history of recipe usage.
		Capture the history of recipe changes.
	Life Cycle Management	Manage approval of recipes for process development and production.
Operations	Protection	Protect recipes from unexpected changes.
		Protect recipes from mistakes in operations.
	Portability	Create, edit, and change recipes outside the <i>equipment</i> which execute them.
	Reusability	Change recipes for a specific piece of <i>equipment</i> .
		Use recipes developed on one piece of <i>equipment</i> for other <i>equipment</i> .
		Share recipes among the same kind of <i>equipment</i> by adjusting for individual differences.
Execution	Safety	Do not execute recipes that are not syntactically correct.
		Execute the specified (selected) recipes without errors.
		Protect recipes being executed from inadvertent change caused by other activities.
	Flexibility	Change a recipe's parameters during or between runs in a systematic way.
System Operation	System Node Operation	Dynamically connect and disconnect <i>equipment</i> to/from the communications network.
	Stand-Alone Operation	Execute recipes with no communications link.
		Create and change recipes with no communications link.
		Manage recipes moved between <i>equipment</i> or from off-line storage through removable media.
		Enable smooth integration of stand-alone <i>equipment</i> into on-line factory systems.



1.7 *Document Structure* — Figure 1.1 depicts the domain of the Recipe Management Standard as consisting of Purpose, Concepts, Behavior, and Message Services. This also reflects an underlying structure of the document.




Figure 1.1
Recipe Management Domain

Purpose provides the motivation for Recipe Management capabilities and is addressed in Section 1.6. Concepts provide a detailed introduction to the *standardized objects* of RMS, their attributes, and their relationships with other objects. Behavior describes the operations that are performed by, or on, these objects. Finally, message services define the messages and their parameters independently of the protocol in which they are implemented. Both concepts and behavior represent an "inside view" of an RMS application, while services provide an interface from an external view.

- Section: 1: the formal introduction to RMS.
- Section 2: an overview of the major objects of RMS and their relationships, to provide a general context for the technical detail that follows.
- Sections 3-6: concepts for recipes, *recipe namespace*, *distributed recipe namespace*, and *recipe executor*.
- Section 7: the concept of the *agent* and of *service resources*, to provide a more complete context for RMS implementations.
- Sections 8-11: behavior (operations) for **recipe management**, **namespace management**, **distributed recipe namespace management**, and the *recipe executor*.

- Sections 12-14: definitions of message services for the *recipe namespace*, *distributed recipe namespace*, and *recipe executor service resources*.
- Section 15: RMS compliance.
- Section 16: a glossary of terms, provided as a convenient reference.
- Related Information: provides background information on the requirements behind RMS and examples of applications.

1.8 *Applicable Documents*

ISO/TR 8509:1987, Information Processing Systems, Open Systems Interconnection — Service Conventions.

2 Overview of RMS

This section provides an introduction to, and overview of, the major objects of RMS.

2.1 *Recipe Management Models* — To provide a clearer understanding of the major entities or areas of functionality specified by RMS, they are portrayed as objects in RMS models, using OMT notation.

Relationship lines in the object-based models do not represent a direct communication link between objects, but rather the knowledge and association that one object has with respect to another.

The models do not indicate the relative location of two objects.

2.2 Major Objects in RMS — Recipe Management is concerned with the major areas of functionality illustrated in Figure 2.1: **recipes**³, **recipe namespaces** providing persistent recipe storage, **recipe namespace managers** that provide management of the *namespace* and access to its recipes, and the **recipe executor** that executes recipes.




Figure 2.1
Major Objects of RMS

³ The recipe object in Figure 2.1 is not itself a standardized object. However, two recipe subtypes defined in Section 3 are standardized objects.

2.2.1 Recipes — Recipes provide a flexible, manipulatable, and re-usable form for users to select sequencing and settings to effect a particular result. Equipment uses recipes in a variety of ways to control the processing environment, sometimes using several types of recipes together. Recipes may also be used for maintenance activities, such as calibration or cleaning.

Recipes can be modified and copied from one environment to another. Their inherent flexibility as a form also becomes an endless source of problems unless they can be managed. Misprocessing, unintentionally running the wrong recipe for a given product, is prohibitively expensive for factories.

RMS defines a recipe as an object with attributes — information about the recipe — as well as content. It is the attributes of the recipe that allow true management of recipes, throughout their lifecycles, to occur.

2.2.2 Recipe Namespace — A *recipe namespace* in RMS is analogous to a smart file directory. The *namespace* provides long-term recipe storage capability.

Basic operations are part of the *namespace* specification. For example, recipes can be copied from one *namespace* to another, copied within a *namespace*, renamed, created, updated, and deleted. They can be downloaded to, and uploaded from, an application called a *recipe executor*. Several recipes can be *linked* together to form a set.

The *namespace* also can serve as a recipe pool that can be shared by a group of equipment of a common type. The *distributed recipe namespace* is a *namespace* able to utilize and manage recipe storage provided by such an equipment group.

2.2.3 Recipe Namespace Manager — The *recipe namespace* itself is passive, a container for recipes. The **recipe namespace manager (manager)** provides the dynamic element that manages the *namespace*. The *manager* represents the interface for the *namespace* to the external world and the internal decision authority within the *namespace*.

2.2.4 Recipe Executor — The **recipe executor (executor)** is the component of an *agent* that understands the contents of a recipe and is able to **verify** their syntactical correctness and to **validate** that it can be executed under the current configuration. It **executes** a recipe by reading its contents and applying them appropriately in order to achieve a desired result.

The *recipe executor* also may provide additional limited storage capability, but it is restricted from modifying existing recipes, except under special conditions, to prevent unexpected and/or unwanted change.

2.3 Agent — As used in RMS, an **agent** is a system, or subsystem, in a factory, that has a physical aspect. It consists of one or more applications that provide and/or use *service resources*, such as a *recipe namespace manager* and/or a *recipe executor*, as illustrated in Figure 2.2.




Figure 2.2
An Example of an Agent and Component Applications

Agent is a generalization that covers traditional *equipment*, supervisory or aggregate systems such as stations, cells, and clusters, and intelligent subsystems within *equipment*, such as process modules within a cluster. Other types of *agents* may also be implemented, such as those dedicated to *recipe namespace management*.

An *agent* that provides a *recipe executor* is called an **executing agent**. **Equipment** is an *executing agent* with associated hardware to which the recipe applies and which it uses to do work.

2.4 Implementations — RMS can be implemented on different platforms within a factory. A host controller may only provide *recipe namespace management*, while a diskless cluster module may provide only *recipe execution*. Traditional *equipment* capable of operating in stand-alone mode is required to provide a *recipe namespace* that is available on powerup as well as a *recipe executor* component. Additional examples of implementations of RMS are provided in Related Information 1.

RMS defines the services provided through an external interface. Communications between or within applications that are internal to an *agent* are not covered by RMS, so long as the attributes and operations comply with RMS requirements.

3 Recipes

This section describes the concepts concerning recipes.

3.1 Motivations — Definitions of recipe structure, attributes, and operations address the following issues:

- maintenance of recipe attributes for management and traceability,
- provision of a standard method for transferring attributes separately, as well as with the recipe body,

- reduction or elimination of the need for dedicated recipe editors,
- support for applications requiring multiple recipe types, such as formats or language syntaxes,
- run-to-run control,
- support for sorting by factory approval and **certification levels**,
- the ability to tune a generic recipe for a specific installation of agent as well as the same agent over time, and
- the ability to share recipes across *executing agents* supporting the same recipe syntax or language(s).

3.2 Basic Concepts — There are two types of recipes addressed by RMS, the **managed recipe** that is stored in a *recipe namespace* and the **execution recipe** that is stored in **recipe execution storage**. This section describes the basic concepts concerning both types of recipes.

3.2.1 Types of Recipes — There are three types of recipe objects shown in Figure 3.1. Most of Recipe Management is concerned with the management of recipes stored in a *recipe namespace*. These recipes are called **managed recipes**. Recipes are also stored by the *recipe executor* in the **recipe execution storage**, and these are called **execution recipes**.




Figure 3.1
Recipe Types

The recipe supertype is an abstract type that shows the common elements and attributes of the *managed recipe* and the *execution recipe* and their differences. The rec-

ipe supertype is not itself intended to be implemented. That is, recipe object types formally defined in RMS are either *managed recipes* or *execution recipes*. As indicated by the light lines, the supertype is not a *standardized object*.

3.2.2 Recipe Structure — Both the *managed* and the *execution* recipe have a **body**⁴ and a set of attributes. The *body* (contents) of the recipe contains the data used by the *recipe executor* for its execution process. The *body* is also not a *standardized object* and may not be accessed through Object Services.

Figure 3.2 illustrates all of the attributes defined by RMS that are common to both the *managed recipe* and the *execution recipe*. Not all of these attributes are required for a minimum implementation of RMS. The rules for setting the values of certain attributes, however, are different for the *managed recipe* and the *execution recipe*.




Figure 3.2
Recipe Supertype

The **identification attributes**, those attributes used for purposes of identification, are discussed in Section 3.3. The other attributes are discussed in later sections specific to each recipe type.

The *managed recipe* also may maintain a set of attributes that are specific to a single *executing agent*. The **agent-specific** attributes are technically attributes of the association between the recipe and the *executing agent*. In RMS, however, these attributes are treated as attributes of a component of the recipe called the

⁴ The recipe's body corresponds to a "process program" in SEMI E5 and SEMI E30.



agent-specific dataset, in order to allow them to be managed within the *namespace* along with the other attributes of the recipe. *Agent-specific* attributes allow the recipe to be personalized for particular *executing agents* (i.e., particular *recipe executors*).

Because the *managed* recipe may be applied to multiple *agents*, it may have multiple *agent-specific datasets*. When a *recipe* is downloaded from a *recipe namespace* to a *recipe executor*, the *managed recipe* is transformed into an *execution recipe* by taking the *agent-specific* attributes of the *agent-specific dataset* for that *agent* (its *recipe executor*) and merging them into the attributes of the *execution recipe*.

Methods of storing a recipe or the components of a recipe are not dictated by RMS. However, the association between a recipe's *body* and its attributes shall be carefully maintained.

3.2.2.1 Recipe Body — The recipe **body** contains the reusable instructions, *settings*, *parameters*, and other data that the *recipe executor* reads to *execute* the recipe and control its operation.

A recipe *body* may be in one of two basic forms: **source** (text) or **object**. The purpose of the *source form* is to encourage the use of human-readable and human-editable text. The *object form* supports the types intended only to be read and applied by machines and other automated systems, such as those produced by vision systems or CAD programs, as well as proprietary formats created and used by the *recipe executor*.

3.2.2.1.1 Source Form — Recipes that are created and modified by a form of editor should be available in *source form*. The **source form** of a recipe *body* is equivalent to a text file, such as ASCII or JIS-8, that can be created, read, printed, or modified with any "standard" text editor⁵. The *executing agent* is not required to supply such a text editor, but its *recipe executor* should be able both to read and write recipes in *source form*. Editing is not covered by RMS.

The *source form* of a recipe may be copied from one *namespace* to another *namespace* and from a *namespace* to a *recipe executor's recipe execution storage*, and all *recipe executors* with access to a given

namespace are assumed to use the same recipe language and the same functionality.

Two restrictions are placed on the interpretation of the text transferred between agents:

- In ASCII, the valid characters are from space (20₁₆) through "~" (7D₁₆), plus the tab (09₁₆), linefeed (0A₁₆), carriage return (0D₁₆), and form feed (0C₁₆) characters. The space and tab are considered as **horizontal whitespace** characters, and the carriage return and form feed characters are considered **vertical whitespace** characters. The line-feed character is used to define the end of a line of text (eol).
- The meaning of a recipe should not depend on embedded control characters, or specific *whitespace* characters, or the existence of *whitespace* at the end of a line. *Horizontal whitespace* characters should be treated alike for purposes of interpretation, as should *vertical whitespace* characters. All other characters, aside from those specified here, are discouraged and, if found, should be removed from the text prior to transfer.

In general, the recipe language and syntax used within a recipe *body* are beyond the scope of recipe management. However, to reduce arbitrary variability for the *user*, the following "rules" are recommended:

- Case is not significant for the purpose of comparison or meaning but should be retained as encountered, since it is often used to enhance readability. For example, the tokens `alarmlevel`, `ALARM-LEVEL`, and `AlarmLevel` should all have the same meaning.
- Comments may be included in the text and are ignored in recipe interpretation. Two types of comments are defined: comments that begin with the character pair `"//"` continue to the end of the line, while comments that begin with the character pair `"/*"` are terminated only by the character pair `"*/"`. A comment of the second type is terminated by the first character pair `"*/"` that follows. Comments may not be nested. Nested comments are not honored and may generate a syntax error.
- The minimum line consists of a single line-feed character.

⁵ The intent is not to specify the text editor, but rather to allow a user to modify text on the basis of clear guidelines and restrictions that may apply. *Source form* recipes are highly desirable to reduce the proliferation of, and need for, proprietary editors that are dedicated and language-specific.



In addition, recipe languages that support **external references** to other recipes should use the text form defined for the recipe *identifier* (Section 3.2.3.4) and for the recipe *specifier* (Section 3.2.4.1.2).

The supplier of the *recipe executor* shall provide documentation that formally defines each recipe language that the *agent* supports. Two forms are required, one for the *user*, who must be able to write correct recipes, and one that covers both the syntactical and lexical structures of the language using a formal descriptive protocol such as the Backus-Naur Form (BNF). The second form of documentation allows parsers to be built that are capable of pre-checking a recipe for syntactic and lexical correctness prior to downloading it. Pre-checking improves efficiency and performance by reducing the demand for recipe checking by the *recipe executor*.

The final responsibility for determining the correctness of a recipe shall belong to the *recipe executor*.

3.2.2.1.2 Object Form — Recipes that cannot be translated into a meaningful *source form* include datasets derived from vision systems or mechanical systems through a special hardware-dependent "teach" operation. Such recipes may remain in a proprietary *object form* and may or may not be applicable to other *recipe executors* because of their hardware dependencies.

For greater efficiency, *recipe executors* may rewrite (e.g., **tokenize**) a *source* recipe to generate a proprietary *object form* of the recipe. This is called the **derived object form**, to distinguish it from *object form* recipes that never exist in *source form*. The *agent's* manufacturer may choose to make this form available for transfer as well. From the point of view of the *supervisor*, the *body* of a recipe that is in *object form* is an unstructured binary vector or string. In general, it cannot be edited on a *supervisor*, and the *source form* of the recipe is still required. *Derived object form* recipes that cannot be used by multiple *recipe executors* should not be stored in *recipe namespaces* that are shared.

Where both the *source form* and the *derived object form* exist, the *identifier* for the *object form* of the recipe shall be different from that of the *source form* in some recognizable way. The supplier of the *recipe executor* shall provide a documented method for distinguishing one from the other within the *identifier*.

This method must be such that the relationship of the *derived object form* recipe from the original *source* recipe is obvious to the user. (See Section 3.2.3.4 for examples.) An attribute of the recipe, SrcRepID, is provided to retain the relationship between the *source form* recipe and the *derived object form* recipe.

3.2.3 Recipe Identification — The ability to properly and unambiguously identify a recipe is critical to both recipe management and processing. It is also important for the user that identification be logical. RMS defines three elements used to uniquely identify a recipe within any *recipe namespace*. These three elements are the recipe's **class**, **name**, and version. Each element is expressed as a text string, and they are concatenated together to form the **recipe identifier** (see Section 3.2.3.4).

3.2.3.1 Recipe Name — The **recipe name** is a user-defined text string that may be used to encode the technology, layer, manufacturing area, and other characteristics. The *recipe name* alone is not necessarily unique within either a given class or a given *namespace*.

The *recipe name* is subject to the conventions for text as defined in Section 1.5.1.

3.2.3.2 Recipe Class — A **recipe class** is a formal grouping of recipes that have a common syntax. A class may contain **subclasses** where the recipes within the *subclasses* operate in different environments or have different syntaxes from the parent *class* and/or from one another. A *subclass* is itself a *class* and may contain further *subclasses*.

The overall model for recipes is a set of *class* hierarchies or "*class trees*" defined by the supplier of the *recipe executor*. The major *classes* of recipes are the **PROCESS class**, the **SERVICE class**, and possibly other *agent-specific classes*. These are called **primary classes**. A **primary class** is a *class* that is not a *subclass* of another *class*.

The **PROCESS class** is a required *primary class* and is composed of recipes whose primary purpose is to increase the manufactured value of the production material⁶.

⁶ Material handling systems, wafer metrology, wafer measurement, wafer inspection systems, as well as process equipment which directly changes the characteristics of the material, increase the value of the material in some way.



The SERVICE class is an optional *primary class* used for recipes whose purpose is to maintain, prepare, calibrate, or test the operation of *equipment*. Typical *subclasses* of the SERVICE *class* might be CALIBRATION and CLEANING. SERVICE recipes that use the same recipe language as normal process recipes are not required to be placed into a separate *primary class*. The SERVICE class allows separation and use of a different syntax for special non-process purposes.

Other *primary classes* that do not fit the PROCESS or SERVICE may be provided by the supplier of the *recipe executor*. Such additional *classes*, and their function, shall be documented.

Recipes that are not in the PROCESS *class* or one of its *subclasses* shall not be used for production purposes. However, recipes within the PROCESS class may be used for purposes other than production.

Figures 3.3 and 3.4 show examples of object models of possible *classes* and *subclasses* for a wire-bonder and for a furnace with an automatic boat-loader subsystem.




Figure 3.3
Wire-Bonder Recipe Classes

The text form of a recipe *class* shall conform to the convention for text (Section 1.5.1). Within the text form of a recipe *identifier*, a class is delimited on both ends by the forward slash character "/". The complete *class* specification within recipe *identifier* is formed by the concatenation of *classes* and *subclasses*, starting with the *primary class* and ending with the particular *class* of the recipe in question:

"/PRIMARY CLASS/SUBCLASS₁/SUBCLASS₂/...".



For example, a recipe named DryOx, version 4, within the NORMAL CYCLE *class* in Figure 3.4., has an *identifier* of

"/PROCESS/FURNACE/DIFFUSION/NORMAL CYCLE/DryOx;4".




Figure 3.4
Recipe Classes for Furnace

The complete *class* specification, which always starts with the *primary class*, shall always be accepted. If *class names* within a *namespace* are all unique, however, the *service* provider may also accept a *class* designator consisting of a single *class*. The *identifier* in the example above then becomes

"/NORMAL CYCLE/DryOx;4".

3.2.3.3 Version — A recipe may evolve over time and exist in several versions. This allows the user to retain a recipe *name* over multiple versions and show the recipe's heritage. Different versions are identified by the **version** portion of the recipe's *identifier*. NOTE: More than one version of a recipe may be in use for production at the same time, not only within the factory, but also within a single equipment.

A **version** is a text string consisting of at least one character. A new *version* is either generated automatically by the *recipe namespace* or is assigned by the *user*. A *user* may assign any combination of text characters and punctuation marks to the *version*, except for:

- the characters prohibited by the convention for text usage and the object identifier (Section 1.5.1),
- *whitespace characters*.

It is recommended, but not required, that only upper case be used.



Numeric versions (version numbers) consist only of the digits "0" through "9" and one decimal point character "." and can be translated to a pure number. To avoid confusion and multiple *versions* with the same numeric value, *numeric versions* are further restricted as follows:

- whole numbers (with no decimal point) may not start with a zero "0" followed by another digit, and
- decimal numbers (with a decimal point) may not start or end with the decimal point or end with a zero following the decimal point.

For example, *version numbers* "09", "1.", ".5", and "1.670" are prohibited. The proper forms with the same numeric values are "9", "1", "0.5", and "1.67". The *user* may assign a *version* of "0" but not of "00".

Versions that are assigned automatically have additional restrictions:

- They shall be *numeric versions*, excluding the decimal point and with a minimum value of "1".
- *Versions* shall be assigned incrementally. When assigning a *version* for a recipe with a given *class* and *name*, if no other recipe exists within the *namespace* with that *class* and *name*, then a *version* of "1" is assigned. Otherwise, the highest *existing version* already in use for that *class* and *name* is determined, and the new *version* is assigned a value equal to that value plus 1. For example, if the highest *version* in use has a numeric value of 5, then the next *version* assigned would be "6".

To compare *two versions*, they are converted to uppercase and compared character by character.

3.2.3.4 Recipe Identifiers — The **recipe identifier** is formed from the concatenation of the recipe's *class*, *name*, and *version*, in that order:

" /CLASS₁/CLASS₂/.../CLASS_n/NAME;VERSION"

where CLASS₁ is a *primary class* and CLASS_{i+1} is a *subclass* of CLASS_i. The *recipe name* follows *class*. *Name* and *version* are always separated by a semicolon ";" (3B₁₆).

Where CLASS_n is a unique class name, the form becomes

/CLASS/NAME;VERSION".

The *recipe identifier* is used for the OSS-required attribute **ObjID** and shall conform to restrictions imposed on **ObjID** (see SEMI E39). The total length of

ObjID may have additional restrictions imposed by the protocol.

NOTE: There is no necessary relationship between a recipe's *identifier* or *name* and any file name(s) under which the *body* and attributes may be stored internally. The *recipe identifier* is a logical reference to the recipe that is independent of specific platforms and implementations. In particular, file services provided by operating systems may have naming restrictions, such as length, that are incompatible with the requirements of RMS. A recipe may be stored internally in different ways, such as in a single flat file, a set of related files, or a database. Actual storage methods shall be invisible to RMS.

Where suppliers supporting recipes in *source form* also rewrite them, as discussed in Section 3.2.2.1.2, a method based on the *recipe identifier* (**ObjID**) is required to both distinguish between the two *forms* and to recognize the relationship between the original *source form* recipe and the derived *object form* recipe. The method may use any of the three elements of the *identifier* for this purpose. For example, recipes in *object form* may be placed in a separate *subclass* called "/OBJ/", or a suffix such as ".obj" might be appended to the user-defined *name* of the original recipe.

3.2.3.4.1 Default Recipe Identifiers — It is always possible to reference a recipe by specifying its *full identifier* (*class*, *name*, and *version number*). In certain cases, described in Section 3.2.4.1.1, it may also be necessary to specify the *namespace* of a recipe. However, it is not always necessary to specify all components of the *identifier*. Specifically, *class* and/or *version* may be omitted when a recipe is *selected* for execution and within an *external reference*. *Namespace* is not normally specified. Rules define default values for *namespace* and *class*, and there are rules for determining the appropriate *version*.

The *recipe name* must always be specified.

For *external references* within a recipe, the default *class* is the *class* of the recipe making the reference. The proper *class* is determined when the recipe is *linked* (see Section 3.2.4.1.3).

Where a *recipe version* is unspecified, the recipe of that *class* and *name* having the highest *approval level* will be used. Where several recipes are found with the same *approval level*, the one with the highest *version* will be used. Where the *version* is unspecified in an



external reference, the correct *version* is determined when a **linked recipe set** is built by the **link** operation.

For references with *managed recipes*, if no *namespace* is specified, the default is *namespace* of the referring recipe.

3.2.4 Advanced Recipe Capabilities

3.2.4.1 Multi-Part Recipes — Use of multi-part recipe is often dictated by the *recipe executor*, which may require a set of recipes of different *classes* for its process. In addition, support for multi-part recipes can be helpful for the user, in particular where recipes are otherwise long. The recipe language may define use of subrecipes analogous to that of subroutines in ordinary programming languages. This enables better re-use where a high degree of similarity exists between different recipes.

Where recipes of more than one *class* are required, it is convenient to allow one recipe to reference another. A recipe in the *PROCESS class*, for example, may refer to recipes in other *classes* to ensure that an entire set of recipes is *executed* together.

It is also convenient to allow one recipe to reference other recipes within the same *class*, where a substantial set of instructions and/or *settings* are common to a variety of other recipes.

A reference within the *body* of one recipe to a different recipe is called an **external reference**. The exact syntax in which this is done is determined by the recipe language used. An *external reference* may also specify values for *variable parameters* defined in the referenced recipe or in one of its *subrecipes*.

3.2.4.1.1 Subrecipes — When one recipe references another recipe, both recipes are required for *execution*. The starting recipe is called the **main** recipe.

Subrecipes are those recipes that are referenced by the *main* recipe, or by another *subrecipe* of the main recipe. A recipe that references other recipes is a **parent** recipe to the referenced *subrecipes*. A parent recipe is not necessarily a *main* recipe. A main recipe must not be referenced by any of its *subrecipes*.

External references to recipes within the same *namespace* may or may not be explicit. That is, references to the current *class* (i.e., the *class* of the recipe being *verified*) may be implicit, and *versions* may be

left to the rules for determining the default *version* at *link-time*.

A recipe may also reference recipes in another *namespace*. This occurs either when a recipe is to be **delegated** or when a recipe in a *namespace* accessed by multiple *agents* must reference a hardware-dependent or other *agent-specific* recipe kept in a **default namespace**. The *default namespace* of a *recipe executor* is always referenced within a recipe as the namespace named "Default". This name is interpreted by the *recipe executor* at execution time, based on its attribute DefaultNamespace (see Section 6.3).

When specified in a text string within a recipe with the subrecipe's *identifier*, the *identifier* is preceded by the value of the *namespace's* identifier, followed by a greater-than symbol ">". This is called the **recipe specifier**. For a recipe ETCH version 5 in a *namespace* named "NS-MOM", this would be specified as

"NS - MOM > /PROCESS/ETCH;5".

According to the rules of OSS, object type may be omitted when it may be otherwise determined. In RMS, determination is made according to order.

3.2.4.1.2 Delegated Recipes — A recipe may reference recipes in another *namespace* that are to be executed by a *component agent* and are termed **delegated recipes**. Although referenced by another recipe, *delegated recipes* are not considered as *subrecipes* of the *parent* (referencing) recipe as they are to be *executed* by a difference *agent* and are themselves required to be a *main* recipe. A reference to a *delegated recipe* is the equivalent to an automated "select" and "start" sequence of commands from the *supervisor*. As such, they are generally subject to the constraints of a *select* operation.

Delegated recipes may or may not require the *name* of the *executing agent* (*recipe executor*), depending on the rules of the particular implementation. For example, a *supervisor* may be able to determine the appropriate *executing agent* at run-time, based on the *namespace* specified. However, the ability to specify a particular *executing agent* is important to the user, as process results may be sensitive to a particular equipment installation. Therefore, recipe languages that support delegation shall allow specification of the



particular *executing agent* to be used. The text format of the full **recipe specifier** in this case is:

"Agent - Name > Namespace - Name > CLASS/Name;Version".

The required order of *identifiers* within the *specifier* is: agent, *namespace*, recipe.

3.2.4.1.3 Linked Recipe Sets — Where one recipe references another, a set of recipes to be executed together is formed by starting at the *main* recipe and collecting the *external references* to identify all the members of the set. This set is called a **linked recipe set**, and the operation of collecting the references is the **link** operation.

3.2.4.2 Variable Parameters — **Variable parameters** are variables that can be assigned values from outside the recipe itself. *Variable parameters* allow recipes to use variables rather than constants for such actions as setting temperature setpoints, time delay intervals, and data-set names. This capability greatly extends the reusability of a recipe.

A *parameter* first is formally **defined** within a recipe body and given a unique **parameter name**, a **parameter initial value** (default value) for use when the recipe is *selected* for execution, unless overridden. Where applicable, the *definition* also includes a **parameter restriction** that represents one or more conditions that any *value* assigned to that *parameter* is required to satisfy to be valid.

The syntax for *parameter definitions* in recipe *bodies* remains unspecified, so long as it conforms to the syntax of the recipe language and contains the required elements. Other restrictions concerning the use of *variable parameters* may be imposed by the *recipe executor's* supplier through the specification of the recipe language.

The **parameter domain** is the set of all possible values of a given *form* (see Section 1.4.3) that fulfill the conditions of the **parameter restriction** (if any).

The form need not be included in the *parameter definition* declared within the recipe *body*, but rather may be derived by the *recipe executor* during the *verification* process from the way in which the *parameter* is applied within the recipe. If temperature is maintained internally as an unsigned integer, for example, those *parameters* that are used in the recipe to assign temperature setpoints and ranges would typically be required by the supplier of the *executing agent* to rep-

resent unsigned integers as well. The supplier shall document the forms and valid *domains* for each of the *parameters* that may be used as a *variable parameter*.

NOTE: The *parameter's value* may be changed from its *initial* (default) *value* within the text of a recipe, and it may also be set as an "argument" passed to a *subrecipe*. These capabilities are a function of the recipe language and are beyond the scope of RMS.

Two categories of *parameters* are defined: **numeric** and **non-numeric**.

3.2.4.2.1 Numeric Parameters — *Numeric parameters* include all *parameters* that can take on any *numeric value* for its *format type* between a **parameter low limit** and **parameter high limit**.

The *parameter restriction* for a *numeric parameter* in any of the attributes that store *parameter definitions* shall be a text string that conforms to one of the following:

- "(a,b)⌋UNITS" the *domain* of numbers x such that $a < x < b$
- "(a,)⌋UNITS" the *domain* of numbers x such that $a < x < +\infty$
- "(,b)⌋UNITS" the *domain* of numbers x such that $-\infty < x < b$
- "UNITS" the *domain* of numbers x such that $-\infty < x < +\infty$.

where:

1. a and b are *numeric values* formed from the digits "0" through "9", the plus and minus signs "+" and "-", the period ".", and the letters "E" or "e" for floating point numbers,
2. ⌋ represents optional *whitespace*, and
3. UNITS is a valid case-sensitive string conforming to the Units of Measure Identifiers (see SEMI E5, Section 9).

For example, both of the following strings are valid: "(500,2000)⌋degC" and "(0,100)".

To include equality, parentheses are replaced by the square brace "[" and/or "]" on the left or right end, respectively. For example, "[a,b)" represents the domain of numbers x such that $a \leq x < b$.

To specify units only, with no restriction on range, the *low limit* may be set to $-\infty$ and the *high limit* to $+\infty$.



However, the *low* and *high* limit shall not exceed absolute minimum and maximum limits set for that *parameter* by the *recipe executor's* supplier.

Numeric parameters with no *parameter restriction* are pure numbers (with no units) with a domain of $(-\infty, \infty)$.

A *numeric parameter* shall have a *format type* of one of the following:

- text string of a length of no more than 80 characters,
- signed or unsigned integer of an even length,
- floating point number of an even length.

Text strings are restricted to the character set defined above to represent the *low* and *high limits* of the *domain* and shall convert to a *numeric value* within that *domain*.

3.2.4.2.2 Non-Numeric Parameters — A **non-numeric parameter** is any parameter other than a *numeric parameter*, including *parameters* whose *domains* are sets of discrete numbers that cannot be represented by a single mathematical interval, and strings that represent names. *Non-numeric values* may or may not place restrictions on the replacement *value*. A *parameter* that contains the name of a wafer map, for example, cannot be easily restricted by a general rule⁷, whereas a *parameter* that contains a string identifying thermocouple type can be restricted to one of a defined, unordered set of valid strings. Restrictions for *non-numeric parameters* might also consist of, or include, logical expressions that shall evaluate to TRUE before a value may be used.

If a recipe language has been defined for the corresponding *class* of recipes, the syntax for the *restriction* is taken directly from the formal *definition* of that *parameter* within the body of a *source form* recipe (and therefore conforms to the specification for that recipe language). Otherwise, the restriction is expressed in a syntax defined by the *recipe executor's* supplier specifically for *parameter* attributes.

For example, the *parameter* for a furnace recipe may specify which set of PID *values* should be used, where sets A-F exist on the *equipment*. The automatic restric-

tion on the *value* for such a parameter, as specified by the *recipe executor's* supplier, might be "{A,B,C,D,E,F}" with the default for the *value* specified as the character "C" with a *format type* of "ASCII string of length 1".

In another example, a *restriction* might be that the *value* be taken from the enumerated set

"{cassette, lot , batch }".

3.2.4.2.3 Agent-Specific Parameters — In order to "tune" a recipe so that it produces the same result on all *recipe executors* of the same type, it may be necessary to provide a different *initial value* for the *parameter* or a different *parameter restriction* for individual *executing agents*. A special editing facility may be provided by the *recipe manager* to allow the original *value* and/or *restriction* of a *parameter* to be modified for a specific *executing agent*.

NOTE: *Agent-specific parameter definitions* replace the corresponding original *definitions* in the recipe that is *downloaded* to the *recipe executor*. At run-time, they are superseded by *parameter values* specified for the *recipe executor's select* operation.

3.2.5 Attributes — The name of an attribute is a text string required to be unique for its object. The names of the attributes defined in this document are reserved for their standard use.

Identification attributes are those used to identify the object in OSS. These attributes are available through OSS, but in RMS, services are handled separately from the other attributes and are not transferred in parameters for "recipe attributes".

Attributes other than *identification* attributes and **mandatory** attributes discussed below always have a defined **default value**. Attributes with a *default value* not otherwise specified are considered to have a **null value** corresponding to their form as their *default value*. The *null value* for a text string is a zero-length string. The *null value* for a numeric form is zero. The *null value* for a boolean form is FALSE. The *null value* for an empty list or structure is an empty (zero-length) list. The manner of representing *null values* is left to the protocol.

An attribute is **cleared** or **reset** by setting its value to the *default value*.

Mandatory attributes are attributes that are required to always have a *non-default value*.

⁷ Typically, the validity of parameters which are unrestricted names can only be established at run-time.



Certain attributes are **required**, including all *mandatory attributes*. A **required** attribute is one that shall be supported with a *non-default value*. For example, the BodyFormat attribute is *required*. This means that *object form* recipes (for which BodyFormat has a *non-default value*) shall be supported by a *namespace*.

Any attribute defined in this standard shall only be modified according to the rules specified for that attribute.

Attributes of the *managed recipe* and the *execution recipe* may be accessed through Object Services. However, only certain of the attributes may be set through Object Services.

For purposes of OSS, all attributes defined in RMS shall be *recognized*. That is, a response to a GetAttr service request that references any attribute of an object defined in RMS for that object shall not return an "invalid attribute" error. If the attribute is not supported by the application, then it shall show the value of that attribute as having the null value appropriate for its form, and it shall deny attempts to set its value through the SetAttr service.

3.2.5.1 Descriptors — Two types of *mandatory* attributes of particular importance are the **timestamp** and **length** attributes. These are provided for the *body*, for each *agent-specific dataset*, and for the set of attributes of the recipe object itself, including the *timestamp* and *length* attributes of the *body*.

A **descriptor** contains the *length* and *timestamp* attributes (in that order) of one or more of the aspects of a recipe: its attributes, its *body*, or an *agent-specific dataset*. *Descriptors* are used to compare two recipes stored in different *namespaces* or in a *namespace* and in the *recipe execution area* of a *recipe executor*.

NOTE: Provision of a real-time clock or other device capable of counting time in centiseconds greatly improves the value of the *timestamps* for this comparison.

3.3 Full and Minimal Recipe Models — Figure 3.5 contains an object model for the *managed* and *execution* recipes supporting all of the standard *attributes* defined in RMS. *Non-identifier* attributes are shown in alphabetical order. Many of these attributes are not required for minimal applications. Minimal models for the *managed* and *execution* recipes are shown in Sections 3.4.2.3 and 3.5.2, respectively.




Figure 3.5
Full Recipe Object Model

3.4 *Managed Recipes* — This section defines the *managed recipe* and its attributes.

3.4.1 *Generic Attributes* — For clarity, the attributes of the *managed recipe* are sometimes called **generic** attributes to distinguish them from *agent-specific* attributes or attributes of the *execution recipe*.

Attributes that are not specifically defined in this standard may be defined by an application or by the user. Attributes that are defined by a *namespace* or by a *recipe executor* for its own use shall be changed only by rules specified and documented by the originator of the attribute.

To prevent conflict between non-standard attribute names and names which may become standard in the future, non-standard attribute names shall start with the prefix "UD_" (the two characters "U" and "D" followed by the underline character). It is the responsibility of the originator to provide unique attribute names. When using other RMS services, non-standard attributes shall follow all standard attributes.

It is desirable that a method be provided for the user to define new attributes.

3.4.1.1 Timestamp Attributes — **Timestamp** information (date and time of the last change) is important to recipe management. *Timestamp* attributes are maintained by the *recipe namespace* and may not be changed otherwise.

A **timestamp** attribute shall always contain the date and time that the particular aspect of the recipe was last changed. This attribute is a text string of the form "yyyymmddhhmmsscc" for the year yyyy, the month mm, the day dd, the hour hh, the minutes mm, the seconds ss, and the centiseconds cc.

The *timestamp* of the *body* is called EditTime and is set when a recipe is first **created** and updated whenever the *body* is modified in any way.

Because a recipe's attributes may be changed without changing the body, both the recipe's *generic* attributes and each set of *agent-specific* attributes themselves each have a *timestamp* attribute. For the *generic* attributes, the *timestamp* attribute is called AttrChgTime. The set of *agent-specific* attributes (of an *agent-specific dataset*) has its own *timestamp* attribute, called AgentSpec ChgTime.

3.4.1.2 Length Attributes — There are three **length attributes**, one which contains the length of body, one that contains the length of the *generic* attributes, and one that contains the length of the *agent-specific* attributes. *Length* attributes are calculated without regard to either protocol overhead or storage overhead, such as proprietary formats used for internal storage that may change from one implementation to another. This preserves the length across different conventions used for recipe storage and different communications protocol.

The *length attribute* of the *body*, BodyLength, contains the length of the *body* in bytes. The *length* of the *body* is calculated as the number of bytes it will require when transferred, excluding any overhead, such as that which might be required for protocol format information.

The length of an individual attribute is calculated as the sum of the lengths of the attribute name and the attribute value.

The length of a set of attributes shall be calculated as the sum of the lengths of the individual attributes that are set to a *non-default value* at the time the calculation is performed, including the *length attribute* itself. All

attributes set to their *default value*, at the time the *length* is calculated, are excluded from the calculation. This is because only the attributes with a *non-default* value are transferred when a recipe is moved into or out of a *namespace*. It also results in a more significant change to value of the *length* attribute when an attribute is set to a non-default value.

3.4.1.3 Descriptors — The *descriptors* of the *managed recipe* are the **body descriptor** (BodyLength and EditTime), the *generic attribute descriptor* (AttrLength and AttrChgTime), and the *agent-specific descriptor* (AgentSpec AttrLength and AgentSpec ChgTime). The **recipe descriptor** for a *managed recipe* consists of the *attribute descriptor*, the *body descriptor* (in that order), followed by *descriptors* of any existing *agent-specific datasets*.

3.4.2 Managed Recipe Object Attribute Definitions — This section provides the formal definitions for the recipe's *generic* attributes. For Object Services, a recipe is considered as the owner of its *components*, any *agent-specific datasets*. The *body* may not be accessed through Object Services.

Attributes in Tables 3.1 and 3.2 are presented in the following order:

- *identification* attributes, including object type and identifier, appearing above the heavy line in Tables 3.1 and 3.2,
- *mandatory* attributes and the other required attributes, in the order in which they are to appear when a recipe is transferred,
- optional attributes in alphabetical order, and
- non-standard attributes, which are transferred last.

When transferring a recipe with RMS services, *identification* attributes are not included in the list of attributes, and required attributes are sent, in order, before optional attributes. *Non-mandatory* attributes having their *default value* are not transferred, since their absence indicates their value, but they are always available through Object Services. Order of optional attributes is not dictated.

3.4.2.1 Generic Attribute Definitions — Table 3.1 provides the formal definition of the *generic* attributes of the *managed recipe*.



Table 3.1 Managed Recipe Attribute Definition

Attribute Name	Definition	Access	Rqmt	Form
Identification Attributes				
<u>ObjType</u>	The object type.	RO	Y	Text: "MRcp"
<u>ObjID</u>	An identifier derived from <u>Class</u> , <u>Name</u> , and <u>Version</u> . No part of a recipe's identifier shall be changed except through <i>renaming</i> .	RO	Y	Formatted text.
<u>Name</u>	A logical name assigned by the user when the recipe is <i>created</i> or <i>renamed</i> .	RO	Y	Text.
<u>Class</u>	The recipe's class (e.g., "/PROCESS/" or "/PROCESS/LOADER/").	RO	Y	Formatted text: "CLASS/CLASS/./CLASS/"
<u>Version</u>	The version of the recipe.	RO	Y	Text.
Mandatory Attributes				
<u>AttrLength</u>	The total length of the <i>generic</i> attributes, in bytes. <i>Mandatory</i> .	RO	Y	Unsigned integer.
<u>AttrChgTime</u>	Timestamp of the last change to a <i>generic</i> attribute. <i>Mandatory</i> .	RO	Y	Formatted text.
<u>BodyLength</u>	Length of the recipe's body, in bytes. <i>Mandatory</i> .	RO	Y	Unsigned integer.
<u>EditTime</u>	Timestamp of when the <i>body</i> was <i>created</i> or last <i>updated</i> . <i>Mandatory</i> .	RO	Y	Formatted text. <i>Timestamp</i> format.
Required Attributes				
<u>BodyFormat</u>	Indicates the form and format of the recipe's <i>body</i> . <i>Default</i> is zero.	RO	Y	Enumerated unsigned integer: 0 = <i>source</i> , 1 = <i>object</i> , > 1 reserved.
<u>Verified</u>	Indicates whether the recipe's body is syntactically correct. Reset when the recipe is <i>created</i> or <i>updated</i> . <i>Default</i> is FALSE.	RO	Y	Boolean.
<u>Linked</u>	Indicates whether the recipe is <i>linked</i> . Reset when the recipe is <i>originated</i> , <i>verified</i> , or <i>unlinked</i> . <i>Default</i> is FALSE.	RO	Y	Boolean.
Optional Attributes				
<u>ApprovalLevel</u>	Indicates the level of approval assigned by an <i>authorized user</i> . <i>Default</i> is zero. Reset when the recipe is <i>originated</i> or <i>linked</i> . For a <i>linked</i> recipe, may not be higher than any of its <i>subrecipes</i> .	RW	N	Unsigned integer.
<u>Comments</u>	User comments.	RW	N	Text. Maximum length is 80 characters.
<u>EditedBy</u>	The name of the person who last edited the recipe.	RO	N	Text. Maximum length is 40 characters.
<u>EstRunTime</u>	The nominal or estimated execution (run) time of the recipe, in seconds. Reset when the recipe is <i>created</i> or <i>updated</i> . Set when the recipe is <i>verified</i> . May be recalculated to total time for a main recipe when <i>linked</i> . Used for scheduling purposes. Algorithm for calculation shall be documented. <i>Default</i> is 0.	RW	N	Unsigned integer.

Table 3.1 Managed Recipe Attribute Definition (Continued)

Attribute Name	Definition	Access	Rqmt	Form
<u>ExecChgCtrl</u>	Specifies change control requirements for recipe. (See Section 6.5.)	RW	N	Binary. Bitwise (MSB = 8**). 1 - The recipe body may be changed, 2 - Change notification is required, 3 - Recipe may be selected after change, 4 - Most recent parameter settings shall be saved.
<u>ExtRef</u>	A list of all recipe <i>specifiers</i> as referenced within the recipe. Explicit <i>versions</i> not required. Reset when the recipe is <i>created</i> , <i>updated</i> , and <i>verified</i> .	RO	N	List of formatted text.
<u>LinkList</u>	A complete list of recipe <i>specifiers</i> found in the <u>ExtRef</u> attribute of a <i>main</i> recipe and all of its <i>sub-recipes</i> , with duplicates removed and all <i>versions</i> explicitly determined. Set for the <i>main</i> recipe when <i>linked</i> . Reset when the recipe is <i>originated</i> or <i>verified</i> . Required for multi-part recipe support.	RO	N	List of formatted text.
<u>LinkParam</u>	A list of all variable parameter definitions contained in the <u>Parameters</u> attribute of a <i>main</i> recipe and all of its <i>subrecipes</i> , with duplicates removed. Reset when the recipe is <i>created</i> , <i>updated</i> , or <i>verified</i> . Set when the recipe is <i>linked</i> . Required for <i>variable parameter</i> support.	RO	N	Structure composed of parameter name, initial value, and restrictions.
<u>Parameters</u>	A list of variable parameter definitions contained in the recipe. Reset when the recipe is <i>created</i> , <i>updated</i> , and <i>verified</i> . Set when the recipe is <i>verified</i> . Required only for <i>variable parameter</i> support.	RO	N	Structure composed of parameter name, initial value, and restrictions.
<u>SrcRepID</u>	<i>Identifier</i> of the <i>source form</i> recipe from which a <i>derived object form</i> recipe is derived. Value determined by the <i>verifier</i> of the recipe. Required only for support of <i>derived object form</i> recipes.	RO	N	Formatted text.
<u>VerificationID</u>	Identification code set by the <i>verifier</i> of the recipe. May be used to determine out-of-date formats that need to be <i>reverified</i> .	RO	N	Text. Maximum length is 40 characters.
<u>UD *</u>	Non-standard attribute defined by supplier or <i>user</i> . Asterisk indicates the part of the attribute name that is provided in this definition. Shall be preserved exactly except by the entity that defined it.	RO	N	Varies with definition. Text form is limited to 80 characters.

** NOTE: SEMI E4 and E5 number bits 1–8, where Bit 8 = MSB (most significant bit).

3.4.2.2 Agent-Specific Attribute Definitions — The names of user-defined attributes shall start with the prefix "AgentSpec_UD_" and shall be preserved without modification from transferred recipes. These attributes follow all standard attributes when transferred.

An *agent-specific dataset* exists only if an *agent-specific* attribute other than the *timestamp* and *attribute-length* attributes of the *dataset* itself have been set.



Table 3.2 defines the attributes of the *agent-specific dataset object*.

Table 3.2 Agent-Specific Dataset Object Attribute Definition

<i>Attribute Name</i>	<i>Definition</i>	<i>Access</i>	<i>Rqmt</i>	<i>Form</i>
<u>ObjType</u>	The object type.	RO	Y	Text = "MRcpASDS"
<u>ObjID</u>	The object's identifier. Contains the value in <u>AgentSpec Agent</u> .	RO	Y	Text.
<u>AgentSpec Agent</u>	The name of the <i>executing agent</i> to which the other attributes in the <i>dataset</i> apply. <i>Mandatory</i> .	RO	Y	Text.
<u>AgentSpec AttrLength</u>	The length of the <i>agent-specific</i> attributes, in bytes. <i>Mandatory</i> .	RO	Y	Unsigned integer.
<u>AgentSpec ChgTime</u>	Timestamp of when an <i>agent-specific</i> attribute was last changed. <i>Mandatory</i> .	RO	Y	Formatted text.
<u>AgentSpec Comments</u>	Comments specific to the <i>agent</i> entered by the author.	RW	N	Text. Maximum length is 80 characters.
<u>AgentSpec LinkParam</u>	A list of <i>variable parameter definitions</i> modified from the list in <u>LinkParam</u> . Valid only for a <i>linked main</i> recipe. <i>Parameter name</i> and form may not be changed.	RO	N	List of Structure composed parameter name, value, and restrictions.
<u>Certified</u>	The certification-level for the specific <i>agent</i> , assigned by an <i>authorized user</i> . Reset when <u>AgentSpec LinkParam</u> is modified. Required for <i>certification</i> support.	RW	N	Unsigned integer.
<u>AgentSpec UD *</u>	Non-standard attribute defined by the supplier or user. Asterisk indicates the part of the attribute name that is provided in this definition. Must be preserved exactly except by the defining entity.	RO	N	Varies with definition. Text form is limited to 80 characters.

3.4.2.4 Minimal Managed Recipe — The minimal model for a *managed recipe* is shown in Figure 3.6. Only required attributes are supported in this model. This model can only be used for single-part recipes in the dedicated *namespace*, such as the *default namespace* required for stand-alone *equipment*. There are no *agent-specific datasets*.




Figure 3.6
Object Model for Minimal *Managed Recipe*



3.5 Execution Recipes — The type of recipe handled by a *recipe executor* is called an *execution recipe*. This type of recipe has attributes and a *body* but does not have an *agent-specific dataset*. Most of the *generic* attributes of the *managed recipe*, and certain of the *agent-specific* attributes, are also attributes of the *execution recipe*, while in other cases, a *generic* attribute and an *agent-specific* attribute are merged and placed in a new *execution recipe* attribute.

A *managed recipe* is converted to an *execution recipe* when the recipe is *downloaded* to the *recipe executor*, and correspondingly, an *execution recipe* is converted to a *managed recipe* when *uploaded* from the *recipe executor*.

The *execution recipe* is defined in detail in Section 6.3.

4 Recipe Namespace

This section defines the basic conceptual model for the **recipe namespace**, a logical domain for recipe storage and management, within which the *identifier* of a recipe is guaranteed to be unique. The model in the current section is applicable to implementations with centralized storage. Section 5 defines an extension of the basic model for a distributed recipe namespace.

In general, the term "namespace" refers to a domain of unique *identifiers*. The issue of *namespace* boundaries exists for all object types, particularly in a distributed environment, and is not unique to recipes. However, within the context of RMS, **namespace** is used as a synonym for *recipe namespace*.

4.1 Motivations — A recipe's *identifier* may not be unique across different *namespaces* or throughout a factory. (That is, a given identifier may be used by internally different recipes except within a single *namespace*.) A primary role of a *namespace* is to define an area within which the uniqueness of any given recipe *identifier* may be guaranteed.

The requirements for Recipe Management that are addressed in this section include:

- to define the boundaries of specific areas where recipes may be uniquely identified, stored, and retrieved,
- to define the *attributes* and operations for the management domain of recipes, including recipe protection and recipe operations,
- to allow control of removable media,

- to allow stand-alone *equipment* to *execute* recipes,
- to facilitate smooth integration of stand-alone *equipment* into on-line factory systems,
- to allow integrated *equipment* to continue to execute recipes when communications have been lost, and
- to provide the basis for the **distributed** model in Section 5, which allows a *supervisor* to use and manage the storage capacity of its *subordinate agents*.

4.2 Namespace Model — The *namespace* model serves two major purposes. First, it provides a common set of management rules. Second, it allows a set of recipes to be shared among a group of *executing agents* that have a common process type, common functionality, and a common recipe language.

Figure 4.1 shows the basic *namespace* model with four objects: the recipe, the *namespace* itself, a *namespace* component called **recipe namespace segment**, and a **recipe namespace manager**.

The **recipe namespace manager** (or **manager**) represents the interface for the *namespace* to the external world and the internal decision authority within the *namespace*. All services for the *namespace* and its recipes are provided by the *manager*. The *namespace* itself is passive. While it has important *attributes*, the *namespace* has no operations, and it provides no services.

The **recipe namespace segment** (or **segment**) represents both the internal storage element and the actual manipulation of recipes within the *namespace*, under the supervision of the *manager*. In the basic model, it also provides no public services. The relationship between the *segment* and the *manager* in this model is presumed to be internal to an application.

Figure 4.1 represents a *namespace* with centralized storage. However, it can easily be extended to the **distributed recipe namespace**, which may have multiple *segments*, each provided by a different external *agent*. For that case, the *manager* and *segment* must provide public services for one another.




Figure 4.1
Recipe Namespace Model

All operations performed on a recipe through Recipe Management *services* are performed within the *namespace* by, or under the supervision of, the **namespace manager**, with two exceptions. A *namespace manager* is not required to understand the syntax of the recipe languages. For this reason, recipe **editing** (modifying a recipe's *body*) is not part of the *namespace* definition, and *verification* shall be provided by a **recipe executor** of an appropriate *executing agent* (see Section 6). These two activities require an in-depth understanding of the particular recipe languages used. Their separation from the other activities allows generic *namespace* capabilities to be provided.

4.3 Namespace Specifications — The combination of the *recipe namespace*, its *segment* component, and its *manager* provides storage, retrieval, and management of recipes conforming to RMS.

A *namespace manager* is responsible for maintaining the integrity of the *namespace*, the integrity of the recipes within the *namespace*, and the integrity of the recipe *identifiers*. It understands the rules regarding recipes, their *attributes*, and their components, and it is responsible for enforcing those rules. It will not allow a *read-only* recipe to be changed or deleted, for example. Therefore, it will not accept a recipe with an *identifier* already used by a *read-only* recipe.

A *namespace* has no restrictions on the **read access** of a recipe as a whole or of its *attributes*, nor is it concerned with the uses to which they might be put outside the *namespace*. The ability for multiple *agents* to share the same recipes is determined solely by the ability of the *agents* to access the same *namespace*. Issues of security and authentication are beyond the scope of RMS.

The term *recipe namespace*, or *namespace*, is used inclusively to refer to those *attributes*, operations, and other requirements common to both the *namespace* described in this section and to its *distributed* subtype. The term **centralized namespace** is used in references to an instance (implementation) of the basic model and to clarify statements that do not apply to instances of the *distributed* subtype.

A *centralized namespace* is analogous to a single directory of files, where duplicate file names are not allowed. It may exist in several different configurations that are incidental to *namespace* requirements.

It is possible to provide a *namespace* that uses removable media for its physical storage. Recipes may then be transferred to and from this *namespace* to any other *namespace*.

In no case shall it be possible to transfer recipes into a *namespace* except according to the requirements for *namespace management*. For example, a *read-only* recipe may not be replaced.

The *centralized namespace* may be applied several ways. An *agent* with *execution* capabilities that may be operated in stand-alone mode shall provide itself with a *centralized namespace* to be used when operating in stand-alone mode. Other *agents*, such as diskless process modules in a cluster, may expect to rely upon the *namespace* capabilities provided by the cluster *supervisor*.

4.4 *Member Agents* — A *namespace* that contains recipes used by multiple *agents* is called **shared**. Otherwise, it is called **dedicated** or **non-shared**. Figure 4.2 illustrates two *dedicated namespaces*, one provided by an *etcher* and one by its host. This terminology is introduced for clarification and descriptive purposes only.




Figure 4.2

Host Backup of Equipment Namespace

The *Members* attribute of the *namespace* contains the names (object identifiers) of the *agents* that use the recipes in the *namespace*, and these *agents* are called **member agents (members)** of the *namespace*. The *namespace manager* uses this information as it may require assistance from a *member agent* to *verify* a recipe.

If the *Members* attribute of the *namespace* contains multiple *agent* names, then the *namespace* is *shared*. Otherwise, it may be *non-shared* or not yet completely set up. If it is empty (null), the *namespace* has not been completely set up and is not fully functional, as certain operations, including recipe *verification*, require it to have content.

4.5 *Illustrations* — This section provides illustrations of several possible configurations of a *centralized namespace*.

Figure 4.2 illustrates an *etcher* that has its own local *namespace* and that communicates with its *supervisor* over an RS-232 line. The *supervisor* in this example maintains a separate *namespace*, providing backup for recipes of particular significance.

Figure 4.3 illustrates a *supervisor* with four diskless subordinates. In this configuration, the *supervisor* provides recipes for all four subordinates from a *centralized namespace*.




Figure 4.3

Namespace Provided by Supervisor

Figure 4.4 illustrates four *executing agents* on a common network with a *supervisory agent* and a sixth *agent* providing a *centralized namespace*. Each *executing agent* is able to access the *namespace* independently of the *supervisor*.




Figure 4.4

Shared Network Access

Figures 4.3 and 4.4 both represent examples of *shared centralized namespaces*.

4.6 *Attribute Definition Tables* — Table 4.1 defines the *attributes* of the *recipe namespace*.

Table 4.1 Recipe Namespace Attribute Definition

<i>Attribute Name</i>	<i>Definition</i>	<i>Access</i>	<i>Rqmt</i>	<i>Form</i>
<u>ObjType</u>	The object type.	RO	Y	Text = "RNS"
<u>ObjID</u>	The <i>name</i> of the <i>namespace</i> .	RO	Y	Text. A name of "Default" is prohibited.
<u>RecipeReadOnly-Level</u>	The level of <i>approval</i> at which recipes are <i>read-only</i> .	RW	Y	Unsigned integer.
<u>Members</u>	The <i>names</i> of <i>agents</i> capable of <i>verifying</i> and <i>executing</i> the recipes in the <i>namespace</i> .	RW	Y	List of <i>agent names</i> (identifiers).
<u>SynchOn</u>	Level of synchronization (see Section 9.5). Required if synchronization is supported.	RW	N	Unsigned integer: Either 0 = Disabled, or any combination (sum) of: 1 = changes in body 2 = new <i>execution recipe</i> 8 = changes in <i>last value</i> 16 = new <i>derived object form execution recipe</i>
<u>SynchFail</u>	Specifiers for <i>execution recipes</i> for which an attempt to <i>synchronize</i> failed. Required if <i>synchronization</i> is supported.	RW	N	List of formatted text.

Table 4.2 defines the attributes of the *recipe namespace manager* objects. The *name (identifier)* of the *manager* is not generally of interest, as the *namespace specifier* is more commonly used. However, it is important for the *manager* to be accessible through Object Services.

Table 4.2 Recipe Namespace Manager Attribute Definition

<i>Attribute Name</i>	<i>Definition</i>	<i>Access</i>	<i>Rqmt</i>	<i>Form</i>
<u>ObjType</u>	The object type.	RO	Y	Text = "RNS_Mgr"
<u>ObjID</u>	The <i>manager's</i> name.	RO	Y	Text.
<u>NamespaceName</u>	The <i>name</i> of the <i>namespace</i> managed.	RO	Y	Text.

5 Distributed Recipe Namespace

A **distributed recipe namespace** is a *recipe namespace* that utilizes the storage capacity of multiple *agents* for recipe storage. Recipes are stored in special **recipe namespace segments** provided by the different *agents*. These *segments* are **distributed recipe namespace segments** and are supervised by the **distributed recipe namespace manager**.

This section defines the different objects that together provide the *distributed recipe namespace* capability: the *distributed recipe namespace*, the *distributed recipe namespace manager*, the *distributed recipe namespace segment*, and the **distributed recipe namespace recorder**. Detailed descriptions of operations are contained in Section 10.

Throughout this document, the acronym **DRNS** refers to the term "distributed recipe namespace" and is used primarily to differentiate a *DRNS* subtype object from its supertype.

5.1 Motivations — The *distributed recipe namespace* capability provides a method for using the storage capacity of multiple *agents*, as illustrated in Figure 5.1. This reduces the storage requirements of a centralized factory system, improves performance by allowing recipes to be used by those *agents*, and provides a centralized management of the *namespace* to ensure that the uniqueness of the recipe *identifier* is properly maintained.

Figure 5.1 illustrates a *distributed recipe namespace* with four *segments*, each provided by a different *agent*, and each *agent* also having a *local namespace*.




Figure 5.1
Illustration of Agents, Segments, and Local Namespaces



5.3 Distributed Recipe Namespace Issues — This section introduces issues that are applied to more than one object within the *distributed recipe namespace* capability.

5.3.1 Object Services — Attachment is a relationship between a managed object and its manager that is defined in OSS (SEMI E39). One object may be attached, detached, and reattached to a manager. When attached, the managed object is able to recognize that certain critical services have been requested by its manager. The *segment* and *recorder* objects shall comply with requirements for the operations and services defined in SEMI E39 (OSS) to attach, detach, and reattach to and from a *DRNS manager*. They shall also allow their *manager* to modify specified attributes that are otherwise read-only.

The ability to create and delete a *segment* or *recorder* is optional.

The *authorized user* may request a *DRNS manager* to attach or detach a specified *segment* or *recorder*.

5.3.2 Logical Recipe — A **logical recipe** is defined as a recipe with a specific *body* and a specific set of *generic attribute values*. Every *managed* recipe stored is an instance of a *logical recipe*.

In a *distributed recipe namespace*, it is normal for instances (copies) of a given *logical recipe* to exist in more than one *DRNS segment* at any time. In other words, multiple duplicate copies of a recipe with the same *identifier* may co-exist within the *distributed recipe namespace*. To retain the integrity of the *namespace*, this can be allowed if, and only if, each copy of a recipe with a given *identifier* is an instance of the same *logical recipe*. It is the responsibility of the *DRNS manager* to ensure this logical identity is maintained.

Agent-specific datasets are not included in the definition of the *logical recipe* because multiple *agent-specific datasets* may exist independently of one another, and only the *dataset* specific to the *agent* providing the *DRNS segment* is normally kept in that *segment*.

5.3.3 Change Requests — A **change request** occurs whenever a *user*, an external application, or an attached *DRNS segment* requests the *DRNS manager* to make, or permit, any change in a recipe. Information concerning *change requests* is kept in the form of logical **change request records**. Once a *change request* is

made, a *change request record* is created and maintained until the change has been either completed or discarded. Because the information represented by *change request records* is publicly available, they provide a degree of diagnostic capability.

The *change request record* is not a formal or standardized object, but it represents the information that is available through services.

Change management is the most critical issue of the *DRNS* capability and is discussed in detail in different sections below.

5.4 Distributed Recipe Namespace Segment — The **distributed recipe namespace segment** (Figure 5.3) is responsible for all of the activities that directly manipulate recipes, including storage, retrieval, deletion, and operations that change a recipe's *attributes* or *body*. This requires a micro-level knowledge of the recipe's *identifier*, its structure, the inter-relationships between the various recipe *attributes*, how the recipe is stored, and how it is transferred. It also includes a macro-level knowledge of all of the *identifiers* of the recipes that it has stored. The *DRNS segment* shall ensure that only one recipe with a given *identifier* exists within that *DRNS segment*. The contents of an unattached *DRNS segment* shall be *read-only*.

A *DRNS segment* and its storage are provided by an *agent*, which could be equipment, an independent "recipe server", or other factory systems.




Figure 5.3
Distributed Recipe Namespace Segment

5.4.1 *Master and Dedicated Segments* — A **master segment** is a specialization of a *DRNS segment* that is capable of storing multiple *agent-specific datasets* per recipe. *Master segments* are not dedicated to a single equipment and are used to store a **full copy** (including all existing *agent-specific datasets*) of all *logical recipes* within the *namespace*. It is required that every *distributed recipe namespace manager* be capable of supporting at least one attached *master segment*. Additional *master segments* may be desirable for further backup protection.

The term **dedicated segment** is used to refer to *segments* that are not *master segments*. *Dedicated segments* store at most one *agent-specific dataset* per recipe. Because a *DRNS segment* provided by equipment would not expect to keep *agent-specific datasets* for other *equipment*, equipment normally does not provide *master segment* capability.

Master and *dedicated* segments have the same attributes and support the same message services, but they respond differently to *agent-specific datasets* and are used for different purposes. A *master segment* requires approval prior to changing an *agent-specific dataset*.

5.4.2 *Change Restrictions* — The *DRNS segment* may provide access to recipes within its storage to the other

components of the *agent* providing the storage. However, neither the generic attributes nor the body of recipes stored by the *segment* shall be changed except with the explicit approval of the *distributed recipe namespace manager*.

Dedicated DRNS segments that are attached to a *DRNS manager* may change the contents of an *agent-specific dataset* without first asking permission. However, they are required to notify the *manager* of any change as soon as it occurs.

The attribute RecipeReadOnlyLevel is set by the *DRNS manager* to the value of the corresponding *namespace* RecipeReadOnlyLevel attribute when the *segment* is first attached and whenever the *namespace* attribute is changed. This attribute has the same function as the *namespace* attribute and allows the *segment* to prohibit changes based on a recipe's *approval level* as defined in Section 8.2.5. For example, the *segment* shall deny requests to *modify* a *write-protected* recipe.

The *DRNS segment* is prohibited from changing its contents whenever it is unattached. However, recipes and recipe *attributes* may be read at any time by other entities, including other *components* of the owner *agent*.

Any changes to a logical recipe stored by the *DRNS segment* shall first be approved by the *DRNS manager* before the change is made. This includes any changes to the body or to any generic attribute. This is required for two reasons. First, two different *DRNS segments* may attempt to change a recipe at the same time, and this activity must be coordinated. Second, to protect the integrity of the recipe *identifier* where multiple instances of a recipe exist, the *distributed recipe namespace* is required to ensure that all such instances have been updated appropriately with that change before other changes to the same recipe are allowed.

NOTE: The prohibition against unauthorized change does not preclude the saving of such changes external to the *DRNS segment* while waiting for authorization. However, changes made by other *DRNS segments*, subsequent to such changes and prior to authorization, may invalidate these changes.

Communications between the different *components* of an *agent* that do not require or use the formal *services* defined by RMS are considered as proprietary to the *agent* and are neither covered nor excluded by RMS, subject to the above restrictions against change.




Figure 5.4

Distributed Recipe Namespace Recorder

5.5 Distributed Recipe Namespace Recorder — The **distributed recipe namespace recorder** provides a method of externally storing and retrieving information critical to automated rebuilding of a damaged *distributed recipe namespace*. It contains two types of information: a list of the *DRNS segments* that are attached to the *DRNS manager* and the current *change request record* in process per recipe.

A *DRNS recorder* may be attached and detached from a *DRNS manager*, and reattached to that *manager*. Information that it contains is available to anyone but may only be changed by its *manager*.

The *DRNS recorder* is able to store the *DRNS segment specifiers* (the *object specifiers*) of the attached *DRNS segments* in its Segments attribute. This information allows a *distributed recipe namespace* to be automatically rebuilt in the event that the *namespace*, or the *DRNS manager*, is damaged and its information becomes lost or unavailable.

The *DRNS recorder* is also used to store the current *change request* per recipe, to delete a *change request*, and to return the set of *change requests* for one or more recipes. The attribute LockedRecipes provides a list of *identifiers* for recipes with *change requests*.

The *DRNS manager* shall use the services of a *DRNS recorder* when one is attached to the *namespace* by a service user. However, use (attachment) of a *recorder* is optional for the user. The *DRNS recorder* is provided

for remote storage of critical information and is not intended as a general source of information for the user. It is not able to provide *inactive change request* information.




Figure 5.5

Distributed Recipe Namespace

5.6 Distributed Recipe Namespace Management Information — The *DRNS management information* object is private and proprietary to the *DRNS manager*. It is included in Figure 5.2 as an emphasis on the importance of the information that a *DRNS manager* requires for management. This includes, but is not limited to, the recipe *identifiers* stored in each attached *segment* and all existing *change request records*.

The *DRNS manager* is required to know the contents of all of its attached *segments* at all times. It shall be able to uniquely identify each instance of a recipe within the *distributed recipe namespace*. It is responsible for tracking the current status of each instance of a logical recipe within the *distributed recipe namespace* when a change to a recipe within one *segment* is being updated to other *segments*. The *DRNS management information* is important for these purposes.

5.7 Distributed Recipe Namespace — The *distributed recipe namespace* has three additional attributes, as shown in Figure 5.5, that are read-only, set by the *DRNS manager*. The attribute Segments contains a list of the *object specifiers* of the *DRNS segments* currently attached. The attribute Recorder contains the object specifier of an assigned *DRNS recorder*. The attribute LockedRecipes contains a list of recipe *identifiers* of recipes with existing *change request records*.

5.8 Distributed Recipe Namespace Manager — The *distributed recipe namespace manager* (Figure 5.6) is responsible for ensuring that the *distributed recipe namespace* and *DRNS segments* operate together and for maintaining *namespace* integrity as a whole. It is required to know the identities of its attached *DRNS recorder* and *DRNS segments* and the contents (*recipe identifiers*) of each attached *DRNS segment* at all times.




Figure 5.6
Distributed Recipe Namespace Manager

The current attachments of the *DRNS manager* are maintained in the *DRNS namespace* attributes Recorder and Segments and through the services of the *DRNS recorder* described in Section 5.5. Whenever a *DRNS segment* is attached, the *DRNS manager* shall add the *segment specifier* to the *namespace* attribute Segments and to the current list maintained by the *DRNS recorder*. Whenever a *DRNS segment* is detached, the *DRNS manager* shall remove its *specifier* from the Segments attribute and from the *DRNS recorder's* list.

The *DRNS manager* is able to know which recipes are stored in each *DRNS segment* through use of Object Services provided for or by the individual *DRNS segments*. It shall maintain the integrity of the *namespace* through ensuring that an *identifier* of any recipe stored within the *namespace* represents exactly one *logical recipe*. This is achievable because the *manager* must give explicit approval of any change of *logical recipes* (generic attributes or body) stored by the *DRNS segments*.




Figure 5.7
Change Request Record

5.8.1 Change Management — A *change request* occurs whenever a user or an attached *segment* requests the *distributed recipe namespace manager* to make, or permit, a specific type of change. A change request record (Figure 5.7) contains information about who requested what type of change. This allows later diagnostics when necessary.

A *change request record* is created for a recipe whenever an attached *DRNS segment* requests a change and is deleted when the requested change has been completed or discarded. A *change request* is either **inactive** or **active**. Once a *change request* is made, it is considered as inactive until the change has been approved by the *DRNS manager*. Once approved, it is considered as **active**.

The *change request record* is not a standardized object and is not accessed directly through public services. However, the information represented by the *change request record* shall be available, upon request, from the *DRNS manager* for all *inactive* and *active change requests*. For this reason, it is convenient to model the *change request record* as an object.

A recipe for which a *change request* exists is called **locked**. Otherwise, it is **unlocked**.

The *DRNS manager* is responsible for ensuring that only one *change request* per recipe is *active* at any time. The results of each approved *change request* shall be updated appropriately to each of the other

DRNS segments with a copy of the same *logical recipe* before a subsequent change to the recipe is approved.

5.9 Building a Distributed Recipe Namespace — A *distributed recipe namespace* is built up in several stages. The *distributed recipe namespace* and its associated *DRNS manager* are created separately in the first stage. The individual *DRNS segments* are first *created* and then attached to the *DRNS manager*. At least one *DRNS segment* must be attached to the *DRNS manager* before it can accept recipes, as the *distributed recipe namespace* has no other means of storage. Figure 5.8 illustrates the attachment relationship.




Figure 5.8
The Attachment Relationship

The attached *DRNS segments* may later be detached from the *namespace* and then either attached again to the same *namespace* or to a different *namespace* without affecting the recipes stored within the *DRNS segment*.

NOTE: When a *DRNS segment* with existing recipes is detached and then attached to a different *namespace*, its recipe *identifiers* fall within the domain of the new *namespace*. This may require some recipes to be *renamed* before the attach process is complete.

5.10 Rebuilding a Damaged Distributed Recipe Namespace — A *distributed recipe namespace*, or its *manager*, may become damaged or unavailable. If a *DRNS recorder* was attached to the *damaged namespace*, then rebuilding the *namespace* can be automated.

A new *distributed recipe namespace* and new *DRNS manager* are created, assigning the old *namespace* name as the object identifier *ObjID* for the new *namespace*. The new *DRNS manager* should be assigned a different *identifier*, however, as a security measure. The *user* may then request the new *DRNS manager* to rebuild the *namespace* with the old *DRNS recorder*. A *namespace* may also be rebuilt without a *recorder* if the list of *segment specifiers* can be provided by the user. However, the user is not expected to provide the information retained through *change request records*.

5.11 Object Attribute Definition Tables — This section contains the formal attribute definitions for the *distributed recipe namespace* capability. Except for the *object identifier* attributes, attributes are listed alphabetically.

Objects that are subtypes of objects introduced in Section 4 inherit the attributes of the supertype objects. These attributes are not repeated in this section.

5.11.1 Distributed Recipe Namespace Segment Attribute Definition — Table 5.1 defines the attributes required for a *distributed recipe namespace segment*, and Table 5.2 defines the attributes of the subtype *master segment*.



Table 5.1 Distributed Recipe Namespace Segment Attribute Definition

<i>Attribute Name</i>	<i>Description</i>	<i>Access</i>	<i>Rqmt</i>	<i>Form</i>
<u>ObjType</u>	The object type.	RO	Y	Text = "RNSDSegment"
<u>ObjID</u>	The object name (identifier).	RO	Y	Text.
<u>Namespace</u>	The name (<u>ObjID</u>) of the <i>namespace</i> to which the <i>segment</i> belongs. May be set by the manager.	RO	Y	Text.
<u>NamespaceManager</u>	Identifies the <i>distributed recipe namespace manager</i> . May be set by the manager.	RO	Y	Text.
<u>RecipeReadOnlyLevel</u>	Used to track the corresponding attribute of the <i>namespace</i> to which the <i>segment</i> belongs. May be set by the manager.	RO	Y	Unsigned integer.

Table 5.2 Distributed Recipe Namespace Master Segment Attribute Definition

<i>Attribute Name</i>	<i>Description</i>	<i>Access</i>	<i>Rqmt</i>	<i>Form</i>
<u>ObjType</u>	The object type.	RO	Y	Text = "RNSDMaster"
<u>ObjID</u>	The object name (identifier).	RO	Y	Text.

5.11.2 *Distributed Recipe Namespace Recorder Attribute Definition* — Table 5.3 defines the attributes required for a *distributed recipe namespace recorder*.

Table 5.3 Distributed Recipe Namespace Recorder Attribute Definition

<i>Attribute Name</i>	<i>Description</i>	<i>Access</i>	<i>Rqmt</i>	<i>Form</i>
<u>ObjType</u>	The object type.	RO	Y	Text = "RNSDRecorder"
<u>ObjID</u>	Text.	RO	Y	Text.
<u>LockedRecipes</u>	List of <i>identifiers</i> of recipes with existing <i>change request records</i> .	RO	Y	List of formatted text.
<u>Namespace</u>	Identifies the <i>namespace</i> to which the recorder is attached. May be set by the manager.	RO	Y	Text.
<u>NamespaceManager</u>	Identifies the <i>distributed recipe namespace manager</i> . May be set by the manager.	RO	Y	Text.
<u>Segments</u>	List of <i>specifiers</i> of currently attached <i>segments</i> .	RO	Y	List of formatted text.



5.11.3 *Distributed Recipe Namespace Attribute Definition* — Table 5.4 defines the attributes required for a *distributed recipe namespace*.

Table 5.4 Distributed Recipe Namespace Attribute Definition

Attribute Name	Description	Access	Rqmt	Form
<u>ObjType</u>	The object type.	RO	Y	Text = "RNSD"
<u>ObjID</u>	Text.	RO	Y	Text.
<u>LockedRecipes</u>	A list of <i>identifiers</i> of all recipes with existing <i>change request records</i> .	RO	Y	List of formatted text.
<u>Recorder</u>	The <i>recorder specifier</i> of the attached <i>distributed recipe namespace recorder</i> .	RO	Y	Text.
<u>Segments</u>	A list of <i>specifiers</i> of the <i>distributed namespace segments</i> attached to the <i>namespace</i> .	RO	Y	List of formatted text.

5.11.4 *Distributed Recipe Namespace Manager Attribute Definition* — Table 5.5 defines the attributes required for a *distributed recipe namespace manager*.

Table 5.5 Distributed Recipe Namespace Manager Attribute Definition

Attribute Name	Definition	Access	Rqmt	Form
<u>ObjType</u>	The object type.	RO	Y	Text = "RNS_MgrD"
<u>ObjID</u>	The <i>manager's name</i> .	RO	Y	Text.

6 Recipe Executor

This section describes the basic concepts for the *recipe executor* and the *execution recipe* that it stores.

The **recipe executor** is the component of an *executing agent* that reads and comprehends the contents of a recipe (its *body*) and puts into effect its instructions, settings, and/or other data. The object model of the *recipe executor* is shown in Figure 6.1.




Figure 6.1
Recipe Executor

The *recipe executor* is able to temporarily store recipes for execution purposes, and it may also be able to store *execution recipes* for later execution.

The execution process is beyond the scope of RMS.

6.1 *Motivations* — Specification of the *recipe executor* and the *execution recipe* that it stores is necessary to complete the management of recipes in the factory. The *recipe executor* provides limited storage for recipes and minimum capability to manage them. The storage provided is intended to be temporary only.

Some *recipe executors* have the ability to purposefully change a recipe's body or create new recipes. To ensure that *execution recipes* remain synchronized with the *managed recipes* in a *recipe namespace*, additional rules are required for such cases.

The requirements for RMS that are addressed by the *recipe executor* include the following:

- The ability to create, edit, and change recipes outside the *executing agents* that execute them,
- The ability for *executing agents* to use recipes developed externally,
- The ability to share recipes among equipment of the same type,
- Protection of stored *execution recipes* and of currently *selected* recipes from unexpected or unauthorized change,
- Execution of recipes without errors, including errors detected during *verification* and **validation** errors caused by improper settings or parameters or incompatibility with the current configuration,
- The ability to change a recipe's *parameters* between runs in a systematic way without changing the recipe itself,
- The ability to dynamically connect and disconnect the *executing agent* from the communications environment, and
- The ability of standalone equipment to execute recipes without a communications link.

6.2 *Description* — The *recipe executor* is able to receive a downloaded *execution recipe* and temporarily store it. The *executor* can *verify* the recipe *body* (Section 11.2.2) both at the time of the download and after the recipe has been stored. It is able to store at least as many recipes as it requires for a single process cycle, which is determined by its own requirements.

Recipes stored by the *recipe executor* may be downloaded from more than one *namespace*. To prevent ambiguity or conflict between recipe *identifiers* from

different *namespaces*, the name of the namespace from which the recipe was originally downloaded is retained in the *identifier* of the *execution recipe*.

The *recipe executor* **selects** one or more specified recipes by *validating* them and preparing them for execution. This may include moving the recipes into a separate **recipe execution area** to create an **executable copy** recipe. The *executable copy* recipe, if it exists separately, shall be protected from inadvertent change caused by other activities, such as downloading a new recipe. The *executable copy* recipe is not otherwise addressed by RMS. If a separate *copy* is *selected*, the stored *execution recipe* shall be protected from change. Protection from change and permission to change are discussed in Section 6.6.

Validation of a recipe consists of checking the values of its settings and variable parameters against existing supplier-defined and/or user-defined restrictions, and ensuring the recipe, or the *linked recipe set*, is valid for the current configuration of the *executor* (e.g. equipment or attached module).

Additionally, the *recipe executor* is able to calculate the amount of available storage, to delete recipes from its storage to make room for new recipes, to **de-select** recipes by preventing them from being re-executed without another explicit *select*, to **rename** an *execution recipe*, and to provide requested information about itself and its stored *execution recipes*, in conformance with OSS.⁸

The *recipe executor* may wish to rewrite *source form* recipes into a proprietary *derived object form* that is more efficient for execution or storage purposes. Where this type of recipe is to be stored for re-use, a new *identifier* is required for the *object form*, as described in Section 3.2.2.1.2. *Derived object form* recipes are discussed in detail in Section 11.2.2.1.

6.3 *The Execution Recipe* — An *execution recipe* is a type of recipe, as shown in Section 3, Figure 3.5. The *recipe executor* stores recipes as *execution recipes*. *Execution recipes* are created in one of two ways: they are either **downloaded** from a *recipe namespace*, or they are **created** by the *recipe executor*.

⁸ SEMI E39 (Object Services Standard)

6.3.1 Comparison of Managed and Execution Recipes

— An *execution recipe* differs from a *managed recipe* in two ways: its attributes and its lack of *agent-specific datasets*.

The differences in attributes between a *managed recipe* and an *execution recipe* consist of:

- The addition of Namespace as an *identification* attribute of the *execution recipe*,
- The intermediate parameters, ExtRef and Parameters, of a *verified managed recipe*, required for the *namespace link* operation, are not used by the *recipe executor* and are not retained in the *execution recipe*,
- The attributes LinkParam and AgentSpec LinkParam of the *managed recipe* are merged into the single attribute ExecLinkParam,
- The attributes of the remaining *agent-specific dataset* are absorbed into the attributes of the *execution recipe*, and
- The addition of the attribute *length* ExecAttrLength and attribute *timestamp* ExecAttrChgTime attributes, and
- The additional attribute ChangedBody, which is required for *execution recipes* where the *recipe executor* is capable of changing the *recipe body* or of creating new recipes and is not otherwise used.

The conversion of a *managed recipe* to and from an *execution recipe* is the responsibility of the *namespace manager* and is discussed in Sections 9.4.8 and 9.4.9.

Most of these differences are invisible for an *unverified recipe* or minimal recipe implementations. For an *unverified recipe*, all of the above attributes are *cleared* to their *default value* and are not transferred with the recipe. Attributes needed for multi-part recipes and *variable parameters* are not required for minimal implementations. Figure 6.2 provides a model of the minimal implementation of an *execution recipe* to meet RMS requirements.




Figure 6.2

Object Model for Minimal Execution Recipe

6.3.2 *Downloaded Recipes* — To *download* a recipe from a *namespace*, the *namespace manager* is responsible for first converting the attributes of the specified *managed recipe* into those of the *execution recipe*. The *recipe executor* stores the downloaded recipe as an *execution recipe*. The body of the recipe is unchanged by the download operation and may be in *source form* or *object form*, including the *derived object form*. (See Section 11.2.2.1.)

The *namespace manager* may optionally request that the recipe being downloaded replace any pre-existing recipe with the same *identifier*. Otherwise, if such a recipe exists, the download request is denied.

Upon request, the *recipe executor* shall *verify* a recipe that it has previously stored and return to the requestor the information required for recipe management.

Attributes of the *execution recipe* that have the same attribute *name* as either a *generic attribute* or an *agent-specific attribute* of the *managed recipe* are not changed except in specific cases.

The attributes giving the *timestamp* and *length* of the attributes of the *execution recipe* are different from those of the *managed recipe*. This allows the *generic* and *body descriptors* of a *managed recipe* to be compared with the corresponding attributes of an *execution recipe* for traceability. The *timestamp* of the *execution recipe's* attributes is ExecChgTime, and the *length* of its attributes is ExecAttrLength. These are calculated by the *recipe executor* at the time the *execution recipe* is stored and updated whenever other attributes of the recipe change.