

Network Working Group
Request for Comments: 3216
Category: Informational

C. Elliott
Cisco Systems
D. Harrington
Enterasys Networks
J. Jason
Intel Corporation
J. Schoenwaelder
F. Strauss
TU Braunschweig
W. Weiss
Ellacoya Networks
December 2001

SMIng Objectives

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This document describes the objectives for a new data definition language, suitable for the modeling of network management constructs, that can be directly mapped into SNMP and COPS-PR protocol operations.

The purpose of this document is to serve as a set of objectives that a subsequent language specification should try to address. It captures the results of the working group discussions towards consensus on the SMIng objectives.

Table of Contents

1.	Introduction	3
2.	Motivation	3
3.	Background	4
4.	Specific Objectives for SMIng	4
4.1	Accepted Objectives	5
4.1.1	The Set of Specification Documents	6
4.1.2	Textual Representation	6
4.1.3	Human Readability	6

4.1.4	Rigorously Defined Syntax	6
4.1.5	Accessibility	7
4.1.6	Language Extensibility	7
4.1.7	Special Characters in Text	7
4.1.8	Naming	8
4.1.9	Namespace Control	8
4.1.10	Modules	8
4.1.11	Module Conformance	9
4.1.12	Arbitrary Unambiguous Identities	9
4.1.13	Protocol Independence	9
4.1.14	Protocol Mapping	10
4.1.15	Translation to Other Data Definition Languages	10
4.1.16	Base Data Types	10
4.1.17	Enumerations	11
4.1.18	Discriminated Unions	11
4.1.19	Instance Pointers	11
4.1.20	Row Pointers	12
4.1.21	Constraints on Pointers	12
4.1.22	Base Type Set	12
4.1.23	Extended Data Types	12
4.1.24	Units, Formats, and Default Values of Defined Types and Attributes	13
4.1.25	Table Existence Relationships	13
4.1.26	Table Existence Relationships (2)	14
4.1.27	Attribute Groups	14
4.1.28	Containment	14
4.1.29	Single Inheritance	15
4.1.30	Reusable vs. Final Attribute Groups	15
4.1.31	Events	16
4.1.32	Creation/Deletion	16
4.1.33	Range and Size Constraints	16
4.1.34	Uniqueness	16
4.1.35	Extension Rules	17
4.1.36	Deprecate Use of IMPLIED Keyword	17
4.1.37	No Redundancy	17
4.1.38	Compliance and Conformance	17
4.1.39	Allow Refinement of All Definitions in Conformance Statements	18
4.1.40	Categories	18
4.1.41	Core Language Keywords vs. Defined Identifiers	19
4.1.42	Instance Naming	19
4.1.43	Length of Identifiers	19
4.1.44	Assign OIDs in the Protocol Mappings	20
4.2	Nice-to-Have Objectives	20
4.2.1	Methods	20
4.2.2	Unions	21
4.2.3	Float Data Types	21
4.2.4	Comments	22

4.2.5	Referencing Tagged Rows	22
4.2.6	Arrays	22
4.2.7	Internationalization	23
4.2.8	Separate Data Modelling from Management Protocol Mapping .	23
4.3	Rejected Objectives	24
4.3.1	Incomplete Translations	24
4.3.2	Attribute Value Constraints	24
4.3.3	Attribute Transaction Constraints	25
4.3.4	Method Constraints	25
4.3.5	Agent Capabilities	25
4.3.6	Relationships	26
4.3.7	Procedures	26
4.3.8	Associations	26
4.3.9	Association Cardinalities	27
4.3.10	Categories of Modules	27
4.3.11	Mapping Modules to Files	27
4.3.12	Simple Grammar	28
4.3.13	Place of Module Information	28
4.3.14	Module Namespace	29
4.3.15	Hyphens in Identifiers	29
5.	Security Considerations	30
6.	Acknowledgements	30
7.	References	30
8.	Authors' Addresses	31
9.	Full Copyright Statement	33

1. Introduction

This document describes the objectives for a new data definition language that can be mapped into SNMP [1], [2] and COPS-PR [3] protocol operations. It may also be translated into SMIv2 [4], [5], [6] MIBs and SPPI [7] PIBs. Concepts such as attributes, attribute groups, methods, conventions for organization into reusable data structures, and mechanisms for representing relationships are discussed.

2. Motivation

As networking technology has evolved, a diverse set of technologies has been deployed to manage the resulting products. These vary from Web based products, to standard management protocols and text scripts. The underlying systems to be manipulated are represented in varying ways including implicitly in the system programming, via proprietary data descriptions, or with standardized descriptions using a range of technologies including MIBs, PIBs, or LDAP schemas. The result is that management interfaces for network protocols, services, and applications such as Differentiated Services may be represented in many different, inconsistent fashions.

The SMIng working group has been chartered to define a new data definition language that will eliminate the need for a separate SMIV2 and SPPI language. That is, the new language should address the needs for the current SMIV2 and SPPI languages so that over time we can all use the new language instead.

Another motivation is to permit a more expressive and complete representation of the modeled information. Examples of additional expressiveness and completeness that are considered are the ability to formally define table existence relationships, the expression of instance creation/deletion capabilities, and the ability to define attribute groups using inheritance. These additional features are discussed in subsequent sections.

It has been recognized that the two main goals of (a) merging SMIV2/SPPI and (b) enhancing the state of art in network management data modeling can lead to conflicts. In such cases, the SMIng working group's consensus is to focus on enhancing the state of art in network management data modeling.

3. Background

The Network Management Research Group (NMRG) of the Internet Research Task Force (IRTF) has researched the issues of creating a protocol-independent data definition language that could be used by multiple protocols. Because SMIV2 and SPPI are very similar, the NMRG focused on merging these two languages, but also researched ways to abstract the objectives to produce a language that could be used for other protocols, such as LDAP and Diameter. The NMRG has published the results of their work in a meanwhile expired Internet Draft, but has submitted their specification as one proposal to consider in the development of the SMIng language.

The SMIng Working Group has accepted their submission for consideration, and to use their proposal to better understand the objectives and possible obstacles to be overcome. Where useful, the NMRG proposal has been referenced in the details below.

4. Specific Objectives for SMIng

The following sections define the objectives for the definition of a new data definition language. The objectives have been organized as follows: accepted objectives (Section 4.1), nice-to-have objectives (Section 4.2), and rejected objectives (Section 4.3). Each objective has the following information:

- o Type: a field that identifies the type of objective, using one of the following values:

- * `basic`: considered a basic objective for SMIng and is contained in SMIV2 and/or SPPI.
 - * `align`: supported in different ways in SMIV2 and SPPI and they must be aligned.
 - * `fix`: considered a fix for a known problem in SMIV2 and/or SPPI.
 - * `new`: considered a new feature.
- o `From`: a field that defines the origin of the objective and that contains one or more of the following values:
 - * `SMI`: exists in SMIV2.
 - * `SPPI`: exists in SPPI.
 - * `NMRG`: exists in the NMRG proposal, but not in SMIV2 or SPPI.
 - * `Charter`: exists in working group charter.
 - * `WG`: proposed during working group discussions.
 - o `Description`: a quick description of the objective.
 - o `Motivation`: rationale for the objective.
 - o `Notes`: optional notes about an objective. For example, for nice-to-have or rejected this may contain reasoning why this objective is not required by the SMIng working group, but justification why it should be considered anyway. Notes may be the opinions of the participants in the discussion on objectives and as such should not be taken as consensus of the working group or the recommendation of the objectives editing team.

4.1 Accepted Objectives

This section represents the list of objectives that have been accepted by the SMIng working group as worthwhile and therefore deserving of further consideration. Each of these objectives must be evaluated by the working group to determine if the benefit incurs an acceptable level of cost. An accepted objective may subsequently be rejected if the cost/benefit analysis determines that the benefit does not justify the cost or that the objective is in direct conflict with one or more other accepted objectives that are deemed more important.

4.1.1 The Set of Specification Documents

Type: new

From: NMRG

Description: SMIV2 is defined in three documents, based on an obsolete ITU ASN.1 specification. SPPI is defined in one document, based on SMIV2. The core of SMIng must be defined in one document and must be independent of external specifications.

Motivation: Self-containment.

4.1.2 Textual Representation

Type: basic

From: SMI, SPPI, WG

Description: SMIng definitions must be represented in a textual format.

Motivation: General IETF consensus.

4.1.3 Human Readability

Type: basic

From: WG

Description: The syntax must make it easy for humans to directly read and write SMIng modules. It must be possible for SMIng module authors to produce SMIng modules with text editing tools.

Motivation: The syntax must make it easy for humans to read and write SMIng modules.

4.1.4 Rigorously Defined Syntax

Type: new

From: NMRG

Description: There must be a rigorously defined syntax for the SMIng language.

Motivation: An unambiguous language promotes consistency across vendors so that different parsers produce the same results. It also provides authoritative rules to SMIng modules designers.

4.1.5 Accessibility

Type: align

From: SMI, SPPI

Description: Attribute definitions must indicate whether attributes can be read, written, created, deleted, and whether they are accessible for notifications, or are not accessible. Align PIB-ACCESS and MAX-ACCESS, and PIB-MIN-ACCESS and MIN-ACCESS.

Motivation: Alignment of SMIV2 and SPPI.

4.1.6 Language Extensibility

Type: new

From: NMRG

Description: The language must have characteristics, so that future modules can contain information of future syntax without breaking original SMIng parsers.

E.g., when SMIV2 introduced REFERENCES it would have been nice if it would not have broken SMIV1 parsers.

Motivation: Achieve language extensibility without breaking core compatibility.

4.1.7 Special Characters in Text

Type: new

From: WG

Description: Allow an escaping mechanism to encode special characters, e.g. double quotes and new-line characters, in text such as DESCRIPTIONS or REFERENCES.

Motivation: ABNF can contain literal characters enclosed in double quotes; to provide the ABNF grammar, there must be the ability to escape special characters.

4.1.8 Naming

Type: basic

From: SMI, SPPI

Description: SMIng must provide mechanisms to uniquely identify attributes, groups of attributes, and events. It is necessary to specify how name collisions are handled.

Motivation: Already in SMIV2 and SPPI.

4.1.9 Namespace Control

Type: basic

From: SMI, SPPI

Description: There must be a hierarchical, centrally-controlled namespace for standard named items, and a distributed namespace must be supported to allow vendor-specific naming and to assure unique module names across vendors and organizations.

Motivation: Need to unambiguously identify definitions of various kinds. Some SMI implementations have problems with different objects from multiple modules but with the same name. Furthermore, the probability of module name clashes rises over time (for example, different vendors defining their own SYSTEM-MIB).

Notes: An example naming scheme is the one employed by the Java programming language with a central naming authority assigning the top-level names.

4.1.10 Modules

Type: basic

From: SMI, SPPI

Description: SMIng must provide a mechanism for uniquely identifying a module, and specifying the status, contact person, revision information, and the purpose of a module.

SMIng must provide mechanisms to group definitions into modules and it must provide rules for referencing definitions from other modules.

Motivation: Modularity and independent advancement of documents.

Notes: Text about module conformance has been moved to Section 4.1.11.

4.1.11 Module Conformance

Type: basic

From: SMI, SPPI

Description: SMIng must provide mechanisms to detail the minimum requirements implementers must meet to claim conformance to a standard based on the module.

Motivation: Ability to convey conformance requirements.

4.1.12 Arbitrary Unambiguous Identities

Type: basic

From: SMI

Description: SMI allows the use of OBJECT-IDENTITIES to define unambiguous identities without the need of a central registry. SMI uses OIDs to represent values that represent references to such identities. SMIng needs a similar mechanism (a statement to register identities, and a base type to represent values).

Motivation: SMI Compatibility.

Notes: This is an obvious objective. Additionally, everything not on the wire, such as modules, will still be assigned OIDs.

It is yet to be determined whether the assignment of the OID occurs within the core or within a protocol-specific mapping.

4.1.13 Protocol Independence

Type: basic

From: Charter

Description: SMIng must define data definitions in support of the SNMP and COPS-PR protocols. SMIng may define data definitions in support of other protocols.

Motivation: So data definitions may be used with multiple protocols and multiple versions of those protocols.

4.1.14 Protocol Mapping

Type: basic

From: Charter

Description: The SMIng working group, in accordance with the working group charter, will define mappings of protocol independent data definitions to protocols based upon installed implementations. The SMIng working group can define mappings to other protocols as long as this does not impede the progress on other objectives.

Motivation: SMIng working group charter.

4.1.15 Translation to Other Data Definition Languages

Type: basic

From: Charter

Description: SMIng language constructs must, wherever possible, be translatable to SMIV2 and SPPI. At the time of standardization of a SMIng language, existing SMIV2 MIBs and SPPI PIBs on the standards track will not be required to be translated to the SMIng language. New MIBs/PIBs will be defined using the SMIng language.

Motivation: Provide best-effort backwards compatibility for existing tools while not placing an unnecessary burden on MIBs/PIBs that are already on the standards track.

4.1.16 Base Data Types

Type: basic

From: SMI, SPPI

Description: SMIng must support the base data types Integer32, Unsigned32, Integer64, Unsigned64, Enumeration, Bits, OctetString, and OID.

Motivation: Most are already common. Unsigned64 and Integer64 are in SPPI, must fix in SMI. Note that Counter and Gauge types can be regarded as derived types instead of base types.

4.1.17 Enumerations

Type: basic

From: SMI, SPPI

Description: SMIng must provide support for enumerations. Enumerated values must be a part of the enumeration definition.

Motivation: SMIV2 already has enumerated numbers.

Notes: Enumerations have the implicit constraint that the attribute is constrained to the values for the enumeration.

4.1.18 Discriminated Unions

Type: new

From: WG

Description: SMIng must support discriminated unions.

Motivation: Allows to group related attributes together, such as InetAddressType (discriminator) and InetAddress, InetAddressIPv4, InetAddressIPv6 (union). The lack of discriminated unions has also lead to relatively complex sparse table work-around in some DISMAN mid-level manager MIBs.

Notes: Discriminated unions have the property that the union attribute type is constrained by the value of the discriminator attribute.

4.1.19 Instance Pointers

Type: basic

From: SMI, SPPI

Description: SMIng must allow specifying pointers to instances (i.e., a pointer to a particular attribute in a row).

Motivation: It is common practice in MIBs and PIBs to point to other instances.

4.1.20 Row Pointers

Type: align

From: SMI, SPPI

Description: SMIng must allow specifying pointers to rows.

Motivation: It is common practice in MIBs and PIBs to point to other rows (see RowPointer, PIB-REFERENCES).

4.1.21 Constraints on Pointers

Type: align

From: SPPI

Description: SMIng must allow specifying the types of objects to which a pointer may point.

Motivation: Allows code generators to detect and reject illegal pointers automatically. Can also be used to automatically generate more reasonable implementation-specific data structures.

Notes: Pointer constraints are a special case of attribute value constraints (Section 4.3.2) in which the prefix of the OID (row or instance pointer) value is limited to be only from a particular table.

4.1.22 Base Type Set

Type: basic

From: SMI, SPPI

Description: SMIng must support a fixed set of base types of fixed size and precision. The list of base types must not be extensible unless the SMI itself changes.

Motivation: Interoperability.

4.1.23 Extended Data Types

Type: align

From: SMI, SPPI

Description: SMIng must support a mechanism to derive new types, which provide additional semantics (e.g., Counters, Gauges, Strings, etc.), from base types. It may be desirable to also allow the derivation of new types from derived types. New types must be as restrictive or more restrictive than the types that they are specializing.

Motivation: SMI uses application types and textual conventions. SPPI uses derived types.

4.1.24 Units, Formats, and Default Values of Defined Types and Attributes

Type: fix

From: NMRG

Description: In SMIv2 OBJECT-TYPE definitions may contain UNITS and DEFVAL clauses and TEXTUAL-CONVENTIONS may contain DISPLAY-HINTS. In a similar fashion units and default values must be applicable to defined types and format information must be applicable to attributes.

Motivation: Some MIBs introduce TCs such as KBytes and every usage of the TC then specifies the UNITS "KBytes". It would simplify things if the UNITS were attached to the type definition itself.

Notes: The SMIng WG must clarify the behavior if an attribute uses a defined type and both, the attribute and the defined type, have units/default/format information.

4.1.25 Table Existence Relationships

Type: align

From: SMI, SPPI

Description: SMIng must support INDEX, AUGMENTS, and EXTENDS in the SNMP/COPS-PR protocol mappings.

Motivation: These three table existence relationships exist either in the SMIv2 or the SPPI.

4.1.26 Table Existence Relationships (2)

Type: new

From: NMRG

Description: SMIng must support EXPANDS and REORDERS relationships in the SNMP/COPS-PR protocol mappings.

Motivation: A REORDERS statement allows indexing orders to be swapped. An EXPANDS statement formally states that there is a 1:n existence relationship between table rows.

4.1.27 Attribute Groups

Type: new

From: NMRG

Description: An attribute group is a named, reusable set of attributes that are meaningful together. It can be reused as the type of attributes in other attribute groups (see also Section 4.1.28). This is similar to 'structs' in C.

Motivation: Required to map the same grouping of attributes into SNMP and COPS-PR tables. Allows to do index reordering without having to redefine the attribute group. Allows to group related attributes together (e.g. InetAddressType, InetAddress).

The ability to group attributes provides an indication that the attributes are meaningful together.

4.1.28 Containment

Type: new

From: NMRG

Description: SMIng must provide support for the creation of new attribute groups from attributes of more basic types and potentially other attribute groups.

Motivation: Simplifies the reuse of attribute groups such as InetAddressType and InetAddress pairs.

Notes: Containment has the implicit existence constraint that if an instance of a contained attribute group exists, then the corresponding instance of the containing attribute group must also exist.

4.1.29 Single Inheritance

Type: new

From: NMRG

Description: SMIng must provide support for mechanisms to extend attribute groups through single inheritance.

Motivation: Allows to extend attribute groups, like a generic DiffServ scheduler, with attributes for a specific scheduler, without cut&paste.

Notes: Single inheritance with multiple levels (e.g., C derives from B, and B derives from A) must be allowed.

Inheritance has the implicit existence constraint that if an instance of a derived attribute group exists, then the corresponding instance of the base attribute group must also exist.

Inheritance could help to add attributes to an attribute group that are specific to a certain protocol mapping and do not appear in the protocol-neutral attribute group.

4.1.30 Reusable vs. Final Attribute Groups

Type: new

From: NMRG, WG

Description: SMIng must differentiate between "final" and reusable attribute groups, where the reuse of attribute groups covers inheritance and containment.

Motivation: This information gives people more information how attribute groups can and should be used. It hinders them from misusing them.

Notes: This objective attempts to convey the idea that some attribute groups are not meant to stand on their own and instead only make sense if contained within another attribute group.

4.1.31 Events

Type: basic

From: SMI, SPPI

Description: SMIng must provide mechanisms to define events which identify significant state changes.

Motivation: These represent the protocol-independent events that lead to SMI notifications or SPPI reports.

4.1.32 Creation/Deletion

Type: align

From: SMI, SPPI

Description: SMIng must support a mechanism to define creation/deletion operations for instances. Specific creation/deletion errors, such as INSTALL-ERRORS, must be supported.

Motivation: Available for row creation in SMI, and available in SPPI.

4.1.33 Range and Size Constraints

Type: basic

From: SMI, SPPI

Description: SMIng must allow specifying range and size constraints where applicable.

Motivation: The SMI and SPPI both support range and size constraints.

4.1.34 Uniqueness

Type: basic

From: SPPI

Description: SMIng must allow the specification of uniqueness constraints on attributes. SMIng must allow the specification of multiple independent uniqueness constraints.

Motivation: Knowledge of the uniqueness constraints on attributes allows to verify protocol specific mappings (e.g. INDEX clauses). The knowledge can also be used by code generators to improve generated implementation-specific data structures.

4.1.35 Extension Rules

Type: basic

From: SMI, SPPI

Description: SMIng must provide clear rules how one can extend SMIng modules without causing interoperability problems "over the wire".

Motivation: SMIV2 and SPPI have extension rules.

4.1.36 Deprecate Use of IMPLIED Keyword

Type: fix

From: WG

Description: The SMIng SNMP mapping must deprecate the use of the IMPLIED indexing schema.

Motivation: IMPLIED is confusing and most people don't understand it. The solution (IMPLIED) is worse than the problem it is trying to solve and therefore for the sake of simplicity, the use of IMPLIED should be deprecated.

4.1.37 No Redundancy

Type: fix

From: NMRG

Description: The SMIng language must avoid redundancy.

Motivation: Remove any textual redundancy for things like table entries and SEQUENCE definitions, which only increase specifications without providing any value.

4.1.38 Compliance and Conformance

Type: basic

From: SMI, SPPI

Description: SMIng must provide a mechanism for compliance and conformance specifications for protocol-independent definitions as well as for protocol mappings.

Motivation: This capability exists in SMIV2 and SPPI. The NMRG proposal has the ability to express much of this information at the protocol-dependent layer. Some compliance or conformance information may be protocol-independent, therefore there is also a need to be able to express this information protocol-independent part.

4.1.39 Allow Refinement of All Definitions in Conformance Statements

Type: fix

From: WG

Description: SMIV2, RFC 2580, Section 3.1 says:

The OBJECTS clause, which must be present, is used to specify each object contained in the conformance group. Each of the specified objects must be defined in the same information module as the OBJECT-GROUP macro appears, and must have a MAX-ACCESS clause value of "accessible-for-notify", "read-only", "read-write", or "read-create".

The last sentence forbids to put a not-accessible INDEX object into an OBJECT-GROUP. Hence, you can not refine its syntax in a compliance definition. For more details, see <http://www.ibr.cs.tu-bs.de/ietf/smi-errata/>

Motivation: This error should not be repeated in SMIng.

4.1.40 Categories

Type: basic

From: SPPI

Description: SMIng must provide a mechanism to group definitions into subject categories. Concrete instances may only exist in the scope of a given subject category or context.

Motivation: To scope the categories to which a module applies. In SPPI this is used to allow a division of labor between multiple client types.

4.1.41 Core Language Keywords vs. Defined Identifiers

Type: fix

From: NMRG

Description: In SMI and SPPI modules some language keywords (macros and a number of basetypes) have to be imported from different SMI language defining modules, e.g. OBJECT-TYPE, MODULE-IDENTITY, Integer32 must to be imported from SNMPv2-SMI and TEXTUAL-CONVENTION must be imported from SNMPv2-TC, if used. MIB authors are continuously confused about these import rules. In SMIng only defined identifiers must be imported. All SMIng language keywords must be implicitly known and there must not be a need to import them from any module.

Motivation: Reduce confusion. Clarify the set of language keywords.

4.1.42 Instance Naming

Type: align

From: SMI, SPPI

Description: Instance naming in SMIV2 and SPPI is different. SMIng must align the instance naming (either in the protocol neutral model or the protocol mappings).

Motivation: COPS-PR and SNMP have different instance identification schemes that must be handled.

Notes: A solution requires to investigate how close the naming schemes dictated by the protocols are. Perhaps it is feasible to have a single instance naming scheme in both SNMP and COPS-PR, even though the current SPPI and SMIV2 are different.

4.1.43 Length of Identifiers

Type: fix

From: NMRG

Description: The allowed length of the various kinds of identifiers must be extended from the current 'should not exceed 32' (maybe even from the 'must not exceed 64') rule.

Motivation: Reflect current practice of definitions.

Notes: The 32-rule was added back in the days where compilers could not deal with long identifiers. This rule is continuously violated these days and it does not make sense to keep it.

4.1.44 Assign OIDs in the Protocol Mappings

Type: new

From: NMRG

Description: SMIng must not assign OIDs to reusable definition of attributes, attribute groups, events, etc. Instead, SNMP and COPS-PR mappings must assign OIDs to the mapped items.

Motivation: Assignment of OIDs in protocol neutral definitions can complicate reuse. OIDs of synonymous attributes are not the same in SMI and SPPI definitions. MIBs and PIBs are already registered in different parts of the OID namespace.

4.2 Nice-to-Have Objectives

This section represents the list of recommended objectives that would be nice to have. However, these are not automatically thought of as accepted objectives as, for example, they may entail a non-trivial amount of work in underlying protocols to support or they may be regarded as less important than other contradicting objectives that are accepted.

4.2.1 Methods

Type: new

From: WG

Description: SMIng should support a mechanism to define method signatures (parameters, return values, exception) that are implemented on agents.

Motivation: Methods are needed to support the definition of operational interfaces such as found in [RFC2925] (ping, traceroute and lookup operations). Also, the ability to define constructor/destructor interfaces could address issues such as encountered with SNMP's RowStatus solution.

Notes: Is it possible to do methods without changing the underlying protocol? There is agreement that methods are useful, but disagreement upon the impact - one end of the spectrum sees this as a documentation tool for existing SNMP capabilities, while the

other end sees this as a protocol update, moving forward, to natively support methods. The proposal is to wait and see if this is practical to implement as a syntax that is useful and can map to the protocol.

4.2.2 Unions

Type: new

From: WG

Description: SMIng should support a standard format for unions.

Motivation: Allows an attribute to contain one of many types of values. The lack of unions has also lead to relatively complex sparse table work-around in some DISMAN mid-level managers. Despite from discriminated unions (see Section 4.1.18), this kind of union has no accompanied explicit discriminator attribute that selects the union's type of value.

Notes: The thought is that SNMP and COPS-PR can already support unions because they do not care about what data type goes with a particular OID.

4.2.3 Float Data Types

Type: new

From: WG, NMRG

Description: SMIng should support the base data types Float32, Float64, Float128.

Motivation: Missing base types can hurt later on, because they cannot be added without changing the language, even as an SMIng extension. Lesson learned from the SMIV1/v2 debate about Counter64/Integer64/...

Notes: There is no mention as to whether or not the underlying protocols will have to natively support float data types. This is left to the mapping. However, it seems imperative that the float data type needs to be added to the set of intrinsic types in the SMIng language at the creation of the language as it will be impossible to add them later without changing the language.

4.2.4 Comments

Type: fix

From: NMRG

Description: The syntax of comments should be well defined, unambiguous and intuitive to most people, e.g., the C++/Java '//' syntax.

Motivation: ASN.1 Comments (and thus SMI and SPPI comments) have been a constant source of confusion. People use arbitrary lengthy strings of dashes ('-----') in the wrong assumption that this is always treated as a comment. Some implementations try to accept these syntactically wrong constructs which even raises confusion. We should get rid of this problem.

Notes: If the SMIng working group adopts a C-like syntax, then the C++/Java single-line comment should be adopted as well.

4.2.5 Referencing Tagged Rows

Type: align

From: SPPI

Description: PIB and MIB row attributes reference a group of entries in another table. SPPI formalizes this by introducing PIB-TAG and PIB-REFERENCES clauses. This functionality should be retained in SMIng.

Motivation: SPPI formalizes tag references. Some MIBs also use tag references (see SNMP-TARGET-MIB in RFC2573) even though SMIV2 does not provide a formal notation.

4.2.6 Arrays

Type: new

From: WG

Description: SMIng should allow the definition of a SEQUENCE OF attributes or attribute groups (Section 4.1.27).

Motivation: The desire for the ability to have variable-length, multi-valued objects.

Notes: Some issues with arrays are still unclear. As long as there are no concepts to solve the problems with access semantics (how to achieve atomic access to arbitrary-sized arrays) and their mappings to SNMP and COPS-PR protocol operations, arrays cannot be more than a nice to have objective.

4.2.7 Internationalization

Type: new

From: WG

Description: Informational text (DESCRIPTION, REFERENCE, ...) should allow i18nized encoding, probably UTF-8.

Motivation: There has been some demand for i18n in the past. The BCP RFC 2277 demands for internationalization.

Notes: Although English is the language of IETF documents, SMIng should allow other languages for private use.

4.2.8 Separate Data Modelling from Management Protocol Mapping

Type: new

From: NMRG

Description: It should be possible to separate the domain specific data modelling work from the network management protocol specific work.

Motivation: Today, working groups designing new protocols are forced to care about the design of SNMP MIBs and maybe COPS-PR PIBs to manage the new protocol. This means that experts in a specific domain are faced with details of at least one foreign (network management) technology. This leads to hard work and long revision processes. It would be a win to separate the task of pure data modelling which can be done by the domain experts easily from the network management protocol specific mappings. The mapping to SNMP and/or COPS-PR can be done (a) later separately and (b) by network management experts. This required NM expertise no longer hinders the progress of the domain specific working groups.

4.3 Rejected Objectives

This section represents the list of objectives that were rejected during the discussion on the objectives. Those objectives that have been rejected need not be addressed by SMIng. This does not imply that they must not be addressed.

4.3.1 Incomplete Translations

Type: basic

From: WG

Description: Reality sucks. All information expressed in SMIng may not be directly translatable to a MIB or PIB construct, but all information should be able to be conveyed in documentation or via other mechanisms.

Motivation: SMIng working group requires this to ease transition.

Notes: The SMIng language itself cannot require what compilers do that translate SMIng into something else. So this seems to fall out of the scope of the current working group charter.

4.3.2 Attribute Value Constraints

Type: new

From: WG

Description: SMIng should provide mechanisms to formally specify constraints between values of multiple attributes.

Motivation: Constraints on attribute values occur where one or more attributes may affect the value or range of values for another attribute. One such relationship exists in IPsec, where the type of security algorithm determines the range of possible values for other attributes such as the corresponding key size.

Notes: This objective as is has been rejected as too general, and therefore virtually impossible to implement. However, constraints that are implicit with discriminated unions (Section 4.1.18), enumerated types (Section 4.1.17), pointer constraints (Section 4.1.21)), etc., are accepted and these implicit constraints are mentioned in the respective objectives.

4.3.3 Attribute Transaction Constraints

Type: new

From: WG

Description: SMIng should provide a mechanism to formally express that certain sets of attributes can only be modified in combination.

Motivation: COPS-PR always does operations on table rows in a single transaction. There are SMIV2 attribute combinations that need to be modified together (such as InetAddressType, InetAddress).

Notes: Alternative is to either use Methods (Section 4.2.1) or assume that all attributes in an attribute group (Section 4.1.27) are to be considered atomic.

4.3.4 Method Constraints

Type: new

From: WG

Description: Method definitions should provide constraints on parameters.

Motivation: None.

Notes: Unless methods (Section 4.2.1) are done, there is no use for this. Furthermore, this objective has not been motivated by any proponent.

4.3.5 Agent Capabilities

Type: basic

From: SMI

Description: SMIng should provide mechanisms to describe agent implementations.

Motivation: To permit manager to determine variations from the standard for an implementation.

Notes: Agent capabilities should not be part of SMIng, but should instead be a separate capabilities table.

4.3.6 Relationships

Type: new

From: NMRG, WG

Description: Ability to formally depict existence dependency, value dependency, aggregation, containment, and other relationships between attributes or attribute groups.

Motivation: Helps humans to understand the conceptual model of a module. Helps implementers of MIB compilers to generate more 'intelligent' code.

Notes: This objective was deemed too general to be useful and instead the individual types of relationship objectives (e.g., pointers, inheritance, containment, etc.) are evaluated on a case-by-case basis with the specific relationships deemed useful being included as accepted objectives.

4.3.7 Procedures

Type: new

From: WG

Description: SMIng should support a mechanism to formally define procedures that are used by managers when interacting with an agent.

Motivation: None.

Notes: This objective has not been motivated by any proponent.

4.3.8 Associations

Type: new

From: WG

Description: SMIng should provide mechanisms to explicitly specify associations.

Motivation: None.

Notes: This objective has not been motivated by any proponent.

4.3.9 Association Cardinalities

Type: new

From: WG

Description: Cardinalities between associations should be formally defined.

Motivation: If you have an association between attribute groups A and B, the cardinality of A indicates how many instances of A may be associated with a single instance of B. Our discussions in Minneapolis indicated that we want to convey "how many" instances are associated in order to define the best mapping algorithm - whether a new table, a single pointer, etc. For example, do we use RowPointer or an integer index into another table? Do we map to a table that holds instances of the association/relationship itself?

Notes: Without associations (Section 4.3.8), this has no use.

4.3.10 Categories of Modules

Type: new

From: WG

Description: The SMIng documents should give clear guidance on which kind of information (with respect to generality, type/attribute group/extension/..) should be put in which kind of a module.

E.g., in SMIV2 we don't like to import Utf8String from SYSAPPL-MIB, but we also do not like to introduce a redundant definition.

A module review process should probably be described that ensures that generally useful definitions do not go into device or service specific modules.

Motivation: Bad experience with SMIV2.

Notes: It is not clear how this can be done with the language to be created by SMIng WG.

4.3.11 Mapping Modules to Files

Type: new

From: NMRG

Description: There should be a clear statement how SMIng modules are mapped to files (1:1, n:1?) and how files should be named (by module name in case of 1:1 mapping?).

Motivation: SMI implementations show up a variety of filename extensions (.txt, .smi, .my, none). Some expect all modules in a single file, others don't. This makes it more difficult to exchange modules.

Notes: This is just an implementation detail and is best left to a BCP and not made a part of the language definition.

4.3.12 Simple Grammar

Type: new

From: NMRG

Description: The grammar of the language should be as simple as possible. It should be free of exception rules. A measurement of simplicity is shortness of the ABNF grammar.

Motivation: Ease of implementation. Ease of learning/understanding.

Notes: This seems like an obvious objective, however shortness of the ABNF grammar is not necessarily a reflection of the simplicity of the grammar.

4.3.13 Place of Module Information

Type: fix

From: NMRG

Description: Module specific information (organization, contact, description, revision information) should be bound to the module itself and not to an artificial node (like SMIV2 MODULE-IDENTITY).

Motivation: Simplicity and design cleanup.

Notes: This does not seem to be a problem with the current SMI. Although simplification is a good thing, this detail is not considered an objective.

4.3.14 Module Namespace

Type: new

From: WG

Description: Currently the namespace of modules is flat and there is no structure in module naming causing the potential risk of name clashes. Possible solutions:

- * Assume module names are globally unique (just as SMIV1/v2), just give some recommendations on module names.
- * Force all organizations, WGs and vendors to apply a name prefix (e.g. CISCO-GAGA-MIB, IETF-DISMAN-SCRIPT-MIB?).
- * Force enterprises to apply a prefix based on the enterprise number (e.g. ENT2021-SOME-MIB).
- * Put module names in a hierarchical domain based namespace (e.g. DISMAN-SCRIPT-MIB.ietf.org).

Motivation: Reduce risk of module name clashes.

Notes: Some aspects of this objective overlap with other objectives (namespace control (Section 4.1.9)) and other aspects were thought best left to a BCP.

4.3.15 Hyphens in Identifiers

Type: fix

From: NMRG

Description: There has been some confusion whether hyphens are allowed in SMIV2 identifiers: Module names are allowed to contain hyphens. Node identifiers usually are not. But for example 'mib-2' is a frequently used identifier that contains a hyphen due to its SMIV1 origin, when hyphen were not disallowed. Similarly, a number of named numbers of enumeration types contain hyphens violating an SMIV2 rule.

SMIng should simply allow hyphens in all kinds of identifiers. No exceptions.

Motivation: Reduce confusion and exceptions. Requires, however, that implementation mappings properly quote hyphens where appropriate.

Notes: This nit-picking is not worth to be subject to the discussion on objectives. However, SMIng should care about the fact that compilers have to map SMIng to programming languages where a hyphen is a minus and thus not allowed in identifiers.

5. Security Considerations

This document defines objectives for a language with which to write and read descriptions of management information. The language itself has no security impact on the Internet.

6. Acknowledgements

Thanks to Dave Durham, whose work on the original NIM (Network Information Model) draft was used in generating this document.

Thanks to Andrea Westerinen for her contributions on the original NIM requirements and SMIng objectives drafts.

7. References

- [1] Case, J., Fedor, M., Schoffstall, M. and J. Davin, "Simple Network Management Protocol (SNMP)", STD 15, RFC 1157, May 1990.
- [2] McCloghrie, K., Case, J., Rose, M. and S. Waldbusser, "Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1905, January 1996.
- [3] Chan, K., Seligson, J., Durham, D., Gai, S., McCloghrie, K., Herzog, S., Reichmeyer, F., Yavatkar, R. and A. Smith, "COPS Usage for Policy Provisioning (COPS-PR)", RFC 3084, March 2001.
- [4] McCloghrie, K., Perkins, D., Schoenwaelder, J., Case, J., Rose, M. and S. Waldbusser, "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.
- [5] McCloghrie, K., Perkins, D., Schoenwaelder, J., Case, J., Rose, M. and S. Waldbusser, "Textual Conventions for SMIv2", STD 58, RFC 2579, April 1999.
- [6] McCloghrie, K., Perkins, D. and J. Schoenwaelder, "Conformance Statements for SMIv2", STD 58, RFC 2580, April 1999.
- [7] McCloghrie, K., Fine, M., Seligson, J., Chan, K., Hahn, S., Sahita, R., Smith, A. and F. Reichmeyer, "Structure of Policy Provisioning Information (SPPI)", RFC 3159, August 2001.

8. Authors' Addresses

Chris Elliott
Cisco Systems
7025 Kit Creek Road
Research Triangle Park, NC 27709
USA

EMail: chelliot@cisco.com

David Harrington
Enterasys Networks
35 Industrial Way
P.O. Box 5005
Rochester, NH 03866-5005
USA

EMail: dbh@enterasys.com

Jamie Jason
Intel Corporation
MS JF3-206
2111 NE 25th Ave.
Hillsboro, OR 97124
USA

EMail: jamie.jason@intel.com

Juergen Schoenwaelder
TU Braunschweig
Muehlenpfordtstr. 23
38106 Braunschweig
Germany

EMail: schoenw@ibr.cs.tu-bs.de
URI: <http://www.ibr.cs.tu-bs.de/>

Frank Strauss
TU Braunschweig
Muehlenpfordtstr. 23
38106 Braunschweig
Germany

E-Mail: strauss@ibr.cs.tu-bs.de
URI: <http://www.ibr.cs.tu-bs.de/>

Walter Weiss
Ellacoya Networks
7 Henry Clay Dr.
Merrimack, NH. 03054
USA

E-Mail: wweiss@ellacoya.com

9. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

