Input Pattern Types

arg-type

```
Pattern:
'[^\*].*|..+'; // must not be single '*'
```

```
Pattern:
'\*'
```

This statement can be used to hide a node from some, or all, northbound interfaces. All nodes with the same value are considered a hide group and are treated the same with regards to being visible or not in a northbound interface.

A node with an hidden property is not shown in the northbound user interfaces (CLI and Web UI) unless an 'unhide' operation is performed in the user interface.

The hidden value 'full' indicates that the node must be hidden from all northbound interfaces, including programmatical interfaces such as NETCONF. The value '*' is not valid. A hide group can be unhidden only if this is explicitly allowed in the confd.conf(5) daemon configuration.

Multiple hide groups can be specified by giving this statement multiple times. The node is shown if any of the specified hide groups is given in the 'unhide' operation. If a mandatory node is hidden, a hook callback function (or similar) might be needed in order to set the element

crypt-hash

```
Pattern:
'$0$.*'
 '|$1$[a-zA-Z0-9./]{1,8}$[a-zA-Z0-9./]{22}'
 '|$5$(rounds=\d+$)?[a-zA-Z0-9./]{1,16}$[a-zA-Z0-9./]{43}'
 '|$6$(rounds=\d+$)?[a-zA-Z0-9./]{1,16}$[a-zA-Z0-9./]{86}'
```

The **crypt-hash** type is used to store passwords using a hash function. The algorithms for applying the hash function and encoding the result are implemented in various UNIX systems as the function crypt(3).

A value of this type matches one of the forms:

- 0<clear text password>

- $<id>$<salt>$<password hash>

- $<id>$<parameter>$<salt>$<password hash>

The '0' prefix signals that the value is clear text. When such a value is received by the server, a hash value is calculated, and the string '$<id>$<salt>$' or $<id>$<parameter>$<salt>$ is prepended to the result. This value is stored in the configuration data store.

If a value starting with '$<id>$', where <id> is not '0', is received, the server knows that the value already represents a hashed value, and stores it as is in the data store.

When a server needs to verify a password given by a user, it finds the stored password hash string for that user, extracts the salt, and calculates the hash with the salt and given password as input. If the calculated hash value is the same as the stored value, the password given by the client is accepted.

This type defines the following hash functions:

Id	Hash Function	Feature
1	MD5	crypt-hash-md5
5	SHA-256	crypt-hash-sha-256
6	SHA-512	crypt-hash-sha-512

The server indicates support for the different hash functions by advertising the corresponding feature.

Reference:

- IEEE Std 1003.1-2008 - crypt() function

- RFC 1321: The MD5 Message-Digest Algorithm

- FIPS.180-3.2008: Secure Hash Standard

date-and-time

```
Pattern:
'\d{4}-\d{2}-\d{2}T\d{2}:\d{2}:\d{2}(\.\d+)?'
 '(Z|[\+\-]\d{2}:\d{2})'
```

The date-and-time type is a profile of the ISO 8601 standard for representation of dates and times using the Gregorian calendar. The profile is defined by the date-time production in Section 5.6 of RFC 3339. The date-and-time type is compatible with the dateTime XML schema type with the following notable exceptions:

1. The date-and-time type does not allow negative years.

2. The date-and-time time-offset -00:00 indicates an unknown time zone (see RFC 3339) while -00:00 and +00:00 and Z all represent the same time zone in dateTime.

3. The canonical format (see below) of data-and-time values differs from the canonical format used by the dateTime XML schema type, which requires all times to be in UTC using the time-offset 'Z'.

This type is not equivalent to the DateAndTime textual convention of the SMIv2 since RFC 3339 uses a different separator between full-date and full-time and provides higher resolution of time-secfrac. The canonical format for date-and-time values with a known time zone uses a numeric time zone offset that is calculated using the device's configured known offset to UTC time.

A change of the device's offset to UTC time will cause date-and-time values to change accordingly. Such changes might happen periodically in case a server follows automatically daylight saving time (DST) time zone offset changes. The canonical format for date-and-time values with an unknown time zone (usually referring to the notion of local time) uses the time-offset -00:00.

Reference:

- RFC 3339: Date and Time on the Internet: Timestamps

- RFC 2579: Textual Conventions for SMIv2

- XSD-TYPES: XML Schema Part 2: Datatypes Second Edition

domain-name

```
Pattern:
'((([a-zA-Z0-9_]([a-zA-Z0-9\-_]){0,61})?[a-zA-Z0-9]\.)*'
 '([a-zA-Z0-9_]([a-zA-Z0-9\-_]){0,61})?[a-zA-Z0-9]\.?)'
 '|\.'
```

The domain-name type represents a DNS domain name. The name must fully qualified whenever possible. Internet domain names are only loosely specified. Section 3.5 of RFC 1034 recommends a syntax (modified in Section 2.1 of RFC 1123). The Pattern above is intended to allow for current practice in domain name use, and some possible future expansion. It is designed to hold various types of domain names, including names used for A or AAAA records (host names) and other records, such as SRV records.

The Internet host names have a stricter syntax (described in RFC 952) than the DNS recommendations in RFCs 1034 and 1123, and that systems that want to store host names in schema nodes using the domain-name type are recommended to adhere to this stricter standard to ensure interoperability.

The encoding of DNS names in the DNS protocol is limited to 255 characters. Since the encoding consists of labels prefixed by a length bytes and there is a trailing NULL byte, only 253 characters can appear in the textual dotted notation.

The description clause of schema nodes using the domain-name type must describe when and how these names are resolved to IP addresses. The resolution of a domain-name value may require to query multiple DNS records. For example, A for IPv4 and AAAA for IPv6. The order of the resolution process and which DNS record takes precedence can either be defined explicitly or may depend on the configuration of the resolver.

Domain-name values use the US-ASCII encoding. Their canonical format uses lowercase US-ASCII characters. Internationalized domain names MUST be A-labels as per RFC 5890.

Reference:

- RFC 952: DoD Internet Host Table Specification

- RFC 1034: Domain Names - Concepts and Facilities

- RFC 1123: Requirements for Internet Hosts -- Application and Support

- RFC 2782: A DNS RR for specifying the location of services (DNS SRV)

- RFC 5890: Internationalized Domain Names in Applications (IDNA): Definitions and Document Framework

dotted-quad

Pattern:
```
'(([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])\.){3}'
 '([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])'
```

An unsigned 32-bit number expressed in the dotted-quad notation, that is, four octets written as decimal numbers and separated with the '.' (full stop) character.

hex-list

Pattern:
`'(([0-9a-fA-F]){2}(:([0-9a-fA-F]){2})*)?'`

DEPRECATED: Use yang:hex-string instead. There are no plans to remove tailf:hex-list. A list of colon-separated hexa-decimal octets, for example '4F:4C:41:71'.

The statement tailf:value-length can be used to restrict the number of octets. Using the 'length' restriction limits the number of characters in the lexical representation

hex-string

Pattern:
'([0-9a-fA-F]{2}(:[0-9a-fA-F]{2})*)?'

A hexadecimal string with octets represented as hex digits separated by colons. The canonical representation uses lowercase characters.

ipv4-address

```
Pattern:
'(([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])\.){3}'
 '([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])'
 '(%[\p{N}\p{L}]+)?'
```

The ipv4-address type represents an IPv4 address in dotted-quad notation. The IPv4 address may include a zone index, separated by a % sign. The zone index is used to disambiguate identical address values. For link-local addresses, the zone index will typically be the interface index number or the name of an interface. If the zone index is not present, the default zone of the device will be used. The canonical format for the zone index is the numerical format.

ipv4-address-and-prefix-length

Pattern:
```
'(([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])\.){3}'
 '([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])'
 '/(([0-9])|([1-2][0-9])|(3[0-2]))'
```

The ipv4-address-and-prefix-length type represents a combination of an IPv4 address and a prefix length. The prefix length is given by the number following the slash character and must be less than or equal to 32.

ipv4-address-no-zone

Pattern:
`'[0-9\.]*'`

An IPv4 address is without a zone index and derived from ipv4-address that is used in situations where the zone is known from the context and hence no zone index is needed.

ipv4-prefix

```
Pattern:
'(([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])\.){3}'
 '([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])'
 '/(([0-9])|([1-2][0-9])|(3[0-2]))'
```

The ipv4-prefix type represents an IPv4 address prefix. The prefix length is given by the number following the slash character and must be less than or equal to 32.

A prefix length value of 'n' corresponds to an IP address mask that has n contiguous 1-bits from the most significant bit (MSB) and all other bits set to 0.

The canonical format of an IPv4 prefix has all bits of the IPv4 address set to zero that are not part of the IPv4 prefix.

ipv6-address

```
Pattern:
'((:|[0-9a-fA-F]{0,4}):)([0-9a-fA-F]{0,4}:){0,5}'
  '((([0-9a-fA-F]{0,4}:)?(:|[0-9a-fA-F]{0,4}))|'
  '(((25[0-5]|2[0-4][0-9]|[01]?[0-9]?[0-9])\.){3}'Pattern:
  '(25[0-5]|2[0-4][0-9]|[01]?[0-9]?[0-9])))'
  '(%[\p{N}\p{L}]+)?'
```

```
Pattern:
'(([^:]+:){6}(([^:]+:[^:]+)|(.*\..*)))|'
  '((([^:]+:)*[^:]+)?::(([^:]+:)*[^:]+)?)'
  '(%.+)?'
```

The ipv6-address type represents an IPv6 address in full, mixed, shortened, and shortened-mixed notation. The IPv6 address may include a zone index, separated by a % sign.

The zone index is used to disambiguate identical address values. For link-local addresses, the zone index will typically be the interface index number or the name of an interface. If the zone index is not present, the default zone of the device will be used.

The canonical format of IPv6 addresses uses the textual representation defined in Section 4 of RFC 5952. The canonical format for the zone index is the numerical format as described in Section 11.2 of RFC 4007.

Reference:

- RFC 4291: IP Version 6 Addressing Architecture

- RFC 4007: IPv6 Scoped Address Architecture

- RFC 5952: A Recommendation for IPv6 Address Text Representation

ipv6-address-and-prefix-length

```
Pattern:
'((:|[0-9a-fA-F]{0,4}):)([0-9a-fA-F]{0,4}:){0,5}'
 '((([0-9a-fA-F]{0,4}:)?(:|[0-9a-fA-F]{0,4}))|'
 '(((25[0-5]|2[0-4][0-9]|[01]?[0-9]?[0-9])\.){3}'
 '(25[0-5]|2[0-4][0-9]|[01]?[0-9]?[0-9])))'
 '(/(([0-9])|([0-9]{2})|(1[0-1][0-9])|(12[0-8])))'

Pattern:
'(([^:]+:){6}(([^:]+:[^:]+)|(.*\..*)))|'
 '((([^:]+:)*[^:]+)?::(([^:]+:)*[^:]+)?)'
 '(/.+)'
```

The ipv6-address-and-prefix-length type represents a combination of an IPv6 address and a prefix length. The prefix length is given by the number following the slash character and must be less than or equal to 128.

ipv6-address-no-zone

Pattern:
'[0-9a-fA-F:\.]*'

An IPv6 address without a zone index. This type, derived from ipv6-address, may be used in situations where the zone is known from the context and hence no zone index is needed.

Reference:

- RFC 4291: IP Version 6 Addressing Architecture

- RFC 4007: IPv6 Scoped Address Architecture

- RFC 5952: A Recommendation for IPv6 Address Text Representation

ipv6-prefix

```
Pattern:
'((:|[0-9a-fA-F]{0,4}):)([0-9a-fA-F]{0,4}:){0,5}'
 '((([0-9a-fA-F]{0,4}:)?(:|[0-9a-fA-F]{0,4}))|'
 '(((25[0-5]|2[0-4][0-9]|[01]?[0-9]?[0-9])\.){3}'Pattern:
 '(25[0-5]|2[0-4][0-9]|[01]?[0-9]?[0-9])))'
 '(/(([0-9])|([0-9]{2})|(1[0-1][0-9])|(12[0-8])))';
```

```
Pattern:
'(([^:]+:){6}(([^:]+:[^:]+)|(.*\..*)))|'
 '((([^:]+:)*[^:]+)?::(([^:]+:)*[^:]+)?)'
 '(/.+)'
```

The ipv6-prefix type represents an IPv6 address prefix. The prefix length is given by the number following the slash character and must be less than or equal to 128.

A prefix length value of n corresponds to an IP address mask that has n contiguous 1-bits from the most significant bit (MSB) and all other bits set to 0.

The IPv6 address should have all bits that do not belong to the prefix set to zero. The canonical format of an IPv6 prefix has all bits of the IPv6 address set to zero that are not part of the IPv6 prefix. Furthermore, the IPv6 address is represented as defined in Section 4 of RFC 5952

Reference:

• RFC 5952: A Recommendation for IPv6 Address Text Representation

mac-address

Pattern:
`'[0-9a-fA-F]{2}(:[0-9a-fA-F]{2}){5}'`

The mac-address type represents an IEEE 802 MAC address. The canonical representation uses lowercase characters. In the value set and its semantics, this type is equivalent to the MacAddress textual convention of the SMIv2.

Reference:

- IEEE 802: IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture

- RFC 2579: Textual Conventions for SMIv2

object-identifier

Pattern:
```
'(([0-1](\.[1-3]?[0-9]))|(2\.(0|([1-9]\d*))))'
 '(\.(0|([1-9]\d*)))*'
```

The object-identifier type represents administratively assigned names in a registration-hierarchical-name tree. The values of this type are denoted as a sequence of numerical non-negative sub-identifier values. Each sub-identifier value MUST NOT exceed $2^{32}-1$ (4294967295). The Sub-identifiers are separated by single dots and without any intermediate whitespace.

The ASN.1 standard restricts the value space of the first sub-identifier to 0, 1, or 2. Furthermore, the value space of the second sub-identifier is restricted to the range 0 to 39 if the first sub-identifier is 0 or 1. Finally, the ASN.1 standard requires that an object identifier has always at least two sub-identifiers. The pattern captures these restrictions.

Although the number of sub-identifiers is not limited, module designers should realize that there may be implementations that stick with the SMIv2 limit of 128 sub-identifiers.

This type is a superset of the SMIv2 OBJECT IDENTIFIER type since it is not restricted to 128 sub-identifiers. Hence, this type SHOULD NOT be used to represent the SMIv2 OBJECT IDENTIFIER type; the object-identifier-128 type SHOULD be used instead.

Reference:

- ISO9834-1: Information technology - Open Systems

- Interconnection - Procedures for the operation of OSI

- Registration Authorities: General procedures and top arcs of the ASN.1 Object Identifier tree

object-identifier-128

Pattern:
`'\d*(\.\d*){1,127}'`

This type represents object-identifiers restricted to 128 sub-identifiers. In the value set and its semantics, this type is equivalent to the OBJECT IDENTIFIER type of the SMIv2.

Reference:

- RFC 2578: Structure of Management Information Version 2 (SMIv2)

octet-list

Pattern:
`'(\d*(.\d*)*)?'`

A list of dot-separated octets, for example '192.168.255.1.0'. The statement tailf:value-length can be used to restrict the number of octets. Using the 'length' restriction limits the number of characters in the lexical representation.

phys-address

Pattern:
'([0-9a-fA-F]{2}(:[0-9a-fA-F]{2})*)?'

Represents media- or physical-level addresses represented as a sequence octets, each octet represented by two hexadecimal numbers. Octets are separated by colons. The canonical representation uses lowercase characters. In the value set and its semantics, this type is equivalent to the PhysAddress textual convention of the SMIv2.

Reference:

- RFC 2579: Textual Conventions for SMIv2

sha-256-digest-string

Pattern:
```
'$0$.*'
 '|$5$(rounds=\d+$)?[a-zA-Z0-9./]{1,16}$[a-zA-Z0-9./]{43}'
```

The sha-256-digest-string type automatically computes a SHA-256 digest for a value adhering to this type. A value of this type matches one of the forms:

- 0<clear text password>

- 5<salt>$<password hash>

- 5rounds=<number>$<salt>$<password hash>

The '0' prefix signals that this is plain text. When a plain text value is received by the server, a SHA-256 digest is calculated, and the string '5<salt>$' is prepended to the

result, where <salt> is a random 16 character salt used to generate the digest. This value is stored in the configuration data store. The algorithm can be tuned through the /confdConfig/cryptHash/rounds parameter, which if set to a number other than the default will cause '5rounds=<number>$<salt>$' to be prepended instead of only '5<salt>$'.

If a value starting with '5' is received, the server knows that the value already represents a SHA-256 digest, and stores it as is in the data store.

If a default value is specified, it must have a '5' prefix.

The digest algorithm used is the same as the SHA-256 crypt function used for encrypting passwords for various UNIX systems.

Reference:

- IEEE Std 1003.1-2008 - crypt() function FIPS.180-3.2008: Secure Hash Standard

sha-512-digest-string

```
Pattern:
'$0$.*'
'|$6$(rounds=\d+$)?[a-zA-Z0-9./]{1,16}$[a-zA-Z0-9./]{86}'
```

The sha-512-digest-string type automatically computes a SHA-512 digest for a value adhering to this type. A value of this type matches one of the forms

- 0<clear text password>

- 6<salt>$<password hash>

- 6rounds=<number>$<salt>$<password hash>

The '0' prefix signals that this is plain text. When a plain text value is received by the server, a SHA-512 digest is calculated, and the string '6<salt>$' is prepended to the

result, where <salt> is a random 16 character salt used to generate the digest. This value is stored in the configuration data store. The algorithm can be tuned through the

/confdConfig/cryptHash/rounds parameter, which if set to a number other than the default will cause '6rounds=<number>$<salt>$' to be prepended instead of only '6<salt>$'.

If a value starting with '6' is received, the server knows that the value already represents a SHA-512 digest, and stores it as is in the data store.

If a default value is specified, it must have a '6' prefix. The digest algorithm used is the same as the SHA-512 crypt function used for encrypting passwords for various UNIX systems.

Reference:

- IEEE Std 1003.1-2008 - crypt() function FIPS.180-3.2008: Secure Hash Standard

size

Pattern:
`'S(\d+G)?(\d+M)?(\d+K)?(\d+B)?'`

A value that represents a number of bytes. An example could be S1G8M7K956B; meaning 1GB + 8MB + 7KB + 956B = 1082138556 bytes.

The value must start with an S. Any byte magnifier can be left out, for example, S1K1B equals 1025 bytes. The order is significant though, that is S1B56G is not a valid byte size.

In ConfD, a 'size' value is represented as an uint64.

uuid

Pattern:
'[0-9a-fA-F]{8}-[0-9a-fA-F]{4}-[0-9a-fA-F]{4}-'
 '[0-9a-fA-F]{4}-[0-9a-fA-F]{12}'

A Universally Unique IDentifier in the string representation defined in RFC 4122. The canonical representation uses lowercase characters. The following is an example of a UUID in string representation: f81d4fae-7dec-11d0-a765-00a0c91e6bf6.

Reference:

- RFC 4122: A Universally Unique IDentifier (UUID) URN Namespace

yang-identifier

Pattern:
'[a-zA-Z_][a-zA-Z0-9\-_.]*'

Pattern:
'.|..|[^xX].*|.[^mM].*|..[^lL].*'

A YANG identifier string as defined by the 'identifier' rule in Section 12 of RFC 6020. An identifier must start with an alphabetic character or an underscore followed by an arbitrary sequence of alphabetic or numeric characters, underscores, hyphens, or dots. A YANG identifier MUST NOT start with any possible combination of the lowercase or uppercase character sequence 'xml'.

Reference:

- RFC 6020: YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)