

NETCONF Interoperability Testing

Ha Manh Tran, Iyad Tumar, and Jürgen Schönwälder

Computer Science, Jacobs University Bremen, Germany
{h.tran,i.tumar,j.schoenwaelder}@jacobs-university.de

Abstract. The IETF has developed a network configuration management protocol called NETCONF which was published as proposed standard in 2006. The NETCONF protocol provides mechanisms to install, manipulate, and delete the configuration of network devices by using an Extensible Markup Language based data encoding on top of a simple Remote Procedure Call layer. This paper describes a NETCONF interoperability testing plan that is used to test whether NETCONF protocol implementations meet the NETCONF protocol specification. The test of four independent NETCONF implementations reveals bugs in several NETCONF implementations. While constructing test cases, a few shortcomings of the specifications were identified as well.

Key words: Network Management, NETCONF, Interoperability Testing

1 Introduction

The NETCONF protocol specified in RFC 4741 [1] defines a mechanism to configure and manage network devices. It allows clients to retrieve configuration from network devices or to add new configuration to these devices. The NETCONF protocol uses a remote procedure call (RPC) paradigm. A client encodes an RPC request in Extensible Markup Language (XML) [2] and sends it to a server using a secure, connection-oriented session. The server returns with an RPC-REPLY response encoded in XML.

The NETCONF protocol supports features required for configuration management that were lacking in other network management protocols, for instance SNMP [3]. NETCONF operates on so called datastores and represents the configuration of a device as a structured document. The protocol distinguishes between running configurations, startup configurations and candidate configurations. In addition, it provides primitives to assist with the coordination of concurrent configuration change requests and to support distributed configuration change transactions over several devices. Finally, NETCONF provides filtering mechanisms, validation capabilities, and event notification support [4].

The aim of this paper is twofold. First, we describe a NETCONF interoperability testing plan that is used to test whether NETCONF protocol implementations meet the NETCONF protocol specification in RFC 4741. The test plan particularly focuses on testing the correctness of NETCONF messages and

operations; it is not our current goal to measure the performance of NETCONF implementations. Second, we will discuss the observations and results that show how the test plan found some NETCONF implementation bugs, and how it revealed a few shortcomings where the specification (RFC 4741 and RFC 4742 [5]) is either somewhat ambiguous or totally silent.

In order to make the paper concise and precise, we use the word `request` when we refer to an `rpc` request message and the word `response` when we refer to an `rpc-reply` response message. We refer to NETCONF operations such as `get-config` by typesetting the operation name in teletype font. The names of test suites are typeset in small caps, e.g., `VACM`.

The rest of the paper is structured as follows: An overview of the NETCONF protocol is presented in Section 2. Section 3 provides information about the systems under test before the test plan is introduced in Section 4. The NETCONF interoperability tool (NIT) is described in Section 5. Preliminary observations are reported in Section 6 before the paper concludes in Section 7.

2 NETCONF Overview

The NETCONF protocol [1] uses a simple remote procedure call (RPC) layer running over secure transports to facilitate communication between a client and a server. The Secure Shell (SSH) [6] is the mandatory secure transport that all NETCONF clients and servers are required to implement as a means of promoting interoperability [5].

The NETCONF protocol can be partitioned into four layers as shown in Figure 1. The transport protocol layer provides a secure communication path between the client and server. The RPC layer provides a mechanism for encoding RPCs. The operations layer residing on top of the RPC layer defines a set of base operations invoked as RPC methods with XML-encoded parameters to manipulate configuration state. The configuration data itself forms the content layer residing above the operations layer.

The NETCONF protocol supports multiple configuration datastores. A configuration datastore is defined as the set of configuration objects required to get a device from its initial default state into a desired operational state. The `running` datastore is present in the base model and provides the currently active configuration. In addition, NETCONF supports a `candidate` datastore, which is a buffer that can be manipulated and later committed to the `running` datastore, and a `startup` configuration datastore, which is loaded by the device as part of initialization when it reboots or reloads [4].

Table 1 shows the protocol operations that have been defined so far by the NETCONF working group of the IETF. The first two operations `get-config` and `edit-config` can be used to read and manipulate the content of a datastore. The `get-config` operation can be used to read all or parts of a specified configuration. The `edit-config` operation modifies all or part of a specified configuration datastore. Special attributes embedded in the config parameter control which parts of the configuration are created, deleted, replaced or merged.

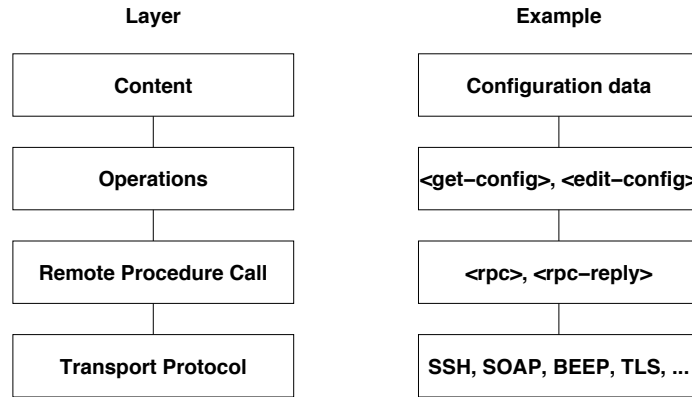


Fig. 1. NETCONF protocol layers [1].

The `test-option` and the `error-option` parameters control the validation and the handling of errors. The `copy-config` operation creates or replaces an entire configuration datastore with the contents of another complete configuration datastore and the `delete-config` operation deletes a configuration datastore (the `running` configuration datastore cannot be deleted).

The `lock` and `unlock` operations do coarse grain locking of a complete datastore and locks are intended to be short lived. More fine grained locking mechanisms are currently being defined in the IETF [4]. The `get` operation can be used to retrieve the running configuration and the current operational state of a device.

Table 1. NETCONF protocol operations (arguments in brackets are optional) [4]

Operation	Arguments
get-config	source [filter]
edit-config	target [default-operation] [test-option] [error-option] config
copy-config	target source
delete-config	target
lock	target
unlock	target
get	[filter]
close-session	
kill-session	session-id
discard-changes	
validate	source
commit	[confirmed confirm-timeout]
create-subscription	[stream] [filter] [start] [stop]

NETCONF sessions can be terminated using either the `close-session` operation or the `kill-session` operation. The `close-session` operation initiates a graceful close of the current session while the `kill-session` operation forces the termination of another session.

The optional `discard-changes` operation clears the candidate configuration datastore by copying the running configuration into the candidate buffer while the optional `validate` operation runs validation checks on a datastore. The optional `commit` operation is used to commit the configuration in the candidate datastore to the running datastore.

A separate specification published as RFC 5277 [7] extends the base NETCONF operations defined in RFC 4741 for notification handling. This is done by adding the `create-subscription` operation and introducing new `notification` messages carrying notification content. By using a notification stream abstraction, it is possible to receive live notifications as well as to replay recorded notifications.

NETCONF protocol introduces the notion of capabilities. A capability is some functionality that supplements the base NETCONF specification. A capability is identified by a uniform resource identifier (URI). The base capabilities are defined using URNs following the method described in RFC 3553 [8]. NETCONF peers exchange device capabilities when the session is initiated: When the NETCONF session is opened, each peer (both client and server) must send a `hello` message containing a list of that peer’s capabilities. This list must include the NETCONF `:base` capability¹. Following RFC 4741, we denote capabilities by the capability name prefixed with a colon, omitting the rest of the URI.

3 Systems Under Test

The systems used for the NETCONF interoperability testing comprise Cisco 1802 integrated services routers, Juniper J6300 routers, the Tail-f ConfD software for configuration management, and the EnSuite software [9] for configuration management. The ConfD software is an extensible development toolkit that allows users to add new components by writing a configuration specification for a data model and loading the generated object and schema files for the components. For the sake of consistency, we refer to the ConfD software as the Tail-f system. The EnSuite software contains a Yencap implementation used to test the NETCONF configuration protocol and extensible features on an experimental network management platform. It also supports web-based configuration management for NETCONF and additional modules and operations for the platform; e.g., the `BGP_Module` for configuring BGP routers and the `Asterisk_Module` for configuring VoIP servers. Table 2 briefly describes the four platforms and the SSH support of the four systems. The ConfD and EnSuite are installed and configured to run on Linux XEN virtual machines [10].

Table 3 presents the NETCONF capabilities announced by the systems under test. The Tail-f system supports all capabilities except the `:startup-ca-`

¹ `urn:ietf:params:netconf:base:1.0`

Table 2. Systems under test

System	Platform	SSH Support
Juniper	JUNOS ver. 9.0	ver. 1.5/2.0
Tail-f	ConfD ver. 2.5.2	ver. 2.0
Cisco	IOS ver. 12.4	ver. 2.0
EnSuite	YencaP ver. 2.1.11	ver. 2.0

pability. The Cisco, Juniper and EnSuite systems support fewer capabilities and apparently the Cisco implementation favours a distinct `startup` datastore while the Juniper implementation favours a `candidate` datastore with commit and rollback support. The EnSuite implementation supports both `startup` and `candidate` datastores. Note that some implementations can be configured to support additional capabilities, but we used the more standard default settings in our tests. In addition to the capabilities listed in Table 3, each system announces several proprietary capabilities.

Table 3. NETCONF capabilities supported by the systems under test

Capability	Juniper	Tail-f	Cisco	EnSuite
<code>:base</code>	✓	✓	✓	✓
<code>:writable-running</code>		✓	✓	✓
<code>:candidate</code>	✓	✓		✓
<code>:confirmed-commit</code>	✓	✓		
<code>:rollback-on-error</code>		✓		
<code>:validate</code>	✓	✓		
<code>:startup</code>			✓	✓
<code>:url</code>	✓	✓	✓	✓
<code>:xpath</code>		✓		✓

The Tail-f and Juniper implementations use an event driven parser. They do not wait for the framing character sequence to respond to a request. The Cisco and EnSuite systems do not seem to have an event driven parser or at least they do not start processing requests until the framing character sequence has been received.

The Juniper implementation is very lenient. For example, it continues processing requests even if the client does not send a `hello` message or the client does not provide suitable XML namespace and message-id attributes. The Juniper implementation supports a large number of vendor-specific operations. In addition, it renders the returned XML content in a tree-structure that is relatively easy to read and it generates XML comments in cases of fatal errors before closing the connection. As a consequence, the Juniper implementation is very easy to use interactively for people who like to learn how things work without using tools other than a scratch pad and a cut and paste device. The EnSuite im-

plementation shares the same characteristics with the Juniper implementation. Moreover, it returns an error message with an explanation of the reason and does not close the connection when the client sends illegal input. It, however, requires message-id attributes for requests.

The Tail-f and Cisco implementations are much less tolerant when processing input not closely following RFC 4741. They also return XML data in a compact encoding, minimizing the embedded white-space and thus reducing message sizes. Without proper tools, it is pretty difficult for humans to read the responses. In some cases, these two implementations close the connection when the client sends illegal input without an indication of the reason for closing the connection. In such cases, it can take some effort to investigate the wrongdoings.

Finally, we like to point out that the Cisco implementation we have tested does not support structured content; i.e., its configuration content is a block of proprietary IOS commands wrapped in an XML element. As a consequence, several of the advanced NETCONF features for retrieving and modifying structured configuration data cannot be applied. The EnSuite implementation still contains bugs and partially supports the `candidate` and `url` capabilities; e.g., several operations on the `candidate` datastore do not seem to work.

4 Test Plan

In this section we describe our NETCONF test plan. To make the execution of the tests efficient and to keep the collection of tests organized, we divided our test plan into five test suites. A test suite is a collection of test cases that are intended to be used to test and verify whether the systems under test meet the NETCONF protocol specification contained in RFC 4741 and RFC 4742.

Table 4 lists the test suites and the current number of test cases in each suite. The total number of test cases is 87. Each test case contains three parts: (i) a pre-configuration prepares the system under test for the test; (ii) a main test sends requests to, and receives responses from, the system under test, and verifies the responses; (iii) a post-configuration brings the system under test to the initial status. Our organization of test cases into test suites is not directly following the vertical layering model show in Figure 1 and the horizontal organization of operations and capabilities in the operations layer as one might expect. The reason is essentially our attempt to reduce the overhead of the pre-configuration and post-configuration parts during the execution of the test suite on the systems under test, e.g., in order to test the `edit-config` operation on a network interface, we describe a sequence of test cases for `create`, `replace`, `merge` and `delete` operations with setting up and cleaning up the interface once. This also led to a more tightly integrated organization of the test cases.

The first test suite is referred as the GENERAL test suite because it includes test cases for individual operations such as `lock`, `unlock`, `close-session`, `kill-session`, `discard-changes`, `validate`, and `commit`. The `lock` and `unlock` test cases verify that the responses do not contain an error or the responses contain a proper error, e.g., a `lock` request to the datastore already locked causes

Table 4. Test suites and current number of test cases

Test Suite	No. Test Cases
GENERAL	19
GET	11
GET-CONFIG	16
EDIT-CONFIG	15
VACM	26

an error. The `kill-session` test case contains a pre-configuration that prepares another running session before terminating it, while the `validate` test case contains a post-configuration that discards a change after validating it. This test suite also checks the format of requests and responses. Few test cases verify whether the responses contain the compulsory attributes and the attribute's value matches the value contained in the requests.

The next two test suites are the GET and GET-CONFIG suites. These suites aim to test the filter mechanism of the `get` and `get-config` operations. While `get` operates on the `running` configuration datastore and the device's state data, `get-config` operates on different sources of the configuration data such as the `running` and `candidate` datastores (depending on the support of capabilities), resulting in additional test cases for the GET-CONFIG suite. Test cases verify several types of subtree filters, e.g., a test case checks whether the system under test returns the entire content of the running configuration data plus the operational state when no filter is used, or another test case checks whether the system under test returns nothing when an empty filter is used.

The EDIT-CONFIG suite involves tests modifying the configuration data in the datastore. This suite includes test cases for the `delete-config`, `copy-config`, and `edit-config` operations. The `edit-config` operation test cases support the `create`, `replace`, `merge` and `delete` operation attributes. Several test cases in this suite are data model specific due to the lack of a common data model, thus we need to implement several tests in different ways. This extra work can be reduced if implementers volunteer to support a common data model.

The last test suite is the VACM suite verifying the NETCONF protocol operations against the VACM data model [11]. This data model is a YANG version of the `SNMP-VIEW-BASED-ACM-MIB` (View-based Access Control Model for the Simple Network Management Protocol [12]). YANG [13] is a data modelling language for NETCONF. Test cases in this suite are generated from this data model focusing on the `group`, `access`, and `view` lists.

5 Test Tool (NIT)

We have implemented a tool called NIT (NETCONF Interoperability Testing tool) to automatically execute the test suites against a system under test. Our NIT tool basically performs the following operations:

1. connecting to a system under test using the SSH

2. verifying the initial `hello` message
3. executing test cases by
 - sending a test request and receiving a response
 - verifying both the request and the response following the criteria defined by RFC 4741.
4. reporting the failure or the success of each test

The tool is equipped with an XML parser to analyze the responses for verification. The parser, upon receiving a response, provides a list of elements with quantity, a list of attributes with quantity, a list of attribute values and a list of text parts. With this information, the tool can detect possible flaws from the responses, such as whether any element or attribute is missing, whether an error must be returned. The following example presents a response without an error or a warning:

```
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply message-id="1007"
          xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <ok/>
</rpc-reply>
]]>]]>
```

The parser provides the following information of the response:

```
-Element Types
  rpc-reply 1
    ok 1
-Attribute Types
  message-id 1
    xmlns 1
-Attribute Values
  message-id ['1007']
    xmlns ['urn:ietf:params:xml:ns:netconf:base:1.0']
-Text Parts
  []
```

We have used the Python unit testing framework [14]. The framework features test automation, shared configuration of setup and shutdown methods, arrangement of tests into collections, and independent reporting of the tests. The tool takes advantage of these features to maintain a single SSH connection for all tests and to group related tests into a collection; e.g., tests concerned with creation, modification and deletion operations are grouped together to reuse and clean the testing environment easily. The tool organizes test cases into several collections of test cases, namely test suites, as discussed in Section 4.

While the tool has been used successfully to test some specific devices (see next section), it possesses several limitations. Firstly, it lacks a resumption mechanism to continue the test run when it encounters connection loss due to the

misbehavior of systems under test. Secondly, while the test cases are believed to comply with RFC 4741, the test scripts, i.e., the piece of code that implements test cases, depends on the specification and configuration of components of the tested systems to produce the requests and to verify the responses. Finally, the framework requires some extra work for complicated test cases; e.g., testing the `lock` operation requires an extra session to lock the database.

6 Preliminary Observations

We have used the NIT tool to test the systems described in Section 3. Since the result of the tests are specific to the different NETCONF implementations, we present the results by referring to system X and we leave out the mapping of X to the systems described in Section 3. Note that we did manually re-check the failed test cases in order to erase bugs in the test scripts. Despite these efforts, several test cases reflect our interpretation of RFC 4741 and there might not be full agreement with our interpretation and thus the numeric results presented below should be taken with a grain of salt.

Table 5. Test result summary organized by the systems under test

System	Success	Failure	Irrelevant
<i>A</i>	47.2%	14.9%	37.9%
<i>B</i>	82.8%	9.2%	8.0%
<i>C</i>	17.3%	10.3%	72.4%
<i>D</i>	17.3%	21.8%	60.9%

Table 5 presents the result of the NIT tool for the systems under test. The “success” and “failure” columns indicate the percentage of passed and failed test cases respectively, while the “irrelevant” column indicates the percentage of test cases that cannot be applied to a specific system due to either system configuration or implementation problems (e.g., the `vacm` data model is not implemented).

We learned that the systems A and B comply reasonably well with the RFCs. The system A fails 14.9% of the test cases and most of them are related to the basic format of request and response messages or the filter mechanism of the `get` operation. The system B performs better with very few failed test cases and most of them are concerned with the validation of XML elements in request messages. The two systems A and B have very few problems with the filter mechanism of the `get-config` operation or the usage of the `edit-config` operation for creating, modifying and deleting configuration elements. The systems C and D perform poorer with 72.4% and 60.9% irrelevant test cases and 10.3% and 21.8% failed test cases, respectively. The failed test cases are related to the format of requests and responses or the filter mechanism of the `get` operation.

Table 6. Test result summary organized by the test suites

Test Suite	Success	Failure	Irrelevant
GENERAL	73.6%	13.2%	13.2%
GET	29.5%	52.3%	18.2%
GET-CONFIG	48.4%	14.1%	37.5%
EDIT-CONFIG	38.3%	1.7%	60%
VACM	19.2%	5.8%	75%

Table 6 reports the passed and failed test cases organized by the test suites over the total number of running test cases for the systems under test. There are two remarks: (i) the GET suite obtains a high percentage of failed test cases 52.3%, and (ii) the EDIT-CONFIG suites obtains low percents of failed test cases 1.7%. We found that the majority of failed test cases from the GET suite is related to the filter mechanism of the `get` operation.

With the failed test cases in mind, we have looked back into the RFCs. There are several things where the RFC is either somewhat ambiguous or totally silent. In general, the RFC should provide more detailed descriptions for error situations and it might be necessary to better constrain the currently open ended format of request and response messages since they for example allow arbitrary values for attributes. Furthermore, the RFC should be updated with clearer examples. Some particular issues are listed below:

- The RFC ignores the XML declaration

```
<?xml version$="1.0" encoding="UTF-8"?>
```

for requests and responses. Some systems do not execute a request without this declaration while other systems do. It seems that the IETF working group favours to have a mandatory XML declaration.

- The examples in RFC 4741 often omit namespace declarations for request and response messages. Only few systems execute a request without a proper namespace declaration and it would help interoperability if the examples would contain namespace declarations where necessary.
- RFC 4741 requires that additional attributes present in the `<rpc>` element of a request message must be returned in the `<rpc-reply>` element of the response message without any change (see section 4.1 of the RFC 4741). This requirement leads to problems when such an attribute conflicts with attributes generated by the implementation. One implementation generated duplicated attributes (and thus invalid XML) while another implementation removes a duplicated attribute resulting in violation of RFC 4741.
- RFC 4741 allows arbitrary strings for the `message-id` attribute. From the tests, we found that implementations terminate the session often without an error indication or return strange results when the `message-id` attribute in a request message contains unexpected content such as the literal string `]]>]]>` or the literal string `</rpc>`. Of course, a proper NETCONF client would not generate such request messages since they are invalid XML. But

on the other hand, one can question whether arbitrary content in request and response attributes is a feature worth to support.

Some of the items listed above are meanwhile actively discussed on the NETCONF working group mailing list and work is underway to revise RFC 4741 in order to fix bugs and to clarify the processing of NETCONF messages [15].

7 Conclusions and Future Work

We have carried out some work on NETCONF interoperability testing. This work aims at observing the compliance of NETCONF implementations with RFC 4741. It also aims at identifying inconsistencies in the RFC. We have proposed a test plan consisting of five test suites. Each test suite contains a number of test cases that involve a single operation or a group of related operations. The test cases exploit several aspects of RFC 4741 including the format of request and response messages, the filter mechanism supported by some operations, NETCONF capabilities, and so on. The test cases have been coded into the NIT tool, which automates the execution of test runs. It should be noted, however, that the test cases so far have not been reviewed and as such there might be disagreement on some test cases whether they are correct or not relative to RFC 4741.

We have used the NIT tool to test four different NETCONF implementations. Our preliminary observations indicate that the number of failed test cases is relatively high for some systems, thus raising the question of the compliance of these systems with RFC 4741. We have also noted some inconsistencies in RFC 4741 that should be addressed in a future revision of this document. It should be mentioned that some test cases are our interpretation of RFC 4741 and it needs to be worked out to what extent our interpretation meets the interpretation of the working group.

While some interesting initial results have been obtained, this work still requires several improvements. First, the coverage of RFC 4741 by the test cases needs to be evaluated and increased by adding additional test cases as needed. Furthermore, it would be nice to reduce the dependency of the test cases on different data models. Third, the NIT tool should be improved to better support more complicated test cases that involve multiple NETCONF sessions. Fourth, it would be nice to have a tool able to generate test suites out of YANG data models. And finally, it would be valuable to repeat the tests with a larger number of different NETCONF implementations and to evaluate how test results impact future software revisions and lead to more interoperability.

Acknowledgment

The work reported in this paper is supported by the EC IST-EMANICS Network of Excellence (#26854).

References

1. R. Enns. NETCONF Configuration Protocol. RFC 4741, December 2006.
2. C. Sperberg-McQueen, J. Paoli, E. Maler, and T. Bray. Extensible Markup Language (XML) 1.0 (Second Edition). <http://www.w3.org/TR/2000/REC-xml-20001006>, October 2000. Last access in July 2008.
3. J. Case, R. Mundy, D. Partain, and B. Stewart. Introduction and Applicability Statements for Internet Standard Management Framework. RFC 3410, December 2002.
4. J. Schönwälder, M. Björklund, and P. Shafer. Network Configuration Management using NETCONF and YANG. *IEEE Communications Magazine*, 2009.
5. M. Wasserman and T. Goddard. Using the NETCONF Configuration Protocol over Secure Shell (SSH). RFC 4742, December 2006.
6. T. Ylonen and C. Lonvick. The Secure Shell (SSH) Protocol Architecture. RFC 4251, January 2006.
7. S. Chisholm and H. Trevino. NETCONF Event Notifications. RFC 5277, July 2008.
8. M. Mealling, L. Masinter, T. Hardie, and G. Klyne. An IETF URN Sub-namespace for Registered Protocol Parameters. RFC 3553, June 2003.
9. V. Cridlig, H. J. Abdelnur, J. Bourdellon, and R. State. A NetConf Network Management Suite: ENSUITE. In *Proc. 5th IEEE International Workshop on IP Operations and Management: Operations and Management in IP-Based Networks (IPOM '05)*, volume 3751 of *LNCS*, pages 152–161. Springer, 2005.
10. P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield. Xen and the Art of Virtualization. In *Proc. 19th ACM Symposium on Operating Systems Principles (SOSP '03)*. ACM, October 2003.
11. J. Schönwälder. VACM Yang Data Model. Jacobs University Bremen, October 2008.
12. B. Wijnen, R. Presuhn, and K. McCloghrie. View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP). RFC 3415, December 2002.
13. M. Björklund. YANG - A data modeling language for NETCONF. Internet draft, January 2009.
14. Python Unit Testing Framework. <http://pyunit.sourceforge.net/>. Last access in November 2008.
15. R. Enns, M. Björklund, and J. Schönwälder. NETCONF Configuration Protocol. Internet Draft <draft-ietf-netconf-4741bis-00.txt>, Juniper Networks, Tail-f Systems, Jacobs University, March 2009.