



syslog-ng Store Box 6.6.0

Performance Whitepaper

Copyright 2021 One Identity LLC.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of One Identity LLC .

The information in this document is provided in connection with One Identity products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of One Identity LLC products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, ONE IDENTITY ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL ONE IDENTITY BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF ONE IDENTITY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. One Identity makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. One Identity does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

One Identity LLC.
Attn: LEGAL Dept
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our Web site (<http://www.OneIdentity.com>) for regional and international office information.

Patents

One Identity is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <http://www.OneIdentity.com/legal/patents.aspx>.

Trademarks

One Identity and the One Identity logo are trademarks and registered trademarks of One Identity LLC. in the U.S.A. and other countries. For a complete list of One Identity trademarks, please visit our website at www.OneIdentity.com/legal. All other trademarks are the property of their respective owners.

Legend

 **WARNING:** A WARNING icon highlights a potential risk of bodily injury or property damage, for which industry-standard safety precautions are advised. This icon is often associated with electrical hazards related to hardware.

 **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

Contents

Preface	4
Log collecting performance	5
Introduction to SSB search algorithms	7
Search performance (SSB T4)	9
Log collection and search	12
Performance of SSB versions	17
Summary	19
About us	20
Contacting us	20
Technical support resources	20

Preface

This white paper enables syslog-ng Store Box (SSB) end-users, integrators, and sales personnel to make predictions about the performance of the SSB appliance based on various environmental and configuration parameters.

The measured data refers to physical SSB appliances.

Log collecting performance

In this simple case we only measured sustained throughput of receiving logs in SSB. This means that there were no other operations like search ongoing at the time of the test. The source of messages was always TCP and we tested with and without encryption enabled on the source. The target logspace was using indexed, compressed log store where all fields were indexed using the default delimiters. We used the default 1024Mb memory limit. There was no pattern database loaded.

Factors with significant effect on performance

On physical hardware, SSB performance is usually limited by available processing power, that is, the performance is CPU bound.

During RAID synchronization the performance can drop significantly, because of the heavy load of the disk subsystem. As a result, a new SSB installation may give misleading performance numbers. Always wait for the RAID synchronization to finish before testing the performance.

If you use DNS for the log source, but the DNS server is unreachable, the performance of SSB will greatly decrease.

Sending logs from one log source into multiple logspaces degrades performance. If possible, send logs from one (or more) log source into a single logspace, and not duplicate the logs. If you need to filter or aggregate log messages in different ways, consider using the filtered logspace and multiple logspace features. (For details, see ["Creating filtered logspaces" in the Administration Guide](#) and ["Creating multiple logspaces" in the Administration Guide](#).)

Parsing syslog headers adds an extra 18% overhead. You can improve the raw performance of SSB by selecting the **Do not parse** option in the log source. However, in this case you cannot search and filter the host, program, pid, and other fields of these messages.

Disabling flow control on a log source will not throttle back clients, and seems to increase performance. However, it may lead to losing messages.

Factors without significant effect on performance

The following factors have no effect, or only limited effect on the performance of SSB.

- Number of plain TCP connections to the log sources of SSB, up to around 5000 connections.
- Number of SSL/TLS TCP connections to the log sources of SSB, up to around 1000 connections.
- Enabling debug logging in SSB has no effect: debug logs are related to tracing web access and related operations.
- The **Trusted, Use DNS, Use FQDN** settings have limited effect on performance, provided that DNS is correctly set up. Internally there is a DNS cache in syslog-ng. (For details on these settings, see [Administration Guide](#).)

Overall performance

Depending on its exact configuration and the mix of log formats received, the largest SSB appliance can collect and index up to 100,000 messages per second (100k EPS) for sustained periods.

Introduction to SSB search algorithms

This section describes how SSB indexes messages and stores data. It gives you insights to understanding and interpreting the search performance of SSB.

Data layout per log store

Level 1:

File system directories organized per year and day of month (YYYY/MM-DD).

Level 2:

Log messages are stored in a single file per day. There is also a file that lists index files (level 3) related to distinct time intervals inside the day.

Level 3:

Index file that holds an ordered list of tokens processed when SSB received the logs. For each token there is a list of unique identifiers that points to the messages that contained the token.

Search algorithm overview

Every SSB search must include the time interval we search in. This information is used to find the days, thus the Level 2 and 3 files that need to be searched. Note SSB stores and searches the log messages based on the time SSB received them (the so-called processed time stamp), and not based on the time stamp included in the log message. The reason is that SSB collects logs in real-time, and also the time stamp in the log messages may not be reliable or complete.

Tokens are the words separated by the delimiters set for the logspace (for details, see ["Configuring the indexer service" in the Administration Guide](#)).

On Level 3, SSB looks up the tokens that match the basic expressions in the search query. Since the tokens are stored in alphabetic order, this lookup is very fast for exact searches. If the token contains wildcards (* or ?), then potential matches are checked individually.

At this point, SSB has the list of message identifiers it needs to calculate AND, OR, NOT expressions and finalize search results per day. Getting the final result simply means repeating the procedure for all the days that are requested in the search interval.

Search performance (SSB T4)

Overview

In this section we describe SSB search algorithm performance, measured from starting a search to returning the first 100 results. When a search is executed, SSB calculates the unique identifiers of every search results, without loading the individual messages. The actual messages are loaded temporarily only when requested on the user interface or the RPC API.

This means that it can easily happen that calculating the results takes under a second, but fetching all the resulting messages takes minutes, because it takes time to read the messages from the disk and return them. This also means that the size of the messages has no impact on the memory usage of search.

We have conducted our tests using a real life logspace containing 200 million log entries (about 9.1Gb compressed). We executed the searches directly on SSB to avoid network and caching effects. The test hardware was an SSB T4 appliance, but the response times are very similar on SSB T10 appliances as well. For SSB T1, response times are higher by a factor of x2.5 on the average.

Search for all

The most basic search expression is the empty (or *) search, which searches for any message stored in a logspace. This search is immediate with little memory usage, since it only involves looking up the unique identifier ranges intersecting with the time range of the search itself.

Simple search expressions

Search for specific token:

Example1: username
Example2: restart

The simplest search expression is a specific token, like login. Tokens are the words separated by the delimiters set for the logspace (for details, see ["Configuring the indexer service" in the Administration Guide](#)).

For 200 million logs, searching for a token takes between 1-5 seconds, and the used memory is roughly the same number of bytes as the number of results.

Wildcards:

```
Example1: user?ame  
Example2: system*  
Example3: *tool
```

You can specify part of a token, or add *, ? characters after and/or in front of the token. For example *pple or appl?.

Search times depend on how many letters are known of the token, especially at the front of the token. The worst case is when the search expression starts with the wildcard, for example *pple, which would take between 30-60 seconds to search for in 200 million messages. Searching for appl* takes around 9 seconds. The absolute worst case is *? where no letter is known, which takes 80 seconds.

Memory consumption is a sum of the number of eventual results plus at most the size of the biggest index file involved in the search. The size of the index file depends on the **Memory limit** setting of the logspace. The higher the limit, the larger the index files. (For details, see ["Configuring the indexer service" in the Administration Guide.](#))

Excluding tokens:

```
Example1: NOT user  
Example2: NOT window*
```

You can exclude tokens from a search using the NOT keyword, as in "NOT apple, NOT *pple, and so on. Such search takes slightly longer than searching for the same expression without NOT. The memory used is roughly the same.

Complex search expressions

Search expressions can be combined with the AND and OR keywords to create complex search expressions. A negation of a complex search expression with NOT keyword remains a complex search expression. Note that SSB automatically optimizes certain search expressions before evaluating them, for example, expression AND * becomes expression, NOT NOT expression becomes expression.

Union of searches (OR):

```
Example1: username OR pass*  
Example2: server1 OR server32 OR server50 OR server3
```

Response time can be calculated by adding up the response times of the searches included in the OR expression. The actual OR operation is extremely efficient, so there is little additional overhead.

Intersection of searches (AND):

Example1: user AND login AND fail
Example2: user AND NOT close*

The maximal response time can be calculated by adding up the response times of included searches in the AND expression. The actual AND operation is extremely efficient, so there is little additional overhead.

Statistics on search results:

Same response times and memory consumption expected as for regular searches.

Log collection and search

Measurement setup:

An indexed logspace was set up with all fields indexed. 25,000,000 log messages were sent via TLS source into the logspace every hour for 24 hours to represent "historical" data. Then we kept sending 25,000,000 log messages every hour, while executing an increasing number of parallel search queries to simulate 1 to 29 concurrent busy users.

User simulation was achieved by executing the same RPC API queries that a front end (i.e. internet browser) would send. These are: splitting the time interval in question into 30 equal parts and calculating the number of results per interval (these are the bars on the search interface) and also fetching the first 200 results (part of this is show at the bottom of the search interface).

Measurement quantities:

- Rate in 1000x message/second, that is, how many log messages we can send into SSB while running searches at the same time.
- Number of searches finished in 3 minutes of the test.

Case 1

The search queries are very simple exact token searches. Memory limit for the logspace is 1024Mb. Search was for X<n> OR Y<m> with variable n,m numbers to avoid effects of the query cache of SSB.

Figure 1: Case 1: T1

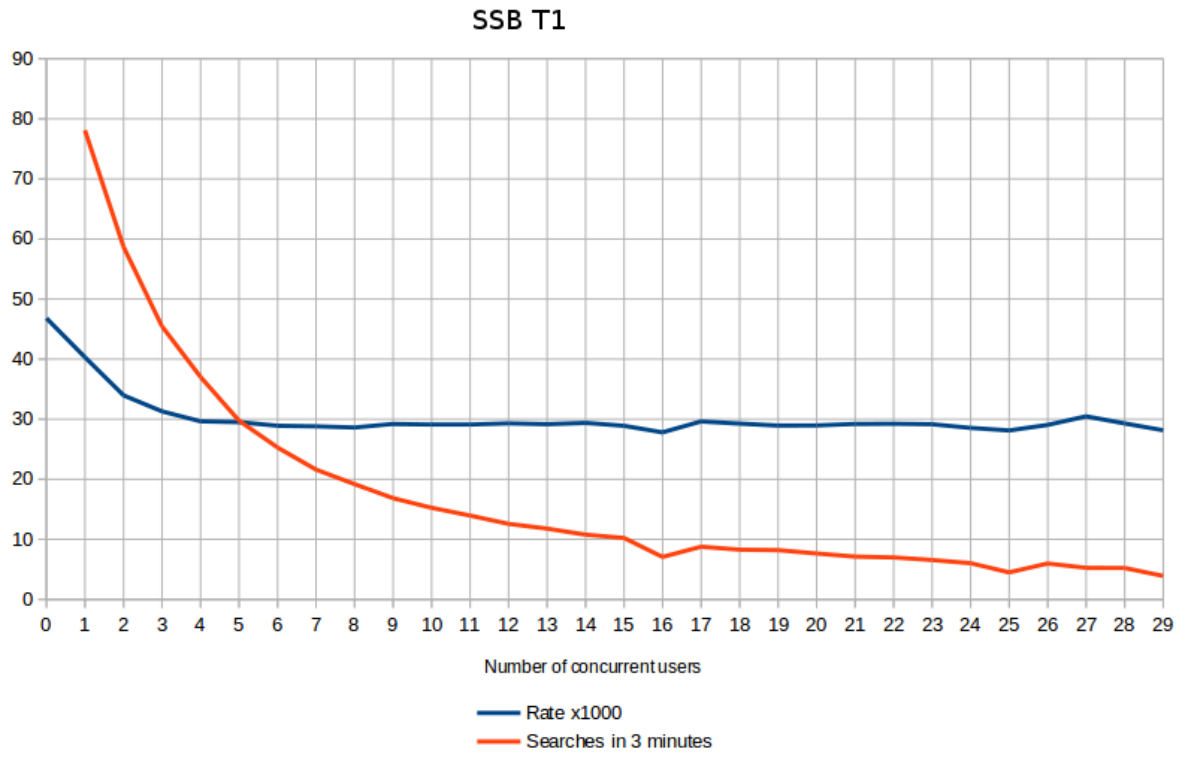


Figure 2: Case 1: T4

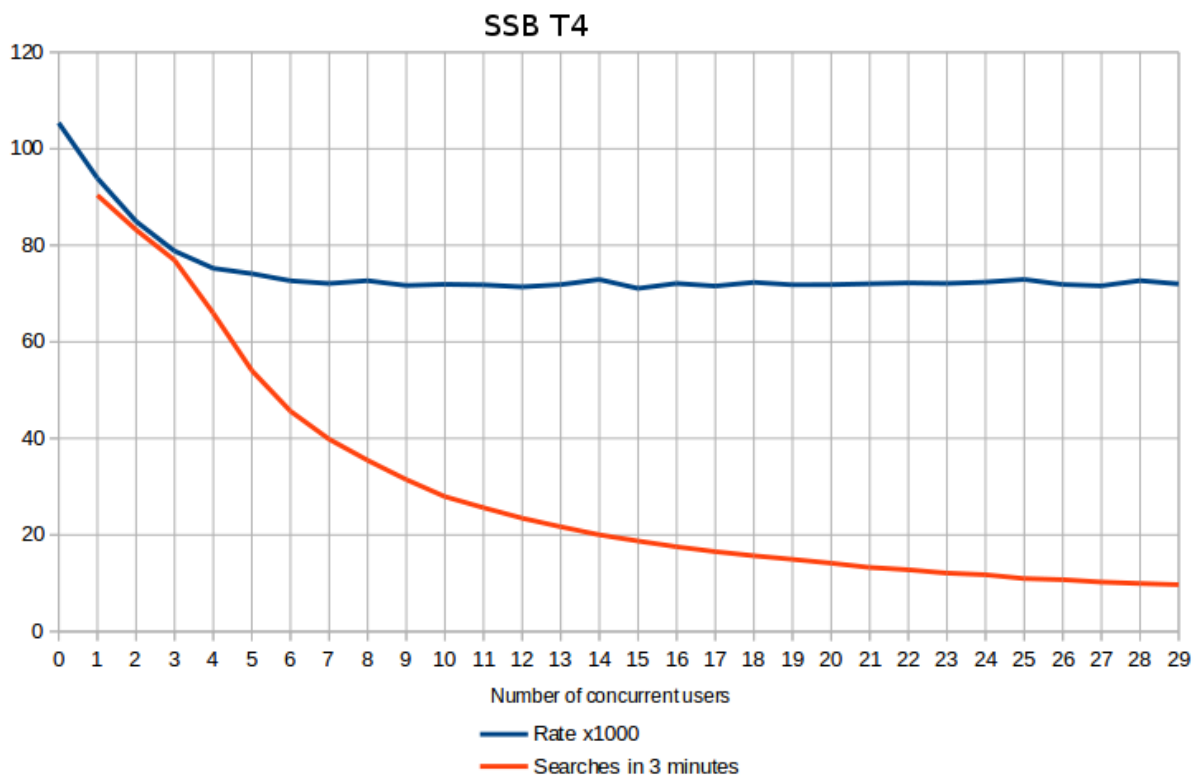
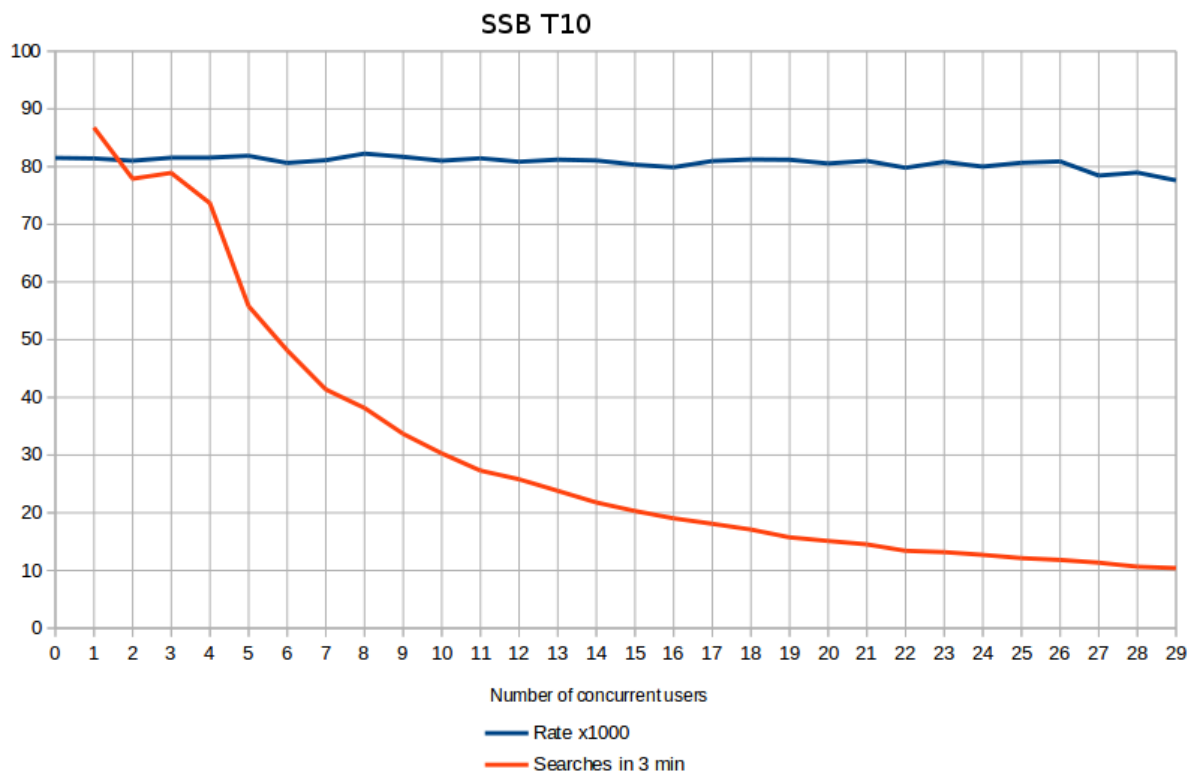


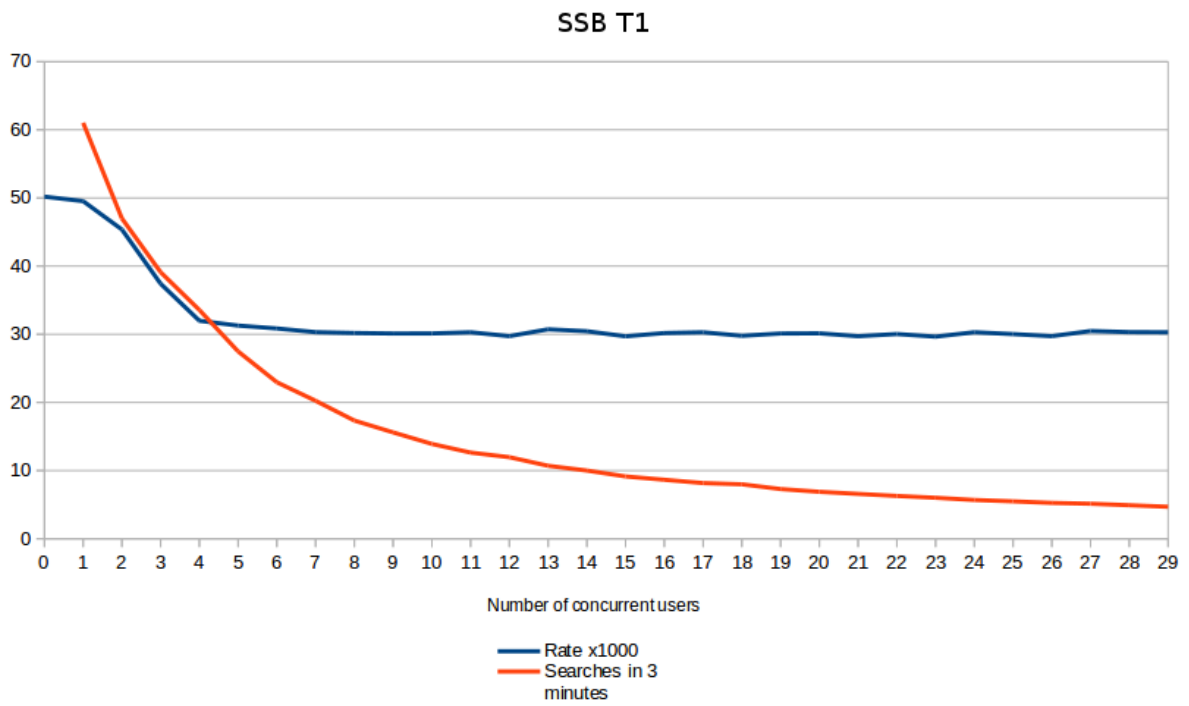
Figure 3: Case 1: T10



Case 2

Addition of *Z to the start of the query to force much more disk IO usage. Results were the same regarding throughput and obviously lower for number of searches.

Figure 4: Case 2: T1



Performance of SSB versions

The following charts show how the search performance of SSB has improved, using SSB 4 LTS (4.0) as a baseline. We used relatively longer queries, where the response time was at least 2 seconds, so that the results are easier to measure and compare. Otherwise, various other factors could have distorted the measurements. Note the following point:

- The structure of index files was optimized in SSB 4 F2, greatly increasing the search performance of SSB. If searching is slow in your SSB and you are using SSB 4 LTS, consider upgrading to a newer version. (Note that this will affect only the index files created after the upgrade.)
- The search algorithms were optimized in SSB 4 F2, decreasing the memory usage of search by an average of 80%. If search causes high memory consumption in your SSB and you are using SSB 4 LTS, consider upgrading to a newer version.

Figure 5: Average response time

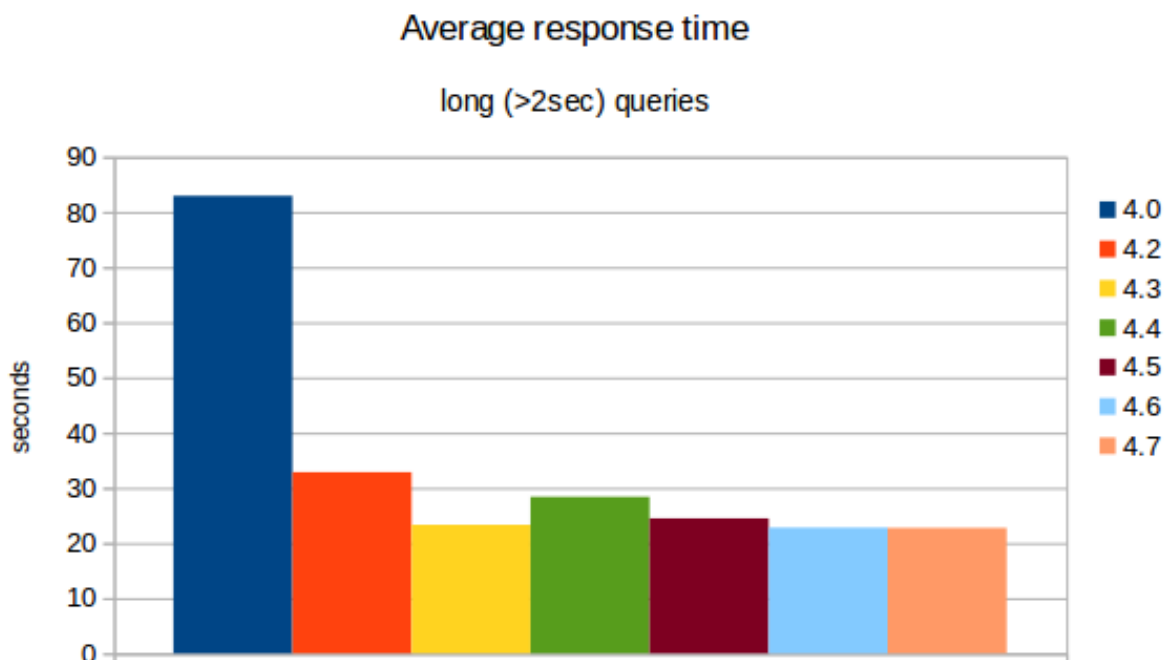
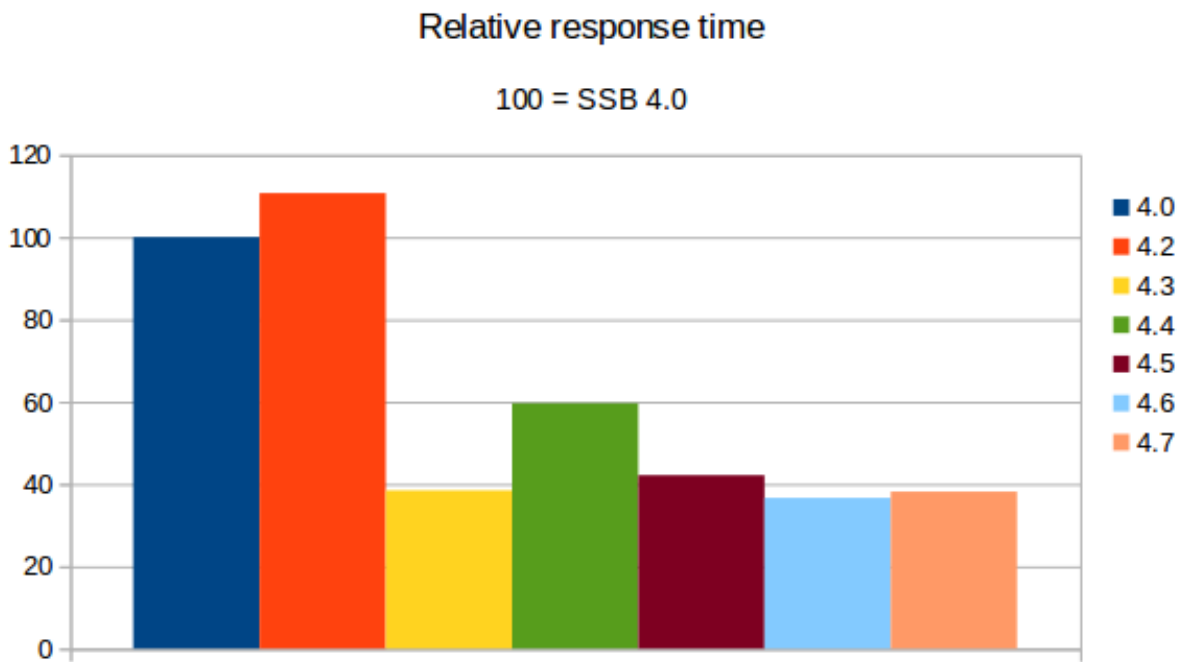


Figure 6: Relative response time



Summary

The test measurements show that the processing capabilities and search performance of syslog-ng Store Box have increased significantly since version 4 LTS, and that SSB is capable of receiving and processing high-volume log traffic. The largest SSB appliance is capable of scaling up to 100,000 event per second (100k EPS).

If the search performance of SSB is not adequate in your environment (search is slow, or greatly increases the memory consumption), check the version of your SSB. If you are using SSB 4 LTS, consider upgrading to a newer version.

If you have questions about the performance of SSB, or need help in optimizing the configuration of your SSB appliance, [contact our Professional Services Team](#).

One Identity solutions eliminate the complexities and time-consuming processes often required to govern identities, manage privileged accounts and control access. Our solutions enhance business agility while addressing your IAM challenges with on-premises, cloud and hybrid environments.

Contacting us

For sales and other inquiries, such as licensing, support, and renewals, visit <https://www.oneidentity.com/company/contact-us.aspx>.

Technical support resources

Technical support is available to One Identity customers with a valid maintenance contract and customers who have trial versions. You can access the Support Portal at <https://support.oneidentity.com/>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to videos at www.YouTube.com/OneIdentity
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product