KSConf Documentation

Release 0.7.9

Lowell Alleman

Sep 23, 2020

Contents

1 Welcome to KSCONF! 3

2 Install 5

3 User Guide 7

3.1 Introduction . . . 7
3.2 Concepts . . . 8
3.3 Installation Guide . . . 10
3.4 Commands . . . 15
3.5 Cheat Sheet . . . 40
3.6 Contributing . . . 44
3.7 Developer setup . . . 46
3.8 Git tips & tricks . . . 47
3.9 Random . . . 51
3.10 Contact . . . 53
3.11 Command line reference . . . 53
3.12 Changelog . . . 65
3.13 Known issues . . . 75
3.14 Advanced Installation Guide . . . 76
3.15 License . . . 86
3.16 API Reference . . . 89

4 Indices and tables 101

Bibliography 103

Python Module Index 105

Index 107



Author Lowell Alleman (Kintyre)

Version 0.7

CHAPTER 1

Welcome to KSCONF!

KSCONF is a modular command line tool for Splunk admins and app developers. It's quick and easy to get started with basic commands and grow into the more advanced commands as needed. Thank you for reviewing our expanding body of documentation to help smooth your transition to a more well-managed Splunk environment and explore ways to integrate Ksconf capabilities into your existing workflow.

We are glad you are here! No matter where you're starting from, Ksconf can help. Let us know if there is anything we can do to help along your journey.

– Kintyre team

CHAPTER 2

Install

Ksconf can be directly installed as a Python (via `pip`) or as a Splunk app. The Splunk app option is often easier.

To install as a **python package**, run the following:

```
pip install kintyre-splunk-conf
```

To install the **Splunk app**, download the latest KSCONF App for Splunk release. Note that a one-time registration command is needed to make `ksconf` executable:

```
splunk cmd python $SPLUNK_HOME/etc/apps/ksconf/bin/install.py
```

CHAPTER 3

User Guide

3.1 Introduction

KSCONF (Kintyre's Splunk Configuration tool) is a command-line tool that helps administrators and developers manage their Splunk environments by enhancing their ability to control configuration files. By design, the interface is modular so that each function (aka subcommand) can be learned quickly and used independently. Most Ksconf commands are simple enough for a quick one-off job, yet reliable enough to integrate into complex app build and deployment workflow.

Ksconf helps manage the nuances of storing Splunk apps in a version control system, such as git. It also supports pointing live Splunk apps to a working tree, merging changes from the live system's (local) folder to the version controlled folder (often 'default'), and in more complex cases, it deals with more than one *layer* of "default", which Splunk can't handle natively.

Note: What KSCONF is not

Ksconf does *not* replace your existing Splunk deployment mechanisms or version control tools. The goal is to complement and extend, not replace, the workflow that works for you.

3.1.1 Design principles

Ksconf is a toolbox. Each tool has a specific purpose and function that works independently. Borrowing from the Unix philosophy, each command should do one thing well and be easily combined to handle higher-order tasks.

When possible, be familiar. Various commands borrow from popular UNIX command line tools such as **`grep`** and **`diff`**. The modular nature of the command and other design features were borrowed from **`git`** and **`splunk`** as well.

Don't impose workflow. Ksconf works with or without version control and independently of your deployment mechanisms. If you are looking to implement these things, Ksconf is a great building block.

Embrace automated testing. It's impractical to check every scenario between each release, but significant work has gone into unit testing the CLI to avoid breakage.

3.1.2 Common uses for Ksconf

- Promote changes from `local` to `default`
- Maintain multiple independent layers of configurations
- Reduce duplicate settings in a local file
- Upgrade apps stored in version control
- Merge or separate configuration files
- Git pre-commit hook for validation
- Git post-checkout hook for workflow automation
- Send *.conf* stanzas to a REST endpoint (Splunk Cloud or no file system access)

3.1.3 Getting started

You're in the right place. If you are a beginner, try checking these out first:

- *Cheat Sheet* - Like jumping in the deep end, or prefer examples of descriptions? Start here.
- *Concepts* - To get a more theoretical background on why these things matter.
- *Commands* - Start here if you'd like a more thorough introduction.

3.2 Concepts

3.2.1 Configuration layers

The idea of configuration layers is shared across multiple actions in Ksconf. Specifically, *combine* is used to merge multiple layers, and the *unarchive* command can be used to install or upgrade an app in a layer-aware way.

What's the problem?

In a typical enterprise deployment of Splunk, a single app can easily have multiple logical sources of configuration:

1. Upstream app releases, often from Splunkbase
2. Organization-specific customizations or fixes added by a local developer
3. Fixes to buggy upstream settings, like `indexes.conf`, requested by your Splunk admin

4. Custom knowledge objects created by subject matter experts

Ideally we would like to version control these, but doing so is complicated because normally you have to manage all four of these logical layers in one 'default' folder.

Note: Isn't that what the **local** folder is for?

Splunk requires that app settings be located either in `default` or `local`; and managing local files with version control leads to merge conflicts. So effectively, all version controlled settings need to be in `default`, or risk merge conflicts. However, making changes to the `default` folder causes issues when you attempt to upgrade an app upstream. See how this is a dilemma?

Let's suppose a new upstream version is released. If you aren't managing layers independently, then you have to manually upgrade the app, being careful to preserve all custom configurations. Compare this to the solution provided by the *combine* functionality. The layered approach provides an advantage because logical sources can be stored separately in their own directories, thus allowing them to be modified independently. Using this approach, changes in the "upstream" layer will only come from an official release, and the organizational layer will contain customizations made solely by your organization. Practically, this means it's no longer necessary to comb through commit logs identifying which custom changes need to be preserved and reapplied.

While this doesn't completely remove the need for a person to review app upgrades, it does lower the overhead enough that updates can be pulled in more frequently, thus minimizing divergence.

3.2.2 Minimizing files

A typical scenario:

To customize a Splunk app or add-on, many admins simply copy the conf file from default to local and then apply changes to the local copy. That's a common practice, but stopping there complicates future upgrades. The next step should be to clean up the local file, deleting all the unmodified entries that were copied from default.

Why does this matter?

If you've copied a default file into the local folder, this means that local file doesn't contain *only* your settings, it contains a copy of *all* of the default settings too. So in the future, fixes published by the app creator are likely to be masked by your local settings. A better approach is to reduce the local conf file leaving only the stanzas and settings that you intended to change. While this is a monotonous to do by hand, it is easily accomplished by *ksconf minimize*. This makes your conf files easier to read and simplifies upgrades.

What does Splunk have to say about this? (From the docs)

> "When you first create this new version of the file, **start with an empty file and add only the attributes that you need to change.** Do not start from a copy of the default directory. If you copy the entire default file to a location with higher precedence, any changes to the default values that occur through future Splunk Enterprise upgrades cannot take effect, because the values in the copied file will override the updated values in the default file." – [SPLKDOC1]

Tip:

It's a good practice to minimize your files right away. If you wait, it may not be obvious what specific version of default that local was copied from. In other words, if you run the **minimize** command *after* you've upgraded the default folder, you may need to do extra work to manually reconcile upgrade differences, because any changes made between the initial version of the default file and the most recent release of the conf file cannot, be automatically addressed in this fashion. If your files are all in git, and you know a REF of the previous version of your default file, you can use some commands like this:

```
# Review the output of the log, and find the revision of the last change
git log --oneline -- default/inputs.conf

# Assuming "e633e6" was identified as the desired baseline ref, based on the 'log'␣
↪output

# Compare what's changed in the 'inputs.conf' file between releases (FYI only)
ksconf diff <(git show e633e6:./default/inputs.conf) default/inputs.conf

# Now apply the 'minimization' based on the original version of inputs.conf
ksconf minimize --target=local/inputs.conf <(git show e633e6:./default/inputs.conf)
```

As always, be sure to double check the results.

3.3 Installation Guide

KSCONF can be installed either as a Splunk app or a Python package. Picking the option that's right for you is fairly easy.

Unless you have experience with Python packaging or are planning on customizing or extending Ksconf, then the *Splunk app* is likely the best place for you to start. The native Python package works well for many developer-centric scenarios, but installation ends up being complicated for the more typical admin-centric use-case. Therefore, most users will find it easier to start with the Splunk app.

Note: The introduction of a Splunk app is a fairly new occurence (as of the 0.6.x release). Originally we resisted this idea, since `ksconf` was designed to manage other apps, not live within one. Ultimately however, the packaging decision was made to ensure users of all levels can utilize the program, as Python packaging is a mess and can be daunting for the uninitiated.

3.3.1 Overview

Install	Advantages	Potential pitfalls
Python package	• Most 'pure' and flexible install • One command install. (ideal) • Easy upgrades • More extendable (plugins) • *Install Python package*	• Lots of potential variations and pitfalls • Many Linux distro's don't ship with `pip` • Must consider/coordinate installation user. • Often requires some admin access. • Too many install options (complexity)
Splunk app	• Quick installation (single download) • Requires one time bootstrap command • Self contained; no admin access require • Fast demo; fight with `pip` later • *Install Splunk App*	• Crippled Python install (no `pip`) • Can't add custom extensions (entrypoints) • No CLI completion (yet) • *Grandfather Paradox*
Offline package	• Security: strict review and change control • *Advanced Installation Guide*.	• Requires many steps. • Inherits 'Python package' pitfalls.

3.3.2 Requirements

Python package install:

- Python Supports Python 2.7, 3.4+
- PIP (strongly recommended)
- Tested on Mac, Linux, and Windows

Splunk app install:

- Splunk 6.0 or greater is installed

3.3.3 Install Splunk App

Download and install the KSCONF App for Splunk. Then open a shell, switch to the Splunk user account and run this one-time bootstrap command.

```
splunk cmd python $SPLUNK_HOME/etc/apps/ksconf/bin/install.py
```

On Windows, open a terminal as Administrator and type:

```
cd "C:\Program Files\Splunk"
bin\splunk.exe cmd python etc\apps\ksconf\bin\install.py
```

This will add `ksconf` to Splunk's `bin` folder, thus making it executable either as `ksconf` or, less optimally, `splunk cmd ksconf`. (If you can run `splunk` without giving it a path, then `ksconf` should work too.)

At some point we may add an option for you to do this setup step from the UI.

Note: Alternate download

You can also download the latest (and pre-release) SPL from the GitHub Releases page. Download the file named like `ksconf-app_for_splunk-ver.tgz`

3.3.4 Install Python package

Quick Install

Using pip:

```
pip install kintyre-splunk-conf
```

System-level install: (For Mac/Linux)

```
curl https://bootstrap.pypa.io/get-pip.py | sudo python - kintyre-splunk-conf
```

Enable Bash Completion

Context-aware autocomplete can be a great time saver. If you're on a Mac or Linux, and would like to enable bash completion, run these commands:

```
pip install argcomplete
echo 'eval "$(register-python-argcomplete ksconf)"' >> ~/.bashrc
```

(This option is not currently available for Splunk App installs due to a lack of documentation and testing available presently. It should be possible. Pull requests are welcome.)

Ran into issues?

If you encounter any issues, please refer to the *Advanced Installation Guide*. Substantial time and effort was placed into the assembly of the information based on various scenarios we encountered. A good place to begin would be in the *Troubleshooting* section.

3.3.5 Install from GIT

If you'd like to contribute to ksconf, or just build the latest and greatest, then installing from the git repository is a good choice. (Technically this is still installing with `pip`, so it's easy to switch between a PyPI install, and a local install.)

```
git clone https://github.com/Kintyre/ksconf.git
cd ksconf
pip install .
```

See *Developer setup* for additional details about contributing to ksconf.

3.3.6 Validate the install

No matter how you install `ksconf`, you can confirm that it's working with the following command:

```
ksconf --version
```

The output should look something like this:

```
                                #
                                ##
 ###  ##     #### ###### #######  ###  ##  #######
 ### ##     ###  ###           ## #### ##
 #####      ###  ###      ##   ## #######  #######
 ### ##     ###  ###      ##   ## ### ###  ##
 ###  ## #####    ######   #####  ###  ##  ##
                                     #

ksconf 0.7.3  (Build 376)
Python: 2.7.15  (/Applications/splunk/bin/python)
Git SHA1 dc94f811 committed on 2019-06-05
Installed at: /Applications/splunk/etc/apps/ksconf/bin/lib/ksconf
Written by Lowell Alleman <lowell@kintyre.co>.
Copyright (c) 2019 Kintyre Solutions, Inc, all rights reserved.
Licensed under Apache Public License v2

  kintyre_splunk_conf  (0.7.3)

    Commands:
      check           (stable)    OK
      combine         (beta)      OK
      diff            (stable)    OK
      filter          (alpha)     OK
      merge           (stable)    OK
      minimize        (beta)      OK
      promote         (beta)      OK
      rest-export     (beta)      OK
      rest-publish    (alpha)     OK   (splunk-sdk 1.6.6)
      snapshot        (alpha)     OK
      sort            (stable)    OK
```

(continues on next page)

(continued from previous page)

```
    unarchive       (beta)      OK
    xml-format      (alpha)     OK   (lxml 4.2.5)
```

Missing 3rd party libraries

Note: *Splunk app for KSCONF* users don't need to worry about this.

As of version 0.7.0, ksconf now includes commands that require external libraries. But to keep the main package slim, these libraries aren't strictly required unless you want the specific commands. As part of this change, **`ksconf --version`** now reports any issues with individual commands in the 3rd column. Any value other than 'OK' indicates a problem. Here's an example of the output if you're missing the `splunk-sdk` package.

```
...
promote         (beta)      OK
rest-export     (beta)      OK
rest-publish    (alpha)     Missing 3rd party module:  No module named splunklib.client
snapshot        (alpha)     OK
...
```

Note that while the `rest-publish` command will not work in the example above, all of the other commands will continue to work fine. If you don't need `rest-publish` then there's no need to do anything about it. If you want the packages, install the "thirdparty" extras using the following command:

```
pip install kintyre-splunk-conf[thirdparty]
```

Other issues

If you run into any issues, check out the *Validate the install* section.

3.3.7 Command line completion

Bash completion allows for a more intuitive and interactive workflow by providing quick access to command line options and file completions. Often this saves time since the user can avoid mistyping file names or be reminded of which command line actions and arguments are available without switching contexts. For example, if the user types `ksconf d` and hits `Tab`, then the `ksconf diff` is completed. Or if the user types `ksconf`, and hits `Tab` twice, the full list of command actions are listed.

This feature uses the argcomplete Python package and supports Bash, zsh, tcsh.

Install via pip:

```
pip install argcomplete
```

Enabling command line completion for ksconf can be done in two ways. The easiest option is to enable it for ksconf only. (However, it only works for the current user; it can break if the ksconf command is referenced in a non-standard way.) The alternate option is to enable global command line completion for all python scripts at once, which is preferable if you use *argparse* for many python tools.

Enable argcomplete for ksconf only:

```
# Edit your bashrc script
vim ~.bashrc

# Add the following line
eval "$(register-python-argcomplete ksconf)"

# Restart you shell, or just reload by running
source ~/.bashrc
```

To enable argcomplete globally, run the command:

```
activate-global-python-argcomplete
```

This adds a new script to your the `bash_completion.d` folder, which can be used for all scripts and all users, but it does add some minor overhead to each completion command request.

OS-specific notes:

- **Mac OS X**: The global registration option may not work as the old version of Bash was shipped by default. So either use the one-shot registration, or install a later version of bash with homebrew: `brew install bash` then. Switch to the newer bash by default with `chsh /usr/local/bin/bash`.
- **Windows**: Argcomplete doesn't work on windows Bash for GIT. See argcomplete issue 142 for more info. If you really want this, use Linux subsystem for Windows instead.

3.4 Commands

The ksconf command documentation is provided in the following ways:

1. A detailed listing of each sub-command is provided in this section. This includes relevant background descriptions, typical use cases, examples, and discussion of relevant topics. An expanded descriptions of CLI arguments and their usage is provided here. If you have not used a particular command before, start here.
2. The *Command line reference* provides a quick and convenient reference when the command line is unavailable. The same information is available by typing `ksconf <CMD> --help`. This is most helpful if you're already familiar with a command, but need a quick refresher.

Warning: Apologies for the dust

The command docs are currently undergoing reorganization. We're considering a topical layout rather than a per-command layout. Feedback and technical writing / organization contributions are highly welcomed.

Table 1: Command Listing

Command	Maturity	Description
ksconf check	stable	Perform basic syntax and sanity checks on .conf files
ksconf combine	beta	Combine configuration files across multiple source directories into a single destination directory. This allows for an arbitrary number of Splunk configuration layers to coexist within a single app. Useful in both ongoing merge and one-time ad-hoc use.
ksconf diff	stable	Compare settings differences between two .conf files ignoring spacing and sort order
ksconf filter	alpha	A stanza-aware GREP tool for conf files
ksconf merge	stable	Merge two or more .conf files
ksconf minimize	beta	Minimize the target file by removing entries duplicated in the default conf(s)
ksconf promote	beta	Promote .conf settings between layers using either batch or interactive mode. Frequently this is used to promote conf changes made via the UI (stored in the `local` folder) to a version-controlled directory, such as `default`.
ksconf rest-export	beta	Export .conf settings as a curl script to apply to a Splunk instance later (via REST)
ksconf rest-publish	alpha	Publish .conf settings to a live Splunk instance via REST
ksconf snapshot	alpha	Snapshot .conf file directories into a JSON dump format
ksconf sort	stable	Sort a Splunk .conf file creating a normalized format appropriate for version control
ksconf unarchive	beta	Install or upgrade an existing app in a git-friendly and safe way
ksconf xml-format	alpha	Normalize XML view and nav files

3.4.1 ksconf

Ksconf: Kintyre Splunk CONFig tool

This utility handles a number of common Splunk app maintenance tasks in a small and easy to deploy package. Specifically, this tool deals with many of the nuances with storing Splunk apps in git and pointing live Splunk apps to a git repository. Merging changes from the live system's (local) folder to the version controlled (default) folder and dealing with more than one layer of "default" are all supported tasks which are not native to Splunk.

```
usage: ksconf [-h] [--version] [--force-color]
              {check,combine,diff,filter,merge,minimize,promote,rest-export,rest-publish,
↪snapshot,sort,unarchive,xml-format}
```

(continues on next page)

(continued from previous page)

```
                    ...
```

Named Arguments

--version show program's version number and exit

--force-color Force TTY color mode on. Useful if piping the output a color-aware pager, like 'less -R'

3.4.2 ksconf check

Provides basic syntax and sanity checking for Splunk's .conf files. Use Splunk's built-in `btool check` for a more robust validation of attributes and values.

Consider using this utility as part of a pre-commit hook.

```
usage: ksconf check [-h] [--quiet] FILE [FILE ...]
```

Positional Arguments

FILE One or more configuration files to check. If '-' is given, then read a list of files to validate from standard input

Named Arguments

--quiet, -q Reduce the volume of output.

See also:

Pre-commit hooks

See *Pre-commit hooks* for more information about how the `check` command can be easily integrated in your git workflow.

How 'check' differs from btool's validation

Keep in mind that idea of *valid* in ksconf is different than within Splunk. Specifically,

- **Ksconf is more picky syntactically.** Dangling stanzas and junk lines are picked up by ksconf in general (the 'check' command or others), but silently by ignored Splunk.
- **Btool handles content validation.** The **`btool check`** mode does a great job of checking stanza names, attribute names, and values. Btool does this well and ksconf tries to not repeat things that Splunk already does well.

Why is this important?

Can you spot the error in this `props.conf`?

```
1 [myapp:web:access]
2 TIME_PREFIX = \[
3 SHOULD_LINEMERGE = false
4 category = Web
5 REPORT-access = access-extractions
6
7 [myapp:total:junk
8 TRANSFORMS-drop = drop-all
```

That's right, line 7 contains the stanza `myapp:total:junk` that doesn't have a closing `]`. How does Splunk handle this? It ignores the broken stanza header completely and therefore `TRANSFORMS-drop` gets added to the `myapp:web:access` sourcetype, which will likely result in the loss of data.

Splunk also ignores entries like this:

```
EVAL-bytes-(coalesce(bytes_in,0)+coalesce(bytes_out,0))
```

Of course here there's no `=` anywhere on the line, so Splunk just assumes it's junk and silently ignores it.

Tip: If you want to see how different this is, run ksconf check against the system default files:

```
ksconf check --quiet $SPLUNK_HOME/etc/system/default/*.conf
```

There's several files that ship with the core product that don't pass this level of validation.

Note: Key concepts

Before diving into the `combine` command, it may be helpful to brush up on the concept of *configuration layers*.

3.4.3 ksconf combine

Merge .conf settings from multiple source directories into a combined target directory. Configuration files can be stored in a `/etc/*.d` like directory structure and consolidated back into a single 'default' directory.

This command supports both one-time operations and recurring merge jobs. For example, this command can be used to combine all users' knowledge objects (stored in 'etc/users') after a server migration, or to merge a single user's settings after their account has been renamed. Recurring operations assume some type of external scheduler is being used. A best-effort is made to only write to target files as needed.

The 'combine' command takes your logical layers of configs (upstream, corporate, Splunk admin fixes, and power user knowledge objects, ...) expressed as individual folders and merges them all back into the single `default` folder that Splunk reads from. One way to keep the 'default' folder up-to-date is using client-side git hooks.

No directory layout is mandatory, but one simple approach is to model your layers using a prioritized 'default.d' directory structure. This idea is borrowed from the Unix System V concept where many services natively read their config files from `/etc/*.d` directories.

```
usage: ksconf combine [-h] [--target TARGET] [--dry-run] [--follow-symlink]
                      [--banner BANNER] [--disable-marker]
                      source [source ...]
```

Positional Arguments

source The source directory where configuration files will be merged from. When multiple source directories are provided, start with the most general and end with the specific; later sources will override values from the earlier ones. Supports wildcards so a typical Unix `conf.d/##-NAME` directory structure works well.

Named Arguments

--target, -t Directory where the merged files will be stored. Typically either 'default' or 'local'

--dry-run, -D Enable dry-run mode. Instead of writing to TARGET, preview changes as a 'diff'. If TARGET doesn't exist, then show the merged file.

--follow-symlink, -l Follow symbolic links pointing to directories. Symlinks to files are followed.

--banner, -b A banner or warning comment added to the top of the TARGET file. Used to discourage Splunk admins from editing an auto-generated file.

For other on-going *combine* operations, it's helpful to inform any .conf file readers or potential editors that the file is automatically generated and therefore could be overwritten again. For one-time *combine* operations, the default banner can be suppressed by passing in an empty string (`''` or `""` on Windows)

--disable-marker Prevents the creation of or checking for the '.ksconf_controlled' marker file safety check. This file is typically used indicate that the destination folder is managed by ksconf. This option should be reserved for well-controlled batch processing scenarios.

You may have noticed similarities between the `combine` and *merge* subcommands. That's because under the covers they are using much of the same code. The combine operation essentially does

a recursive merge between a set of directories. One big difference is that `combine` command will handle non-conf files intelligently, not just conf files.

Note: Mixing layers

Just like all layers can be managed independently, they can also be combined in any way you would like. While this workflow is outside of the scope of the examples provided here, it's certainly feasible. This also allows for different layers to be mixed-and-matched by selectively including which layers to combine.

Examples

Merging a multilayer app

Let's assume you have a directory structure that looks like the following. This example features the Cisco Security Suite.

```
Splunk_CiscoSecuritySuite/
├── README
├── default.d
│   ├── 10-upstream
│   │   ├── app.conf
│   │   ├── data
│   │   │   └── ui
│   │   │       ├── nav
│   │   │       │   └── default.xml
│   │   │       └── views
│   │   │           ├── authentication_metrics.xml
│   │   │           ├── cisco_security_overview.xml
│   │   │           ├── getting_started.xml
│   │   │           ├── search_ip_profile.xml
│   │   │           ├── upgrading.xml
│   │   │           └── user_tracking.xml
│   │   ├── eventtypes.conf
│   │   ├── macros.conf
│   │   ├── savedsearches.conf
│   │   └── transforms.conf
│   ├── 20-my-org
│   │   └── savedsearches.conf
│   ├── 50-splunk-admin
│   │   ├── indexes.conf
│   │   ├── macros.conf
│   │   └── transforms.conf
│   └── 70-firewall-admins
│       ├── data
│       │   └── ui
│       │       └── views
│       │           ├── attacks_noc_bigscreen.xml
│       │           ├── device_health.xml
```

(continues on next page)

(continued from previous page)

```
│       │           └── user_tracking.xml
│       └── eventtypes.conf
```

In this structure, you can see several layers of configurations at play:

1. The `10-upstream` layer appears to be the version of the default folder that shipped with the Cisco app.
2. The `20-my-org` layer is small and only contains tweaks to a few saved search entires.
3. The `50-splunk-admin` layer represents local settings changes to specify index configurations, and to augment the macros and transformations that ship with the default app.
4. And finally, `70-firewall-admins` contains some additional view (2 new, and 1 existing). Note that since `user_tracking.xml` is not a `.conf` file it will fully replace the upstream default version (that is, the file in `10-upstream`)

Here are the commands that could be used to generate a new (merged) `default` folder from all of the layers shown above.

```
cd Splunk_CiscoSecuritySuite
ksconf combine default.d/* --target=default
```

See also:

The *unarchive* command can be used to install or upgrade apps stored in a version controlled system in a layer-aware manor.

Consolidating 'users' directories

The `combine` command can consolidate 'users' directory across several instances after a phased server migration. See *Migrating the 'users' folder*.

3.4.4 ksconf diff

Compares the content differences of two .conf files

This command ignores textual differences (like order, spacing, and comments) and focuses strictly on comparing stanzas, keys, and values. Note that spaces within any given value, will be compared. Multi-line fields are compared in a more traditional 'diff' output so that long saved searches and macros can be compared more easily.

```
usage: ksconf diff [-h] [-o FILE] [--comments] CONF1 CONF2
```

Positional Arguments

CONF1	Left side of the comparison
CONF2	Right side of the comparison

Named Arguments

-o, --output File where difference is stored. Defaults to standard out.

--comments, -C Enable comparison of comments. (Unlikely to work consistently)

Example

Add screenshot here

To use `ksconf diff` as an external diff tool, check out *Ksconf as external difftool*.

3.4.5 ksconf filter

Filter the contents of a conf file in various ways. Stanzas can be included or excluded based on a provided filter or based on the presence or value of a key.

Where possible, this command supports GREP-like arguments to bring a familiar feel.

```
usage: ksconf filter [-h] [-o FILE] [--comments] [--verbose]
                     [--match {regex,wildcard,string}] [--ignore-case]
                     [--invert-match] [--files-with-matches]
                     [--count | --brief] [--stanza PATTERN]
                     [--attr-present ATTR] [--keep-attrs WC-ATTR]
                     [--reject-attrs WC-ATTR]
                     CONF [CONF ...]
```

Positional Arguments

CONF Input conf file

Named Arguments

-o, --output File where the filtered results are written. Defaults to standard out.

--comments, -C Preserve comments. Comments are discarded by default.

--verbose Enable additional output.

--match, -m Possible choices: regex, wildcard, string

Specify pattern matching mode. Defaults to 'wildcard' allowing for * and ? matching. Use 'regex' for more power but watch out for shell escaping. Use 'string' to enable literal matching.

--ignore-case, -i Ignore case when comparing or matching strings. By default matches are case-sensitive.

--invert-match, -v Invert match results. This can be used to show what content does NOT match, or make a backup copy of excluded content.

Output mode

Select an alternate output mode. If any of the following options are used, the stanza output is not shown.

--files-with-matches, -l List files that match the given search criteria

--count, -c Count matching stanzas

--brief, -b List name of matching stanzas

Stanza selection

Include or exclude entire stanzas using these filter options.

All filter options can be provided multiple times. If you have a long list of filters, they can be saved in a file and referenced using the special `file://` prefix. One entry per line.

--stanza Match any stanza who's name matches the given pattern. PATTERN supports bulk patterns via the `file://` prefix.

--attr-present Match any stanza that includes the ATTR attribute. ATTR supports bulk attribute patterns via the `file://` prefix.

Attribute selection

Include or exclude attributes passed through. By default, all attributes are preserved. Allowlist (keep) operations are preformed before blocklist (reject) operations.

--keep-attrs Select which attribute(s) will be preserved. This space separated list of attributes indicates what to preserve. Supports wildcards.

--reject-attrs Select which attribute(s) will be discarded. This space separated list of attributes indicates what to discard. Supports wildcards.

How is this different that btool?

Some of the things filter can do functionally overlaps with **`btool list`**. Take for example:

```
ksconf filter search/default/savedsearches.conf --stanza "Messages by minute last 3 hours"
```

Is essentially the same as:

```
splunk btool --app=search savedsearches list "Messages by minute last 3 hours"
```

The output is the same, assuming that you didn't overwrite any part of that search in `local`. But if you take off the `--app` argument, you'll quickly see that `btool` is merging all the layers together to show the final value of all attributes. That is certainly a helpful thing to do, but not always what you want.

Ksconf is *only* going to look at the file you explicitly pointed it to. It doesn't traverse the tree on it's own. This means that it works on app directory structure that live inside or outside of your Splunk instance. If you've ever tried to run `btool check` on an app that you haven't installed yet, then you'll understand the value of this.

In many other cases, the usage of both `ksconf filter` and `btool` differ significantly.

Examples

Lift and shift

Copy all indexes defined within a specific app.

```
cd $SPLUNK_DB
for idx in $(ksconf filter $SPLUNK_HOME/etc/app/MyApp/default/indexes.conf --brief)
do
    echo "Copy index ${idx}"
    tar -czf "/migrate/export-${idx}" "${idx}"
done
```

Now you'll have a copy all of the necessary indexes in the `/migrate` folder to make *MyApp* work on another Splunk instance. Of course, there's likely other migration tasks to consider, like copying the actual app. This is just one way ksconf can help.

Can I do the same thing with standard unix tools?

Sure, go for it!

Yes, there's significant overlap with the filter command and what you can do with **grep**, **awk**, or **sed**. Much of that is on purpose, and in fact some command line arguments were borrowed.

I used to do these tasks by hand, but it's easy to make mistakes. The idea of **ksconf** is to give you stable and reliable tools that are more suitable for `.conf` file work. Also keep in mind that these features are expanding much more quickly than the unix tools change.

Although, if you've had to deal with BSD vs GNU tools and trying to find a set of common arguments, then you probably already appreciate how awesome a domain-specific-tool like this is.

3.4.6 ksconf merge

Merge two or more .conf files into a single combined .conf file. This is similar to the way that Splunk logically combines the `default` and `local` folders at runtime.

```
usage: ksconf merge [-h] [--target FILE] [--ignore-missing] [--dry-run]
                    [--banner BANNER]
                    FILE [FILE ...]
```

Positional Arguments

FILE The source configuration file(s) to collect settings from.

Named Arguments

--target, -t Save the merged configuration files to this target file. If not provided, the merged conf is written to standard output.

--ignore-missing, -s Silently ignore any missing CONF files.

--dry-run, -D Enable dry-run mode. Instead of writing to TARGET, preview changes in 'diff' format. If TARGET doesn't exist, then show the merged file.

--banner, -b A banner or warning comment added to the top of the TARGET file. Used to discourage Splunk admins from editing an auto-generated file.

Examples

Here is an elementary example that merges all `props.conf` file from *all* of your technology addons into a single output file:

```
ksconf merge --target=all-ta-props.conf etc/apps/*TA*/{default,local}/props.conf
```

See an expanded version of this example here: *Building an all-in one TA for your indexing tier*

3.4.7 ksconf minimize

See also:

See the *Minimizing files* for background on why this is important.

Minimize a conf file by removing any duplicated default settings.

Reduce a local conf file to only your intended changes without manually tracking which entries you've edited. Minimizing local conf files makes your local customizations easier to read and often results in cleaner upgrades.

```
usage: ksconf minimize [-h] [--target TARGET] [--dry-run | --output OUTPUT]
                       [--explode-default] [-k PRESERVE_KEY]
                       CONF [CONF ...]
```

Positional Arguments

CONF The default configuration file(s) used to determine what base settings are. The base settings determine what is unnecessary to repeat in target file.

Named Arguments

--target, -t The local file that you wish to remove duplicate settings from. This file will be read from and then replaced with a minimized version.

--dry-run, -D Enable dry-run mode. Instead of writing and minimizing the TARGET file, preview what would be removed as a 'diff'.

--output Write the minimized output to a separate file instead of updating TARGET.

This option can be used to *preview* the actual changes. Sometimes if `--dry-run` mode produces too much output, it's helpful to look at the actual minimized version of the file in concrete form (rather than a relative format, like a diff.) This may also be helpful in other workflows.

--explode-default, -E Enable minimization across stanzas for special use-cases. Helpful when dealing with stanzas downloaded from a REST endpoint or `btool list` output.

This mode will not only minimize the same stanza across multiple config files, it will also attempt to minimize any default values stored in the `[default]` or global stanza as well. For this to be effective, it's often necessary to include system-level defaults in the CONF list. For example, to trim out cruft in savedsearches.conf, make sure you add `etc/system/default/savedsearches.conf` as an input.

-k, --preserve-key Specify attributes that should always be kept.

Example usage

```
cd Splunk_TA_nix
cp default/inputs.conf local/inputs.conf

# Edit 'disabled' and 'interval' settings in-place
vi local/inputs.conf

# Remove all the extra (unmodified) bits
ksconf minimize --target=local/inputs.conf default/inputs.conf
```

Undoing a minimize

You can use *ksconf merge* to reverse the effect of minimize by running a command like so:

```
ksconf merge default/inputs.conf local/inputs.conf
```

Additional capabilities

For special cases, the `--explode-default` mode reduces duplication between entries in normal stanzas (as normal) and then additionally reduces duplication between individual stanzas and default entries. Typically you only need this mode if your dealing with a conf file that's been fully expanded to include all the layers, which doesn't happen under normal circumstances. This does happen anytime you download a stanza from a REST endpoint or munge together output from `btool list`. If you've ever done this with `savedsearches.conf` stanzas, you'll be painfully aware of how massive they are! This is the exact use case that `--explode-default` was written for.

In such a case, it may be helpful to minimize against the full definition of *default*, which effectively requires looking at all the layers of default. This includes all global app settings, and system-level settings.

There are limitations to this approach.

- You have to manually list out all the layers. (Sometimes just pointing to the system-level defaults is good enough)
- Minimize doesn't take namespace into account. This means ownership, sharing, and ACLs are ignored.

In many ways `minimize` mimics what Splunk does *every* time it updates a conf file, as discussed in *How Splunk writes to conf files*. If you find yourself frequently needing the power of `--explode-default`, at some point a potentially better approach may be to simply post stanzas to the REST endpoint. However, this typically does a good enough job, especially for offline scenarios.

Additionally, this command doesn't strictly require a bloated file. For example, if `disabled = 0` is both a global default, and set on a per-stanza basis, that could be reduced too. However, typically this isn't super helpful.

3.4.8 ksconf promote

Propagate .conf settings applied in one file to another. Typically this is used to move `local` changes (made via the UI) into another layer, such as the `default` or a named `default.d/50-xxxxx`) folder.

Promote has two modes: batch and interactive. In batch mode, all changes are applied automatically and the (now empty) source file is removed. In interactive mode, the user is prompted to select stanzas to promote. This way local changes can be held without being promoted.

NOTE: Changes are *MOVED* not copied, unless `--keep` is used.

```
usage: ksconf promote [-h] [--batch | --interactive | --summary] [--verbose]
                      [--match {regex,wildcard,string}] [--ignore-case]
                      [--invert-match] [--stanza PATTERN] [--force] [--keep]
                      [--keep-empty]
                      SOURCE TARGET
```

Positional Arguments

SOURCE The source configuration file to pull changes from. (Typically the `local` conf file)

TARGET Configuration file or directory to push the changes into. (Typically the `default` folder)

Named Arguments

--batch, -b Use batch mode where all configuration settings are automatically promoted. All changes are removed from source and applied to target. The source file will be removed unless `--keep-empty` is used.

--interactive, -i Enable interactive mode where the user will be prompted to approve the promotion of specific stanzas and attributes. The user will be able to apply, skip, or edit the changes being promoted.

--summary, -s Summarize content that could be promoted.

--verbose Enable additional output.

--force, -f Disable safety checks. Don't check to see if SOURCE and TARGET share the same basename.

--keep, -k Keep conf settings in the source file. All changes will be copied into the TARGET file instead of being moved there. This is typically a bad idea since local always overrides default.

--keep-empty Keep the source file, even if after the settings promotions the file has no content. By default, SOURCE will be removed after all content has been moved into TARGET. Splunk will re-create any necessary local files on the fly.

Automatic filtering options

Include or exclude stanzas to promote using these filter options. Stanzas selected by these filters will be promoted.

All filter options can be provided multiple times. If you have a long list of filters, they can be saved in a file and referenced using the special `file://` prefix. One entry per line.

--match, -m Possible choices: regex, wildcard, string

Specify pattern matching mode. Defaults to 'wildcard' allowing for `*` and `?` matching. Use 'regex' for more power but watch out for shell escaping. Use 'string' to enable literal matching.

--ignore-case Ignore case when comparing or matching strings. By default matches are case-sensitive.

--invert-match, -v Invert match results. This can be used to prevent content from being promoted.

--stanza Promote any stanza with a name matching the given pattern. PATTERN supports bulk patterns via the `file://` prefix.

Warning: The promote command **moves** configuration settings between *SOURCE* and *TARGET* and therefore both files are updated. This is unlike most other commands where only *TARGET* is modified. Using the `--keep` argument will prevent *SOURCE* from being updated.

Modes

Promote has different modes:

Batch mode Changes are applied automatically and the (now empty) source file is removed by default. The source file can be retained by using either the `--keep` or `--keep-empty` arguments, see descriptions above.

Interactive mode Prompts the user to pick which stanzas and attributes to integrate. In practice, it's common that not all local changes will be ready to be promoted and committed at the same time.

Hint: This mode was inspired by **`git add --patch`** command.

Summary mode Shows the user a brief breakdown of what stanzas are available for promotion. This can be used to simply the use of the `--stanza` filtering options (automatic promotion) to show the names of stanzas available for promotion. Note that when `--summary` and `--stanza` are used at the same time, then the summary output will include any output not *already* matched by `--stanza` filter.

Default If you haven't specified either batch or interactive mode, you'll be asked to pick one at startup. You'll be given the option to show a diff, apply all changes, or be prompted to keep or reject changes interactively.

Automated promotions

Ksconf 0.7.8 added support for automatic stanza matching and promotion using a `ksconf filter`-like CLI options.

Key features include:

Automatic promotion of stanzas One or more named stanzas can be promoted automatically using the `--stanza` argument. This argument can be given multiple times to match multiples stanzas at once. In batch mode, only the named stanzas will be promoted; but in interactive mode, the named stanzas will be promoted first, and any content remaining to be promoted can be handled interactively.

Matching mode Like with the `ksconf filter` command, multiple methods of matching are supported. This includes: string matching (default), wildcard (or "glob") matching, and regular expressions.

Inversion The `--invert-match` option allows for the selection to be inverted. In this mode, it's possible to select which stanzas should *not* be promoted. This can be used as a blocklist to prevent accidental promotions.

Safety checks

Moving content between files is a potentially risky operation. Here are some of the safety mechanisms that ksconf has in place to prevent data loss.

Tip: Pairing ksconf with a version control tool like **`git`**, while not required, does provide another layer of protection against loss or corruption. If you promote and commit changes frequently, then the scope of potential loss is reduced.

Syntax checking Strong syntax checking is enabled for both *SOURCE* and *TARGET* to prevent mistakes, such as dangling or duplicate stanzas, which could lead to even more corruption.

File fingerprinting Various attributes of the *SOURCE* and *TARGET* files are captured at startup and compared again before any changes are written to disk. This reduces the possibility of a race-condition on a live Splunk system. This mostly impacts interactive mode because the session lasts longer. If this is a concern, run promote only when Splunk is offline.

Same file check Attempts to promote content from a file to itself are prevented. While logically no one would want to do this, in practice having a clear error message saves time and confusion.

Base name check The *SOURCE* and *TARGET* should share the same base name. In other words, trying to promote from `inputs.conf` into `props.conf` (due to a typo) will be prevented. This matters more in batch mode. In interactive mode, it should be pretty obvious that the type of entries don't make sense and therefore the user can simply exit without saving.

For scripting purposes, there may be times where pushing changes between arbitrary-named files is helpful, so this check can be bypassed by using the `--force` argument.

Note: Unfortunately, the unit testing coverage for the `promote` command is quite low. This is primarily because I haven't yet figured out how to handle unit testing for interactive CLI tools (as this is the only interactive command to date.) I'm also not sure how much the UI may change; Any assistance in this area would be greatly appreciated.

Examples

A simple promotion looks like this.

```
ksconf promote local/props.conf default/props.conf
```

This is equivalent to this minor shortcut.

```
ksconf promote local/props.conf default
```

In this case, ksconf determines that `default` is a directory and therefore assumes that you want the same filename, `props.conf` in this case.

Tip: Using a directory as TARGET may seem like a trivial improvement, but in practice it greatly reduces accidental cross-promotion of content. Therefore, we suggest its use.

Similarly, a shortcut for pushing between metadata files exists:

```
ksconf promote metadata/local.meta metadata
```

A few example of automatic promotion of a named stanza:

```
# Single stanzas
ksconf promote local/savedsearches.conf default --stanza "My fancy search"

# Wildcard promote all prod server alerts
ksconf promote local/savedsearches.conf default --match wildcard --stanza "Server␣
↪PRD* Alert"

# Automatically promote everything except for one search:
ksconf promote local/savedsearches.conf default --batch --invert-match --stanza
↪"Local test"
```

Interactive mode

Keyboard shortcuts

Key	Meaning	Description
y	Yes	Apply changes
n	No	Don't apply
d	Diff	Show the difference between the file or stanza.
q	Quit	Exit program. Don't save changes.

Limitations

- Currently, an attribute-level section has not be implemented. Entire stanzas are either kept local or promoted fully.

- Interactive mode currently lacks "help". In the meantime, see the keyboard shortcuts listed above.
- At present, comments in the *SOURCE* file will not be preserved.
- If *SOURCE* or *TARGET* is modified externally while promote is running, the entire operation will be aborted, thus loosing any custom selections you made in interactive mode. This needs improvement.
- There's currently no way to preserve certain local settings with some kind of "never-promote" flag. It's not uncommon to have some settings in `inputs.conf`, for example, that you never want to promote.
- There is no *dry-run* mode supported. Primarily, this is because it would only work for batch mode, and in interactive mode you explicitly see exactly what will be changed before anything is applied. (If you really need a dry-run for batch mode, use *ksconf merge* to show the result of *TARGET SOURCE* combined.)

3.4.9 ksconf rest-export

Deprecated since version 0.7.0: You should consider using *ksconf rest-publish* instead of this one. The only remaining valid use case for `rest-export` (this command) is for disconnected scenarios. In other words, if you need to push stanzas to a Splunkd instance where you don't (and can't) install `ksconf`, then this command may still be useful to you. In this case, `ksconf rest-export` can create a shell script that you can transfer to the correct network, and then run the shell script. But for **ALL** other use cases, the `rest-publish` command is superior.

Build an executable script of the stanzas in a configuration file that can be later applied to a running Splunk instance via the Splunkd REST endpoint.

This can be helpful when pushing complex props and transforms to an instance where you only have UI access and can't directly publish an app.

```
usage: ksconf rest-export [-h] [--output FILE] [--disable-auth-output]
                          [--pretty-print] [-u | -D] [--url URL] [--app APP]
                          [--user USER] [--owner OWNER] [--conf TYPE]
                          [--extra-args EXTRA_ARGS]
                          CONF [CONF ...]
```

Positional Arguments

CONF Configuration file(s) to export settings from.

Named Arguments

--output, -t Save the shell script output to this file. If not provided, the output is written to standard output.

-u, --update Assume that the REST entities already exist. By default, output assumes stanzas are being created.

-D, --delete Remove existing REST entities. This is a destructive operation. In this mode, stanza attributes are unnecessary and ignored. NOTE: This works for 'local' entities only; the default folder cannot be updated.

--url URL of Splunkd. Default: "https://localhost:8089"

--app Set the namespace (app name) for the endpoint

--user Deprecated. Use –owner instead.

--owner Set the object owner. Typically, the default of 'nobody' is ideal if you want to share the configurations at the app-level.

--conf Explicitly set the configuration file type. By default, this is derived from CONF, but sometimes it's helpful to set this explicitly. Can be any valid Splunk conf file type. Examples include: 'app', 'props', 'tags', 'savedsearches', etc.

--extra-args Extra arguments to pass to all CURL commands. Quote arguments on the command line to prevent confusion between arguments to ksconf vs curl.

Output Control

--disable-auth-output Turn off sample login curl commands from the output.

--pretty-print, -p Enable pretty-printing. Make shell output a bit more readable by splitting entries across lines.

> **Warning:** For interactive use only!
>
> This command is indented for manual admin workflows. It's quite possible that shell escaping bugs exist that may allow full shell access if you put this into an automated workflow. Evaluate the risks, review the code, run as a least-privilege user, and be responsible.

Roadmap

For now, the assumption is that `curl` command will be used. (Patches to support the Power Shell `Invoke-WebRequest` cmdlet would be greatly welcomed!)

Example

```
ksconf rest-export --output=apply_props.sh etc/app/Splunk_TA_aws/local/props.conf
```

3.4.10 ksconf rest-publish

Note: This command effectively replaces *ksconf rest-export* for nearly all use cases. The only thing that `rest-publish` can't do that `rest-export` can, is handle a disconnected scenario. But for **ALL** other use cases, the `rest-publish` (this command) command is far superior.

Note: This commands requires the Splunk Python SDK, which is automatically bundled with the *Splunk app for KSCONF*.

Publish stanzas in a .conf file to a running Splunk instance via REST. This requires access to the HTTPS endpoint of Splunk. By default, ksconf will handle both the creation of new stanzas and the update of existing stanzas.

This can be used to push full configuration stanzas where you only have REST access and can't directly publish an app.

Only attributes present in the conf file are pushed. While this may seem obvious, this fact can have profound implications in certain situations, like when using this command for continuous updates. This means that it's possible for the source .conf to ultimately differ from what ends up on the server's .conf file. One way to avoid this, is to explicitly remove an object using `--delete` mode first, and then insert a new copy of the object. Of course, this means that the object will be unavailable. The other impact is that diffs only compares and shows a subset of attribute.

Be aware, that for consistency, the configs/conf-TYPE endpoint is used for this command. Therefore, a reload may be required for the server to use the published config settings.

```
usage: ksconf rest-publish [-h] [--conf TYPE] [-m META] [--url URL]
                           [--user USER] [--pass PASSWORD] [-k]
                           [--session-key SESSION_KEY] [--app APP]
                           [--owner OWNER] [--sharing {user,app,global}] [-D]
                           CONF [CONF ...]
```

Positional Arguments

CONF Configuration file(s) to export settings from.

Named Arguments

--conf Explicitly set the configuration file type. By default, this is derived from CONF, but sometimes it's helpful to set this explicitly. Can be any valid Splunk conf file type. Examples include: 'app', 'props', 'tags', 'savedsearches', etc.

-m, --meta Specify one or more `.meta` files to determine the desired read & write ACLs, owner, and sharing for objects in the CONF file.

--url	URL of Splunkd. Default: "https://localhost:8089"
--user	Login username Splunkd. Default: "admin"
--pass	Login password Splunkd. Default: "changeme"
-k, --insecure	Disable SSL cert validation.
--session-key	Use an existing session token instead of using a username and password to login.
--app	Set the namespace (app name) for the endpoint
--owner	Set the user who owns the content. The default of 'nobody' works well for app-level sharing.
--sharing	Possible choices: user, app, global Set the sharing mode.
-D, --delete	Remove existing REST entities. This is a destructive operation. In this mode, stanza attributes are unnecessary. NOTE: This works for 'local' entities only; the default folder cannot be updated.

Examples

A simple example:

```
ksconf rest-publish etc/app/Splunk_TA_aws/local/props.conf \
    --user admin --password secret --app Splunk_TA_aws --owner nobody --sharing global
```

This command also supports replaying metdata like ACLs:

```
ksconf rest-publish etc/app/Splunk_TA_aws/local/props.conf \
    --meta etc/app/Splunk_TA_aws/metdata/local.meta \
    --user admin --password secret --app Splunk_TA_aws
```

3.4.11 ksconf snapshot

Build a static snapshot of various configuration files stored within a structured json export format. If the .conf files being captured are within a standard Splunk directory structure, then certain metadata and namespace information is assumed based on typical path locations. Individual apps or conf files can be collected as well, but less metadata may be extracted.

```
usage: ksconf snapshot [-h] [--output FILE] [--minimize] PATH [PATH ...]
```

Positional Arguments

PATH	Directory from which to load configuration files. All .conf and .meta file are included recursively.

Named Arguments

--output, -o Save the snapshot to the named files. If not provided, the snapshot is written to standard output.

--minimize Reduce the size of the JSON output by removing whitespace. Reduces readability.

> **Warning: Output NOT stable!**
>
> The output from this command hasn't really been tested in any kind of serious way for usability. Consider this a proof-of-concept. Anyone interested in this type of functionality should *reach out* to discuss uses cases.

Example

```
ksconf snapshot --output=daily-$(date +%Y-%m-%d).json $SPLUNK_HOME/etc/app/
```

3.4.12 ksconf sort

Sort a Splunk .conf file. Sort has two modes: (1) by default, the sorted config file will be echoed to the screen. (2) the config files are updated in-place when the `-i` option is used.

Manually managed conf files can be protected against changes by adding a comment containing the string `KSCONF-NO-SORT` to the top of any .conf file.

```
usage: ksconf sort [-h] [--target FILE | --inplace] [-F] [-q] [-n LINES]
                   FILE [FILE ...]
```

Positional Arguments

FILE Input file to sort, or standard input.

Named Arguments

--target, -t File to write results to. Defaults to standard output.

--inplace, -i Replace the input file with a sorted version.

WARNING: This a potentially destructive operation that may move/remove comments.

-n, --newlines Number of lines between stanzas.

In-place update arguments

-F, --force Force file sorting for all files, even for files containing the special 'KSCONF-NO-SORT' marker.

-q, --quiet Reduce the output. Reports only updated or invalid files. This is useful for pre-commit hooks, for example.

See also:

Pre-commit hooks

See *Pre-commit hooks* for more information about how the `sort` command can be easily integrated in your git workflow.

Examples

To recursively sort all files

```
find . -name '*.conf' | xargs ksconf sort -i
```

3.4.13 ksconf unarchive

Install or overwrite an existing app in a git-friendly way. If the app already exists, steps will be taken to upgrade it safely.

The `default` folder can be redirected to another path (i.e., `default.d/10-upstream` or other desirable path if you're using the `ksconf combine` tool to manage extra layers).

```
usage: ksconf unarchive [-h] [--dest DIR] [--app-name NAME]
                        [--default-dir DIR] [--exclude EXCLUDE] [--keep KEEP]
                        [--allow-local]
                        [--git-sanity-check {off,changed,untracked,ignored}]
                        [--git-mode {nochange,stage,commit}] [--no-edit]
                        [--git-commit-args GIT_COMMIT_ARGS]
                        SPL
```

Positional Arguments

SPL The path to the archive to install.

Supports tarballs (.tar.gz, .spl), and less-common zip files (.zip)

Named Arguments

--dest Set the destination path where the archive will be extracted. By default, the current directory is used. Sane values include: etc/apps, etc/deployment-apps, and so on.

Often this will be a git repository working tree where Splunk apps are stored.

--app-name The app name to use when expanding the archive. By default, the app name is taken from the archive as the top-level path included in the archive (by convention).

Expanding archives that contain multiple (ITSI) or nested apps (NIX, ES) is not supported.

--default-dir Name of the directory where the default contents will be stored. This is a useful feature for apps that use a dynamic default directory that's created and managed by the 'combine' mode.

--exclude, -e Add a file pattern to exclude from extraction. Splunk's pseudo-glob patterns are supported here. `*` for any non-directory match, `...` for ANY (including directories), and `?` for a single character.

--keep, -k Specify a pattern for files to preserve during an upgrade. Repeat this argument to keep multiple patterns.

--allow-local Allow local/* and local.meta files to be extracted from the archive.

Shipping local files is a Splunk app packaging violation so local files are blocked to prevent customizations from being overridden.

--git-sanity-check By default, `git status` is run on the destination folder to detect working tree or index modifications before the unarchive process starts, but this is configurable. Sanity check choices go from least restrictive to most thorough:

- Use `off` to prevent any 'git status' safety checks.
- Use `changed` to abort only upon local modifications to files tracked by git.
- Use `untracked` (the default) to look for changed and untracked files before considering the tree clean.
- Use `ignored` to enable the most intense safety check which will abort if local changes, untracked, or ignored files are found.

--git-mode Possible choices: nochange, stage, commit

Set the desired level of git integration. The default mode is *stage*, where new, updated, or removed files are automatically handled for you.

To prevent any `git add` or `git rm` commands from being run, pick the 'nochange' mode.

If a git commit is incorrect, simply roll it back with `git reset` or fix it with a `git commit --amend` before the changes are pushed anywhere else. There's no native `--dry-run` or undo for unarchive mode because that's why you're using git in the first place, right? (Plus, such features would require significant overhead and unit testing.)

--no-edit Tell git to skip opening your editor on commit. By default, you will be prompted to review/edit the commit message. (Git Tip: Delete the content of the default message to abort the commit.)

--git-commit-args, -G Extra arguments to pass to 'git'

Note: What if I'm not using version control?

Sanity checks and commit modes are automatically disabled if the app is being installed into a directory that is *not* contained within a git working tree. Ksconf confirms that *git* is present and functional before running sanity checks.

3.4.14 ksconf xml-format

Normalize and apply consistent XML indentation and CDATA usage for XML dashboards and navigation files.

Technically this could be used on *any* XML file, but certain element names specific to Splunk's simple XML dashboards are handled specially, and therefore could result in unusable results.

The expected indentation level is guessed based on the first element indentation, but can be explicitly set if not detectable.

```
usage: ksconf xml-format [-h] [--indent INDENT] [--quiet] FILE [FILE ...]
```

Positional Arguments

FILE One or more XML files to check. If '-' is given, then a list of files is read from standard input

Named Arguments

--indent Number of spaces. This is only used if indentation cannot be guessed from the existing file.

--quiet, -q Reduce the volume of output.

See also:

Pre-commit hooks

See *Pre-commit hooks* for more information about how the `xml-format` command can be integrated in your git workflow.

NOTE: While it may work on other XML files, it hasn't been tested for other files, and therefore is not recommended as a general-purpose XML formatter. Specific awareness of various Simple XML tags is baked into this product.

Note: This command requires the external `lxml` Python module.

This package was specifically selected (over the built-in 'xml.etree' interface) because it (1) supports round-trip preservation of CDATA blocks, and (2) already ships with Splunk's embedded Python.

This is an optional requirement, unless you want to use the `xml-format` command. However, due to packaging limitations and pre-commit hook support, installation of the python package will result in an attempt to install lxml as well. Please *reach out* if this is causing issues for you; I'm looking into other options too.

Why is this important?

TODO: Note the value of using `<!CDATA[[ ]]>` blocks.

Value of consistent indentation.

To recursively format xml files

```
find . -path '*/data/ui/views/*.xml' -o -path '*/data/ui/nav/*.xml' | ksconf xml-format -
```

3.5 Cheat Sheet

Here's a quick rundown of handy `ksconf` commands:

Note: Note that for clarity, most of the command line arguments are given in their long form, but many options also have a short form.

Long commands may be broken across line for readability. When this happens, a trailing backslash (`\`) is added so the command could still be copied verbatim into most shells.

Contents

- *Cheat Sheet*
  - *General purpose*
    - *Comparing files*
    - *Sorting content*
    - *Extract specific stanza*
    - *Remove unwanted settings*

- *Cleaning up*
  - *Reduce cruft in local*
  - *Pushing local changes to default*
- *Advanced usage*
  - *Migrating content between apps*
  - *Migrating the 'users' folder*
  - *Maintaining apps stored in a local git repository*
- *Putting it all together*
  - *Pulling out a stanza defined in both default and local*
  - *Building an all-in one TA for your indexing tier*

3.5.1 General purpose

Comparing files

Show the differences between two conf files using *ksconf diff*.

```
ksconf diff savedsearches.conf savedsearches-mine.conf
```

Sorting content

Create a normalized version of a configuration file, making conf files easier to merge with **git**. Run an in-place sort like so:

```
ksconf sort --inplace savedsearches.conf
```

Tip: Use the `ksconf-sort` *pre-commit* hook to do this for you.

Extract specific stanza

Say you want to *grep* your conf file for a specific stanza pattern:

```
ksconf filter search/default/savedsearches.conf --stanza 'Errors in the last *'
```

Say you want to list stanzas containing `cron_schedule`:

```
ksconf filter Splunk_TA_aws/default/savedsearches.conf --brief \
    --attr-present 'cron_schedule'
```

Remove unwanted settings

Say you want to remove `vsid` from a legacy savedsearches file:

```
ksconf filter search/default/savedsearches.conf --reject-attrs "vsid"
```

To see just to the schedule and scheduler status of scheduled searches, run:

```
ksconf filter Splunk_TA_aws/default/savedsearches.conf \
    --attr-present cron_schedule \
    --keep-attrs 'cron*' \
    --keep-attrs enableSched
    --keep-attrs disabled
```

3.5.2 Cleaning up

Reduce cruft in local

If you're in the habit of copying the *default* files to *local* in the TAs you deploy, here is a quick way to 'minimize' your files. This will reduce the *local* file by removing all the *default* settings you copied but didn't change. (The importance of this is outlined in *Minimizing files*.)

```
ksconf minimize Splunk_TA_nix/default/inputs.conf --target Splunk_TA_nix/local/
↪inputs.conf
```

Pushing local changes to default

App developers can push changes from the `local` folder to the `default` folder:

```
ksconf promote --interactive myapp/local/props.conf myapp/default/props.conf
```

You will be prompted to pick which items you want to promote. Alternatively, use the `--batch` option to promote everything in one step, without reviewing the changes first.

3.5.3 Advanced usage

Migrating content between apps

Say you want to move a bunch of savedsearches from `search` into a more appropriate app. First create a file that lists all the names of your searches (one per line) in `corp_searches.txt`. Next, copy just the desired stanzas, those named in the 'corp_searches' file, over to your new `corp_app` application.

```
ksconf filter --match string --stanza 'file://corp_searches.txt' \
    search/local/savedsearches.conf --output corp_app/default/savedsearches.conf
```

Now, to avoid duplication and confusion, you want to remove that exact same set of searches from the search app.

```
ksconf filter --match string --stanza 'file://corp_searches.txt' \
    --invert-match search/local/savedsearches.conf \
    --output search/local/savedsearches.conf.NEW

# Backup the original
mv search/local/savedsearches.conf \
    /my/backup/location/search-savedsearches-$(date +%Y%M%D).conf

# Move the updated file in place
mv search/local/savedsearches.conf.NEW search/local/savedsearches.conf
```

Note: Setting the matching mode to `string` prevents any special characters that may be present in your search names from being interpreted as wildcards.

Migrating the 'users' folder

Say you stood up a new Splunk server and the migration took longer than expected. Now you have two `users` folders and don't want to loose all the goodies stored in either one. You've copied the users folder to `user_old`. You're working from the new server and would generally prefer to keep whatever is on the new server over what is on the old. (This is because some of your users copied over some of their critical alerts manually while waiting for the migration to complete, and they've made updates they don't want to lose.)

After stopping Splunk on the new server, run the following commands.

```
mv /some/share/users_old  $SPLUNK_HOME/etc/users.old
mv $SPLUNK_HOME/etc/users $SPLUNK_HOME/etc/users.new

ksconf combine $SPLUNK_HOME/etc/users.old $SPLUNK_HOME/etc/users.new \
    --target $SPLUNK_HOME/etc/users --banner ''
```

Now double check the results and start Splunk.

We use the `--banner` option here to essential disable an output banner. Because, in this case, the combine operation is a one-time job and therefore no warning is needed.

Maintaining apps stored in a local git repository

```
ksconf unarchive
```

3.5.4 Putting it all together

Pulling out a stanza defined in both default and local

Say you wanted to count the number of searches containing the word `error`

```
ksconf merge default/savedsearches.conf local/savedsearches.conf \
    | ksconf filter - --stanza '*Error*' --ignore-case --count
```

This is a simple example of chaining two basic **ksconf** commands together to perform a more complex operation. The first command handles the merge of default and local `savedsearches.conf` into a single output stream. The second command filters the resulting stream finding stanzas containing the word 'Error'.

Building an all-in one TA for your indexing tier

Say you need to build a single TA containing all the index-time settings for your indexing tier. (Note: Enterprise Security does something similar when generating the indexer app.)

```
ksconf merge etc/apps/*TA*/{default,local}/props.conf \
    | ksconf filter --output=TA-for-indexers/default/props.conf \
      --include-attr 'TRANSFORMS*' \
      --include-attr 'TIME_*' \
      --include-attr 'MUST_BREAK*' \
      --include-attr 'SHOULD_LINEMERGE' \
      --include-attr 'EVENT_BREAKER*' \
      --include-attr 'LINE_BREAKER*'
```

This example is incomplete because it doesn't list *every* index-time `props.conf` attribute, and leaves out `transforms.conf` and `fields.conf`, but hopefully you get the idea.

3.6 Contributing

Pull requests are greatly welcome! If you plan on contributing code back to the main `ksconf` repo, please follow the standard GitHub fork and pull-request work-flow. We also ask that you enable a set of git hooks to help safeguard against avoidable issues.

3.6.1 Pre-commit hook

The ksconf project uses the pre-commit hook to enable the following checks:

- Fixes trailing whitespace, EOF, and EOLs
- Confirms python code compiles (AST)
- Blocks the committing of large files and keys
- Rebuilds the dynamic portions of the docs related to the CLI.

- Confirms that all unit tests pass. (Currently, this is the same test run by Travis CI, but since tests complete in under 5 seconds, the run-everywhere approach seems appropriate for now. Eventually, the local testing will likely become a subset of the full test suite.)

Note: Multiple uses of pre-commit

Be aware, that the ksconf repo both uses pre-commit for validation of it's own content, and it provides a pre-commit hook service definition for other repos. The first scenario is discussed in this section of the guide. The second scenario is for repositories that house Splunk apps to use *ksconf check* and *ksconf sort* as easy to use hooks against their own `.conf` files which is discussed further in *Pre-commit hooks*.

Installing the pre-commit hook

To ensure your changes comply with the ksconf coding standards, please install and activate pre-commit.

Install:

```
sudo pip install pre-commit

# Register the pre-commit hooks (one time setup)
cd ksconf
pre-commit install --install-hooks
```

Install gitlint

Gitlint will check to ensure that commit messages are in compliance with the standard subject, empty-line, and body format. You can enable it with:

```
gitlint install-hook
```

3.6.2 Refresh module listing

After making changes to the module hierarchy or simply adding new commands, refresh the listing for the autodoc extension by running the following command. Note that this may not remove old packages.

```
sphinx-apidoc -o docs/source/ ksconf --force
```

3.6.3 Create a new subcommand

Checklist:

1. Create a new module in `ksconf.commands.<CMD>`.

- Create a new class derived from `KsconfCmd`. You must, at a minimum, define the following methods:
  - `register_args()` to setup any config parser inputs.
  - `run()` which handles the actual execution of the command.

2. Register a new entrypoint configuration in the `setup_entrypoints.py` script. Edit the `_entry_points` dictionary to add an entry for the new command.
   - Each entry must include command name, module, and implementation class.
3. Create unit tests in `test/test_cli_<CMD>.py`.
4. Create documentation in `docs/source/cmd_<CMD>.rst`. You'll want to build the docs locally to make sure everything looks correct. Part of the documentation is automatically generated from the argparse arguments defined in the `register_args()` method, but other bits need to be spelled out explicitly.

When in doubt, it may be helpful to look back over history in git for other recently added commands and use that as an example.

3.7 Developer setup

The following steps highlight the developer install process.

3.7.1 Tools

If you are a developer, then we strongly suggest installing into a virtual environment to prevent overwriting the production version of ksconf and for the installation of the developer tools. (The virtualenv name `ksconfdev-pyve` is used below, but this can be whatever suites, just make sure not to commit it.)

```
# Setup and activate virtual environment
virtualenv ksconfdev-pyve
. ksconfdev-pyve/bin/activate

# Install developer packages
pip install -r requirements-dev.txt
```

3.7.2 Install ksconf

```
git clone https://github.com/Kintyre/ksconf.git
cd ksconf
pip install .
```

3.7.3 Building the docs

```
cd ksconf
. ksconfdev-pyve/bin/activate

cd docs
make html
open build/html/index.html
```

If you'd like to build PDF, then you'll need some extra tools. On Mac, you may also want to install the following (for building docs, etc.):

```
brew install homebrew/cask/mactex-no-gui
```

3.7.4 Running TOX

Local testing across multiple versions of python can be accomplished with tox and pyenv. See the online docs for theses tools for more details.

Tox and pyenv can be run like so:

```
# Install the necessary python versions
pyenv install 2.7.17
...
pyenv install 3.8.1

# Set specific default version of python for each major/minor release (tab completion is␣
↪your friend here)
pyenv local 2.7.17 ... 3.8.1

# Run tox for ALL python versions
tox

# Run tox for just a specific python version
tox -e py38
```

Some additional information about how to setup and run these tests can be gleaned from the `Vagrantfile` and `Dockerfile` in the root of the git repository, though specific python versions contained there may be quite out of date.

3.8 Git tips & tricks

3.8.1 Pre-commit hooks

Ksconf is setup to work as a pre-commit plugin. To use ksconf in this manner, simply configure the ksconf repo in your pre-commit configuration file. If you haven't done any of this before, it's not difficult to setup but is beyond the scope of this guide. We suggest that you read the pre-commit docs and review this section when you are ready to setup the hooks.

Hooks provided by ksconf

Three hooks are currently defined by the ksconf repository:

ksconf-check Runs *ksconf check* to perform basic validation tests against all files in your repo that end with `.conf` or `.meta`. Any errors will be reported by the UI at commit time and you'll be able to correct mistakes before bogus files are committed into your repo. If you're not sure why you'd need this, check out *Why validate my conf files?*

ksconf-sort Runs *ksconf sort* to normalize any of your `.conf` or `.meta` files which will make diffs more readable and merging more predictable. As with any hook, you can customize the filename pattern of which files this applies to. For example, to manually organize `props.conf` files, simply add the `exclude` setting. *See Example below.*

ksconf-xml-format: Runs *ksconf xml-format* to apply consistency to your XML representations of Simple XML dashboards and navigation files. Formatting includes appropriate indention and the automatic addition of `<![CDATA[ ... ]]>` blocks, as needed, to reduce the need for XML escaping, resulting in more readable source file. By default, this hook looks at standard locations where XML views and navigation typically live. So if you use Advanced XML, proceed with caution, as they share the same path and haven't been tested.

Configuring pre-commit hooks in you repo

To add ksconf pre-commit hooks to your repository, add the following content to your `.pre-commit-config.yaml` file:

```
repos:
- repo: https://github.com/Kintyre/ksconf
  rev: v0.7.7
  hooks:
    - id: ksconf-check
    - id: ksconf-sort
    - id: ksconf-xml-format
```

For general reference, here's a copy of what we frequently use for our repos.

```
- repo: https://github.com/pre-commit/pre-commit-hooks
  rev: v2.0.0
  hooks:
    - id: trailing-whitespace
      exclude: README.md
    - id: end-of-file-fixer
      exclude: README.md$
    - id: check-json
    - id: check-xml
    - id: check-ast
    - id: check-added-large-files
```

(continues on next page)

(continued from previous page)

```
      args: [ '--maxkb=50' ]
    - id: check-merge-conflict
    - id: detect-private-key
    - id: mixed-line-ending
      args: [ '--fix=lf' ]
- repo: https://github.com/Kintyre/ksconf
  rev: v0.7.7
  hooks:
    - id: ksconf-check
    - id: ksconf-sort
      exclude: (props|logging)\.conf
    - id: ksconf-xml-format
```

Tip: You should update `rev` to the most currently released stable version. Upgrading this frequently isn't typically necessary since these two operations are pretty basic and stable. However, it's still a good idea to review the change log to see what, if any, pre-commit functionality was updated.

Note: Sometimes pre-commit can get in the way.

Instead of disabling it entirely, it's often better to disable the specific rule that's causing an issue using the `SKIP` environmental variable. So for example, if intentionally adding a file over 50 Kb, a command like this will allow all the *other* rules to still run.

```
SKIP=check-added-large-file git commit -m "Refresh lookup files for bogus TA"
```

This and other tricks are fully documented in the pre-commit docs. However, this comes up frequently enough that it's worth repeating here.

Should my version of ksconf and pre-commit plugins be the same?

If you're running both `ksconf` locally as well as the ksconf pre-commit plugin, then technically you have ksconf installed twice. That may sound less than ideal, but practically, this isn't a problem. As long as the version of the ksconf CLI tool is *close* to the `rev` listed in `.pre-commit-config.yaml`, then everything should work fine.

Our suggestion:

1. Keep versions in the same *major.minor* release range or bump the version every 6-12 months.
2. Check the changelog for any pre-commit related changes or compatibility concerns.

While keeping `ksconf` CLI versions in sync across your environment is recommended, it doesn't matter as much for the pre-commit plugin. Why?

1. The pre-commit plugin offers a small subset of overall ksconf functionality.
2. The exposed functionality is stable and changes infrequently.

3. Updating pre-commit too frequently may cause unnecessary delays if you have a large team or high number of git clones throughout your environment, as each one will have to wait and upgrade the next time pre-commit is kicked off.

3.8.2 Git configuration tweaks

Ksconf as external difftool

Use *ksconf diff* as an external *difftool* provider for **git**. Edit `~/.gitconfig` and add the following entries:

```
[difftool "ksconf"]
    cmd = "ksconf --force-color diff \"$LOCAL\" \"$REMOTE\" | less -R"
[difftool]
    prompt = false
[alias]
    ksdiff = "difftool --tool=ksconf"
```

Now you can run this new `git` alias to compare files in your directory using the `ksconf diff` feature instead of the default textual diff that git provides. This is especially helpful if the `ksconf-sort` pre-commit hook hasn't been enabled.

```
git ksdiff props.conf
```

Tip: Wonky version of git?

If you find yourself in the situation where `git-difftool` hasn't been fully installed correctly (or the Perl extensions are missing), then here's a workaround option for you.

```
ksconf diff <(git show HEAD:./props.conf) props.conf
```

Take note of the relative path prefix `./`. In practice, this can be problematic.

Stanza aware textual diffs

Make `git diff` show the 'stanza' on the `@@` output lines.

Note: How does git know that?

Ever wonder how `git diff` is able to show you the name of the function or method where changes were made? This works for many programming languages out of the box. If you've ever spent much time looking at diffs, that additional context is invaluable. As it turns out, this is customizable by adding a stanza matching regular expression with a file pattern match.

Simply add the following settings to your git configuration:

```
[diff "conf"]
    xfuncname = "^(\\[.*\\])$"
```

Then register this new ability with specific file patterns using git's `attributes` feature. Edit `~/.config/git/attributes` and add:

```
*.conf diff=conf
*.meta diff=conf
```

Note: Didn't work as expected?

Be aware that the location for your global-level attributes may be different. Use the following command to test if the settings have been applied.

```
git check-attr -a -- *.conf
```

Test to make sure the `xfuncname` attribute was set as expected:

```
git config diff.conf.xfuncname
```

3.9 Random

3.9.1 Typographic and Convention

Pronounced: k·s·knf

Capitalization:

Form	Acceptability factor
`ksconf`	Always lower for CLI. Generally preferred.
KSCONF	Okay for titles.
Ksconf	Title case is okay too.
KSConf	You'll see this, but weird.
KSconf	Just proper nouns capitalized
KsConf	No, except maybe in a class name?
KsconF	Thought about it. No go! Reserved for ASCII art ONLY

I wrote this while laughing at my own lack of consistency.
– Lowell

3.9.2 How Splunk writes to conf files

Splunk does some counter intuitive thing when it writes to local conf files.

For example,

1. All conf file updates are automatically minimized. Splunk never has to write the entire contents because updates *only* happen to "local" files.
2. Modified stanzas are sometimes rewritten in place, and other times removed from the current position and moved to the bottom of the .conf file. This behavior appears to vary based on what REST endpoint is used to initiate the update.
3. New stanzas are written with attributes sorted lexicographically. When a stanza is updated in place, the modified attributes may be updated in place and new entires are typically added at the bottom of the stanza.
4. Sometimes boolean values persist in unexpected ways. Primarily this is because there's more than one way to represent them textually, and that textual representation is different than what's stored in default. Often, literal values are passed through a conf REST POST so they make it to disk, but when read, are translated into booleans.

Essentially, Splunk will always "minimize" the conf file at each update. This is because Splunk internally keeps track of the final representation of the entire stanza (in memory), and only when it's written to disk does Splunk care about the current contents of the local file. In fact, Splunk re-reads the conf file immediately before updating it. This is why, if you've made a local changes and forgot to reload, Splunk will typically not lose your changes. (Unless you've updated the same attribute both places... I mean, it's not magic.)

Tip: Don't believe me? Try it yourself.

To prove that it works this way, simply find a saved search that you modified from any app that you installed. Look at the local conf file and observe your changes. Now, go edit the saved search and restore some attribute to it's original value; the most obvious one here would be the `search` attribute, but that's tricky if it's multiple lines. Now, go look at the local conf file again. If you've updated it with *exactly* the same value, then that attribute will have been completely removed from the local file. This is in fact a neat trick that can be used to revert local changes to allow future updates to "pass-though" unimpeded. In SHC scenarios, this may be your only option to remove local settings.

Okay, so what's the value in having a *minimize* command if Splunk does this automatically every time it's makes a change? Well, simply put, because Splunk can't write to all local file locations. Splunk only writes to the local folders of system, etc/users, and etc/apps (and sometimes to deployment-apps app.conf local file, but that's a different topic).

Also, there are times where boolean values will show up in an unexpected manor because of how Splunk treats them internally. It isn't certain if this is a silly mistake in the default .conf files or a clever workaround to what's essentially a design flaw in the conf system. Either way, we suspect the user benefits. Because Splunk accepts more values as boolean than what it will write out, certain boolean values will always be explicitly stored in the conf files. This means that `disabled` and several other settings in `savedsearches.conf` always get explicitly written. How is that helpful? Well, imagine what would happen if you accidentally changed `disabled = 1` in the global stanzas in savedsearches.conf. Well, *nothing* if all savedsearches have that values explicitly written. The point is this: there are times when repeating yourself isn't a bad thing. (Incidentally, this is the reason for the `--preserve-key` flag on the *minimize* command.)

3.9.3 Grandfather Paradox

The KSCONF Splunk app disadvantageously breaks it's designed paradigm. Ksconf was designed to be the program that manages all your other apps, so by deploying ksconf as an app itself, we open up the possibility that ksconf could upgrade, deploy, or manage itself. Basically, it could cut off the limb that it's standing on. Practically, this can get messy, especially if you're on Windows, where file locking is also likely to cause issues.

So sure, if you want to be picky, "Grandfather paradox" is probably the wrong analogy. Pull requests are welcome.

3.10 Contact

If you have questions, concerns, ideas about the product or how to make it better, please let us know!

Here are some ways to get in contact with us and other KSCONF users:

- Chat about #ksconf on Splunk's Slack channel.
- Email us at hello@kintyre.co for general inquiries, if you're interested in commercial support, or would like to fund new features.
- Ask a question on
  - Splunk Answers
  - GitHub

3.11 Command line reference

KSCONF supports the following CLI options:

3.11.1 ksconf

```
usage: ksconf [-h] [--version] [--force-color]
              {check,combine,diff,filter,promote,merge,minimize,snapshot,sort,rest-
↪export,rest-publish,unarchive,xml-format}
              ...

Ksconf: Kintyre Splunk CONFig tool

This utility handles a number of common Splunk app maintenance tasks in a small
and easy to deploy package.  Specifically, this tool deals with many of the
nuances with storing Splunk apps in git and pointing live Splunk apps to a git
repository.  Merging changes from the live system's (local) folder to the
version controlled (default) folder and dealing with more than one layer of
"default" are all supported tasks which are not native to Splunk.
```

(continues on next page)

(continued from previous page)

```
positional arguments:
  {check,combine,diff,filter,promote,merge,minimize,snapshot,sort,rest-export,rest-
↪publish,unarchive,xml-format}
    check               Perform basic syntax and sanity checks on .conf files
    combine             Combine configuration files across multiple source
                        directories into a single destination directory. This
                        allows for an arbitrary number of Splunk configuration
                        layers to coexist within a single app. Useful in both
                        ongoing merge and one-time ad-hoc use.
    diff                Compare settings differences between two .conf files
                        ignoring spacing and sort order
    filter              A stanza-aware GREP tool for conf files
    promote             Promote .conf settings between layers using either
                        batch or interactive mode. Frequently this is used to
                        promote conf changes made via the UI (stored in the
                        'local' folder) to a version-controlled directory,
                        such as 'default'.
    merge               Merge two or more .conf files
    minimize            Minimize the target file by removing entries
                        duplicated in the default conf(s)
    snapshot            Snapshot .conf file directories into a JSON dump
                        format
    sort                Sort a Splunk .conf file creating a normalized format
                        appropriate for version control
    rest-export         Export .conf settings as a curl script to apply to a
                        Splunk instance later (via REST)
    rest-publish        Publish .conf settings to a live Splunk instance via
                        REST
    unarchive           Install or upgrade an existing app in a git-friendly
                        and safe way
    xml-format          Normalize XML view and nav files

optional arguments:
  -h, --help            show this help message and exit
  --version             show program's version number and exit
  --force-color         Force TTY color mode on. Useful if piping the output a
                        color-aware pager, like 'less -R'
```

3.11.2 ksconf check

```
usage: ksconf check [-h] [--quiet] FILE [FILE ...]

Provides basic syntax and sanity checking for Splunk's .conf files. Use
Splunk's built-in 'btool check' for a more robust validation of attributes and
values. Consider using this utility as part of a pre-commit hook.

positional arguments:
  FILE          One or more configuration files to check. If '-' is given, then
                read a list of files to validate from standard input
```

(continues on next page)

(continued from previous page)

```
optional arguments:
  -h, --help   show this help message and exit
  --quiet, -q  Reduce the volume of output.
```

3.11.3 ksconf combine

```
usage: ksconf combine [-h] [--target TARGET] [--dry-run] [--follow-symlink]
                      [--banner BANNER] [--disable-marker]
                      source [source ...]

Merge .conf settings from multiple source directories into a combined target
directory.  Configuration files can be stored in a '/etc/*.d' like directory
structure and consolidated back into a single 'default' directory.

This command supports both one-time operations and recurring merge jobs.  For
example, this command can be used to combine all users' knowledge objects (stored
in 'etc/users') after a server migration, or to merge a single user's settings
after their account has been renamed.  Recurring operations assume some type
of external scheduler is being used.  A best-effort is made to only write to
target files as needed.

The 'combine' command takes your logical layers of configs (upstream, corporate,
Splunk admin fixes, and power user knowledge objects, ...) expressed as
individual folders and merges them all back into the single 'default' folder
that Splunk reads from.  One way to keep the 'default' folder up-to-date is
using client-side git hooks.

No directory layout is mandatory, but one simple approach is to model your
layers using a prioritized 'default.d' directory structure. This idea is
borrowed from the Unix System V concept where many services natively read their
config files from '/etc/*.d' directories.

positional arguments:
  source                The source directory where configuration files will be
                        merged from. When multiple source directories are
                        provided, start with the most general and end with the
                        specific; later sources will override values from the
                        earlier ones. Supports wildcards so a typical Unix
                        'conf.d/##-NAME' directory structure works well.

optional arguments:
  -h, --help            show this help message and exit
  --target TARGET, -t TARGET
                        Directory where the merged files will be stored.
                        Typically either 'default' or 'local'
  --dry-run, -D         Enable dry-run mode. Instead of writing to TARGET,
                        preview changes as a 'diff'. If TARGET doesn't exist,
                        then show the merged file.
  --follow-symlink, -l  Follow symbolic links pointing to directories.
```

(continues on next page)

(continued from previous page)

```
                        Symlinks to files are followed.
  --banner BANNER, -b BANNER
                        A banner or warning comment added to the top of the
                        TARGET file. Used to discourage Splunk admins from
                        editing an auto-generated file.
  --disable-marker      Prevents the creation of or checking for the
                        '.ksconf_controlled' marker file safety check. This
                        file is typically used indicate that the destination
                        folder is managed by ksconf. This option should be
                        reserved for well-controlled batch processing
                        scenarios.
```

3.11.4 ksconf diff

```
usage: ksconf diff [-h] [-o FILE] [--comments] CONF1 CONF2

Compares the content differences of two .conf files

This command ignores textual differences (like order, spacing, and comments) and
focuses strictly on comparing stanzas, keys, and values.  Note that spaces within
any given value, will be compared. Multi-line fields are compared in a more␣
↪traditional
'diff' output so that long saved searches and macros can be compared more easily.

positional arguments:
  CONF1                 Left side of the comparison
  CONF2                 Right side of the comparison

optional arguments:
  -h, --help            show this help message and exit
  -o FILE, --output FILE
                        File where difference is stored. Defaults to standard
                        out.
  --comments, -C        Enable comparison of comments. (Unlikely to work
                        consistently)
```

3.11.5 ksconf filter

```
usage: ksconf filter [-h] [-o FILE] [--comments] [--verbose]
                     [--match {regex,wildcard,string}] [--ignore-case]
                     [--invert-match] [--files-with-matches]
                     [--count | --brief] [--stanza PATTERN]
                     [--attr-present ATTR] [--keep-attrs WC-ATTR]
                     [--reject-attrs WC-ATTR]
                     CONF [CONF ...]

Filter the contents of a conf file in various ways. Stanzas can be included or
excluded based on a provided filter or based on the presence or value of a
```

(continues on next page)

(continued from previous page)

```
key. Where possible, this command supports GREP-like arguments to bring a
familiar feel.

positional arguments:
  CONF                  Input conf file

optional arguments:
  -h, --help            show this help message and exit
  -o FILE, --output FILE
                        File where the filtered results are written. Defaults
                        to standard out.
  --comments, -C        Preserve comments. Comments are discarded by default.
  --verbose             Enable additional output.
  --match {regex,wildcard,string}, -m {regex,wildcard,string}
                        Specify pattern matching mode. Defaults to 'wildcard'
                        allowing for '*' and '?' matching. Use 'regex' for
                        more power but watch out for shell escaping. Use
                        'string' to enable literal matching.
  --ignore-case, -i     Ignore case when comparing or matching strings. By
                        default matches are case-sensitive.
  --invert-match, -v    Invert match results. This can be used to show what
                        content does NOT match, or make a backup copy of
                        excluded content.

Output mode:
  Select an alternate output mode. If any of the following options are used,
  the stanza output is not shown.

  --files-with-matches, -l
                        List files that match the given search criteria
  --count, -c           Count matching stanzas
  --brief, -b           List name of matching stanzas

Stanza selection:
  Include or exclude entire stanzas using these filter options. All filter
  options can be provided multiple times. If you have a long list of
  filters, they can be saved in a file and referenced using the special
  'file://' prefix. One entry per line.

  --stanza PATTERN      Match any stanza who's name matches the given pattern.
                        PATTERN supports bulk patterns via the 'file://'
                        prefix.
  --attr-present ATTR   Match any stanza that includes the ATTR attribute.
                        ATTR supports bulk attribute patterns via the
                        'file://' prefix.

Attribute selection:
  Include or exclude attributes passed through. By default, all attributes
  are preserved. Allowlist (keep) operations are preformed before blocklist
  (reject) operations.

  --keep-attrs WC-ATTR  Select which attribute(s) will be preserved. This
```

(continues on next page)

(continued from previous page)

```
                        space separated list of attributes indicates what to
                        preserve. Supports wildcards.
  --reject-attrs WC-ATTR
                        Select which attribute(s) will be discarded. This
                        space separated list of attributes indicates what to
                        discard. Supports wildcards.
```

3.11.6 ksconf promote

```
usage: ksconf promote [-h] [--batch | --interactive | --summary] [--verbose]
                      [--match {regex,wildcard,string}] [--ignore-case]
                      [--invert-match] [--stanza PATTERN] [--force] [--keep]
                      [--keep-empty]
                      SOURCE TARGET

Propagate .conf settings applied in one file to another.  Typically this is used
to move 'local' changes (made via the UI) into another layer, such as the
'default' or a named 'default.d/50-xxxxx') folder.

Promote has two modes:  batch and interactive.  In batch mode, all changes are
applied automatically and the (now empty) source file is removed.  In interactive
mode, the user is prompted to select stanzas to promote.  This way local changes
can be held without being promoted.

NOTE: Changes are *MOVED* not copied, unless '--keep' is used.

positional arguments:
  SOURCE                The source configuration file to pull changes from.
                        (Typically the 'local' conf file)
  TARGET                Configuration file or directory to push the changes
                        into. (Typically the 'default' folder)

optional arguments:
  -h, --help            show this help message and exit
  --batch, -b           Use batch mode where all configuration settings are
                        automatically promoted. All changes are removed from
                        source and applied to target. The source file will be
                        removed unless '--keep-empty' is used.
  --interactive, -i     Enable interactive mode where the user will be
                        prompted to approve the promotion of specific stanzas
                        and attributes. The user will be able to apply, skip,
                        or edit the changes being promoted.
  --summary, -s         Summarize content that could be promoted.
  --verbose             Enable additional output.
  --force, -f           Disable safety checks. Don't check to see if SOURCE
                        and TARGET share the same basename.
  --keep, -k            Keep conf settings in the source file. All changes
                        will be copied into the TARGET file instead of being
                        moved there. This is typically a bad idea since local
                        always overrides default.
```

(continues on next page)

(continued from previous page)

```
  --keep-empty          Keep the source file, even if after the settings
                        promotions the file has no content. By default, SOURCE
                        will be removed after all content has been moved into
                        TARGET. Splunk will re-create any necessary local
                        files on the fly.

Automatic filtering options:
  Include or exclude stanzas to promote using these filter options.
  Stanzas selected by these filters will be promoted.

  All filter options can be provided multiple times.
  If you have a long list of filters, they can be saved in a file and
  referenced using the special 'file://' prefix.  One entry per line.

  --match {regex,wildcard,string}, -m {regex,wildcard,string}
                        Specify pattern matching mode. Defaults to 'wildcard'
                        allowing for '*' and '?' matching. Use 'regex' for
                        more power but watch out for shell escaping. Use
                        'string' to enable literal matching.
  --ignore-case         Ignore case when comparing or matching strings. By
                        default matches are case-sensitive.
  --invert-match, -v    Invert match results. This can be used to prevent
                        content from being promoted.
  --stanza PATTERN      Promote any stanza with a name matching the given
                        pattern. PATTERN supports bulk patterns via the
                        'file://' prefix.
```

3.11.7 ksconf merge

```
usage: ksconf merge [-h] [--target FILE] [--ignore-missing] [--dry-run]
                    [--banner BANNER]
                    FILE [FILE ...]

Merge two or more .conf files into a single combined .conf file. This is
similar to the way that Splunk logically combines the 'default' and 'local'
folders at runtime.

positional arguments:
  FILE                  The source configuration file(s) to collect settings
                        from.

optional arguments:
  -h, --help            show this help message and exit
  --target FILE, -t FILE
                        Save the merged configuration files to this target
                        file. If not provided, the merged conf is written to
                        standard output.
  --ignore-missing, -s  Silently ignore any missing CONF files.
  --dry-run, -D         Enable dry-run mode. Instead of writing to TARGET,
                        preview changes in 'diff' format. If TARGET doesn't
```

(continues on next page)

(continued from previous page)

```
                        exist, then show the merged file.
  --banner BANNER, -b BANNER
                        A banner or warning comment added to the top of the
                        TARGET file. Used to discourage Splunk admins from
                        editing an auto-generated file.
```

3.11.8 ksconf minimize

```
usage: ksconf minimize [-h] [--target TARGET] [--dry-run | --output OUTPUT]
                       [--explode-default] [-k PRESERVE_KEY]
                       CONF [CONF ...]

Minimize a conf file by removing any duplicated default settings. Reduce a
local conf file to only your intended changes without manually tracking which
entries you've edited. Minimizing local conf files makes your local
customizations easier to read and often results in cleaner upgrades.

positional arguments:
  CONF                  The default configuration file(s) used to determine
                        what base settings are. The base settings determine
                        what is unnecessary to repeat in target file.

optional arguments:
  -h, --help            show this help message and exit
  --target TARGET, -t TARGET
                        The local file that you wish to remove duplicate
                        settings from. This file will be read from and then
                        replaced with a minimized version.
  --dry-run, -D         Enable dry-run mode. Instead of writing and minimizing
                        the TARGET file, preview what would be removed as a
                        'diff'.
  --output OUTPUT       Write the minimized output to a separate file instead
                        of updating TARGET.
  --explode-default, -E
                        Enable minimization across stanzas for special use-
                        cases. Helpful when dealing with stanzas downloaded
                        from a REST endpoint or 'btool list' output.
  -k PRESERVE_KEY, --preserve-key PRESERVE_KEY
                        Specify attributes that should always be kept.
```

3.11.9 ksconf snapshot

```
usage: ksconf snapshot [-h] [--output FILE] [--minimize] PATH [PATH ...]

Build a static snapshot of various configuration files stored within a
structured json export format. If the .conf files being captured are within a
standard Splunk directory structure, then certain metadata and namespace
information is assumed based on typical path locations. Individual apps or
```

(continues on next page)

(continued from previous page)

```
conf files can be collected as well, but less metadata may be extracted.

positional arguments:
  PATH                  Directory from which to load configuration files. All
                        .conf and .meta file are included recursively.

optional arguments:
  -h, --help            show this help message and exit
  --output FILE, -o FILE
                        Save the snapshot to the named files. If not provided,
                        the snapshot is written to standard output.
  --minimize            Reduce the size of the JSON output by removing
                        whitespace. Reduces readability.
```

3.11.10 ksconf sort

```
usage: ksconf sort [-h] [--target FILE | --inplace] [-F] [-q] [-n LINES]
                   FILE [FILE ...]

Sort a Splunk .conf file.  Sort has two modes:  (1) by default, the sorted
config file will be echoed to the screen.  (2) the config files are updated
in-place when the '-i' option is used.

Manually managed conf files can be protected against changes by adding a comment␣
↪containing the
string 'KSCONF-NO-SORT' to the top of any .conf file.

positional arguments:
  FILE                  Input file to sort, or standard input.

optional arguments:
  -h, --help            show this help message and exit
  --target FILE, -t FILE
                        File to write results to. Defaults to standard output.
  --inplace, -i         Replace the input file with a sorted version. WARNING:
                        This a potentially destructive operation that may
                        move/remove comments.
  -n LINES, --newlines LINES
                        Number of lines between stanzas.

In-place update arguments:
  -F, --force           Force file sorting for all files, even for files
                        containing the special 'KSCONF-NO-SORT' marker.
  -q, --quiet           Reduce the output. Reports only updated or invalid
                        files. This is useful for pre-commit hooks, for
                        example.
```

3.11.11 ksconf rest-export

```
usage: ksconf rest-export [-h] [--output FILE] [--disable-auth-output]
                          [--pretty-print] [-u | -D] [--url URL] [--app APP]
                          [--user USER] [--owner OWNER] [--conf TYPE]
                          [--extra-args EXTRA_ARGS]
                          CONF [CONF ...]

Build an executable script of the stanzas in a configuration file that can be␣
↪later applied to
a running Splunk instance via the Splunkd REST endpoint.

This can be helpful when pushing complex props and transforms to an instance where␣
↪you only have
UI access and can't directly publish an app.

positional arguments:
  CONF                  Configuration file(s) to export settings from.

optional arguments:
  -h, --help            show this help message and exit
  --output FILE, -t FILE
                        Save the shell script output to this file. If not
                        provided, the output is written to standard output.
  -u, --update          Assume that the REST entities already exist. By
                        default, output assumes stanzas are being created.
  -D, --delete          Remove existing REST entities. This is a destructive
                        operation. In this mode, stanza attributes are
                        unnecessary and ignored. NOTE: This works for 'local'
                        entities only; the default folder cannot be updated.
  --url URL             URL of Splunkd. Default: https://localhost:8089
  --app APP             Set the namespace (app name) for the endpoint
  --user USER           Deprecated. Use --owner instead.
  --owner OWNER         Set the object owner. Typically, the default of
                        'nobody' is ideal if you want to share the
                        configurations at the app-level.
  --conf TYPE           Explicitly set the configuration file type. By
                        default, this is derived from CONF, but sometimes it's
                        helpful to set this explicitly. Can be any valid
                        Splunk conf file type. Examples include: 'app',
                        'props', 'tags', 'savedsearches', etc.
  --extra-args EXTRA_ARGS
                        Extra arguments to pass to all CURL commands. Quote
                        arguments on the command line to prevent confusion
                        between arguments to ksconf vs curl.

Output Control:
  --disable-auth-output
                        Turn off sample login curl commands from the output.
  --pretty-print, -p    Enable pretty-printing. Make shell output a bit more
                        readable by splitting entries across lines.
```

3.11.12 ksconf rest-publish

```
usage: ksconf rest-publish [-h] [--conf TYPE] [-m META] [--url URL]
                           [--user USER] [--pass PASSWORD] [-k]
                           [--session-key SESSION_KEY] [--app APP]
                           [--owner OWNER] [--sharing {user,app,global}] [-D]
                           CONF [CONF ...]

Publish stanzas in a .conf file to a running Splunk instance via REST. This
requires access to the HTTPS endpoint of Splunk. By default, ksconf will
handle both the creation of new stanzas and the update of existing stanzas.
This can be used to push full configuration stanzas where you only have REST
access and can't directly publish an app. Only attributes present in the conf
file are pushed. While this may seem obvious, this fact can have profound
implications in certain situations, like when using this command for
continuous updates. This means that it's possible for the source .conf to
ultimately differ from what ends up on the server's .conf file. One way to
avoid this, is to explicitly remove an object using '--delete' mode first, and
then insert a new copy of the object. Of course, this means that the object
will be unavailable. The other impact is that diffs only compares and shows a
subset of attribute. Be aware, that for consistency, the configs/conf-TYPE
endpoint is used for this command. Therefore, a reload may be required for the
server to use the published config settings.

positional arguments:
  CONF                  Configuration file(s) to export settings from.

optional arguments:
  -h, --help            show this help message and exit
  --conf TYPE           Explicitly set the configuration file type. By
                        default, this is derived from CONF, but sometimes it's
                        helpful to set this explicitly. Can be any valid
                        Splunk conf file type. Examples include: 'app',
                        'props', 'tags', 'savedsearches', etc.
  -m META, --meta META  Specify one or more '.meta' files to determine the
                        desired read & write ACLs, owner, and sharing for
                        objects in the CONF file.
  --url URL             URL of Splunkd. Default: https://localhost:8089
  --user USER           Login username Splunkd. Default: admin
  --pass PASSWORD       Login password Splunkd. Default: changeme
  -k, --insecure        Disable SSL cert validation.
  --session-key SESSION_KEY
                        Use an existing session token instead of using a
                        username and password to login.
  --app APP             Set the namespace (app name) for the endpoint
  --owner OWNER         Set the user who owns the content. The default of
                        'nobody' works well for app-level sharing.
  --sharing {user,app,global}
                        Set the sharing mode.
  -D, --delete          Remove existing REST entities. This is a destructive
                        operation. In this mode, stanza attributes are
                        unnecessary. NOTE: This works for 'local' entities
                        only; the default folder cannot be updated.
```

3.11.13 ksconf unarchive

```
usage: ksconf unarchive [-h] [--dest DIR] [--app-name NAME]
                        [--default-dir DIR] [--exclude EXCLUDE] [--keep KEEP]
                        [--allow-local]
                        [--git-sanity-check {off,changed,untracked,ignored}]
                        [--git-mode {nochange,stage,commit}] [--no-edit]
                        [--git-commit-args GIT_COMMIT_ARGS]
                        SPL

Install or overwrite an existing app in a git-friendly way.
If the app already exists, steps will be taken to upgrade it safely.

The 'default' folder can be redirected to another path (i.e., 'default.d/10-
↪upstream' or
other desirable path if you're using the 'ksconf combine' tool to manage extra␣
↪layers).

positional arguments:
  SPL                   The path to the archive to install.

optional arguments:
  -h, --help            show this help message and exit
  --dest DIR            Set the destination path where the archive will be
                        extracted. By default, the current directory is used.
                        Sane values include: etc/apps, etc/deployment-apps,
                        and so on.
  --app-name NAME       The app name to use when expanding the archive. By
                        default, the app name is taken from the archive as the
                        top-level path included in the archive (by
                        convention).
  --default-dir DIR     Name of the directory where the default contents will
                        be stored. This is a useful feature for apps that use
                        a dynamic default directory that's created and managed
                        by the 'combine' mode.
  --exclude EXCLUDE, -e EXCLUDE
                        Add a file pattern to exclude from extraction.
                        Splunk's pseudo-glob patterns are supported here. '*'
                        for any non-directory match, '...' for ANY (including
                        directories), and '?' for a single character.
  --keep KEEP, -k KEEP  Specify a pattern for files to preserve during an
                        upgrade. Repeat this argument to keep multiple
                        patterns.
  --allow-local         Allow local/* and local.meta files to be extracted
                        from the archive.
  --git-sanity-check {off,changed,untracked,ignored}
                        By default, 'git status' is run on the destination
                        folder to detect working tree or index modifications
                        before the unarchive process start. Sanity check
                        choices go from least restrictive to most thorough:
                        'off' prevents all safety checks. 'changed' aborts
                        only upon local modifications to files tracked by git.
```

(continues on next page)

(continued from previous page)

```
                        'untracked' (the default) looks for changed and
                        untracked files. 'ignored' aborts is (any) local
                        changes, untracked, or ignored files are found.
  --git-mode {nochange,stage,commit}
                        Set the desired level of git integration. The default
                        mode is *stage*, where new, updated, or removed files
                        are automatically handled for you. To prevent any 'git
                        add' or 'git rm' commands from being run, pick the
                        'nochange' mode.
  --no-edit             Tell git to skip opening your editor on commit. By
                        default, you will be prompted to review/edit the
                        commit message. (Git Tip: Delete the content of the
                        default message to abort the commit.)
  --git-commit-args GIT_COMMIT_ARGS, -G GIT_COMMIT_ARGS
                        Extra arguments to pass to 'git'
```

3.11.14 ksconf xml-format

```
usage: ksconf xml-format [-h] [--indent INDENT] [--quiet] FILE [FILE ...]

Normalize and apply consistent XML indentation and CDATA usage for XML
dashboards and navigation files. Technically this could be used on *any* XML
file, but certain element names specific to Splunk's simple XML dashboards are
handled specially, and therefore could result in unusable results. The
expected indentation level is guessed based on the first element indentation,
but can be explicitly set if not detectable.

positional arguments:
  FILE             One or more XML files to check. If '-' is given, then a
                   list of files is read from standard input

optional arguments:
  -h, --help       show this help message and exit
  --indent INDENT  Number of spaces. This is only used if indentation cannot
                   be guessed from the existing file.
  --quiet, -q      Reduce the volume of output.
```

3.12 Changelog

Note: Changes in master, but not released yet are marked as *DRAFT*.

3.12.1 Ksconf 0.7.x

New functionality, massive documentation improvements, metadata support, and Splunk app install fixes.

Release v0.7.9 (2020-09-23)

- Fixed bug where empty stanzas could be removed from `.conf` files. This can be detrimental for `capability::*` entries in `authorize.conf`, for example. A big thanks to nebffa for tracking down this bug!

Release v0.7.8 (2020-06-19)

- New automatic `promote` mode is now available using CLI arguments! This allows stanzas to be selected for promotion from the CLI in batch and interactive modes. This implementation borrows (and shares code) with the `ksconf filter` command so hopefully the CLI arguments look familiar. It's possible to promote a single stanza, a stanza wildcard, regex or invert the matching logic and promote everything except for the named stanza (blocklist). Right now `--stanza` is the only supporting matching mode, but more can be added as needed. A huge thanks to mthambipillai for providing a pull-request with an initial implementation of this feature!
- Added a new summary output mode (`ksconf promote --summary`) that will provide a quick summary of what content could be promoted. This can be used along side the new `--stanza` filtering options to show the names of stanzas that can be promoted.
- Replaced insensitive terminology with race-neutral terms. Specifically the terms 'blacklist' and 'whitelist' have been replaced. NOTE: This does *not* change any CLI attributes, but in a few cases the standard output terminology is slighly different. Also terminology in `.conf` files couldn't be updated as that's controlled by Splunk.
- Fixed bug in the `unarchive` command where a `locale` folder was blocked as a `local` folder and where a nested `default` folder (nested under a Python package, for example) could get renamed if `--default-dir` was used, now only the top-most `default` folder is updated. Also fixed an unlikley bug triggered when `default/app.conf` is missing.
- Fixed bug with `minimize` when the required `--target` argument is not given. This now results in a reminder to the user rather than an unhandled exception.
- Splunk app packaging fix. Write access to the app was previously not granted due to a spelling mistake in the metadata file.

Release v0.7.7 (2020-03-05)

- Added new `--follow-symlink` option to the `combine` command so that input directory structures with symbolic links can be treated the same as proper directories.

- Corrected Windows issue where wildcard (glob) patterns weren't expanded by for `check` and `sort`. This is primarily a difference in how a proper shells (e.g., bash, csh, zsh) handle expansion natively vs CMD on Windows does not. However, since this is typically transparently handled by many CLI tools, we'll follow suite. (BTW, running ksconf from the GIT Bash prompt is a great alternative.) Only the most minimalistic expansion rules will be available, (so don't expect `{props,transforms,app}.conf` to work anytime soon), but this should be good enough for most use cases. Thanks to SID800 for reporting this bug.
- Fixed issues with the `unarchive` command when `git` is not installed or an app is being unarchived (installed/upgrade) into a location not managed by Git. Note that additional output is now enabled when the `KSCONF_DEBUG` environmental variable is set (in lieu of a proper verbose mode). Bug report provided by SID800.
- Enhanced `ksconf --version` output to include Git executable path and version information; as well as a platform dump. (Helpful for future bug reporting.)
- Added feature to disable the marker file (safety check) automatically created by the `combine` command for use in automated processing workflows.
- Updated `pre-commit` documentation and sample configurations to use `rev` rather than `sha` as the means of identifying upstream tags or revisions. Recent releases of `pre-commit` will warn you about this during each run.
- Fixed a temporary file cleanup issue during certain in-place file replacement operations. (If you found any unexpected `*.tmp` files, this could have been the cause.)

Release v0.7.6 (2019-08-15)

- Fresh review and cleanup of all docs! (A huge thank you to Brittany Barnett for this massive undertaking)
- Fixed unhandled exception when encountering a global stanza in metadata files.
- Expand some error messages, sanity checks, and added a new session token (`--session-key`) authentication option for `rest-publish`.

Release v0.7.5 (2019-07-03)

- Fixed a long-term bug where the diff output of a single-line attribute change was incorrectly represented in the textual output of 'ksconf diff' and the diff output in other commands. This resolves a combination of bugs, the first half of which was fixed in 0.7.3.
- Allow `make_docs` script to run on Windows, and other internal doc build process improvements.

Release v0.7.4 (2019-06-07)

- Inline the `six` module to avoid elusive bootstrapping cases where the module couldn't be found. This primarily impacts `pre-commit` users. The `ksconf.ext.*` prefix is being used for this, and any other inlined third party modules we may need in the future.

- Other minor docs fixes and internal non-visible changes.

Release v0.7.3 (2019-06-05)

- Added the new *ksconf xml-format* command.
  - The `ksconf xml-format` command brings format consistency to your XML representations of Simple XML dashboards and navigation files by fixing indentation automatically adding `<![CDATA[ ... ]]>` blocks, as needed, to reduce the need for XML escaping, resulting in more readable source.
  - Additionally, a new pre-commit hook named *ksconf-xml-format* was added to leverage this new functionality. It looks specifically for xml views and navigation files based on path. This may also include Advanced XML, which hasn't been tested; So if you use Advanced XML, proceed with caution.
  - Note that this adds `lxml` as a packaging dependency which is needed for pre-commit hooks, but not strictly required at run time for other ksconf commands. This is NOT ideal, and may change in the future in attempts to keep ksconf as light-weight and standalone as possible. One possible alternative is setting up a different repo for pre-commit hooks. Python packaging and distribution tips welcome.
- Fixed data loss bug in `promote` (interactive mode only) and improved some UI text and prompts.
- Fixed colorization of `ksconf diff` output where certain lines failed to show up in the correct color.
- Fixed bug where debug tracebacks didn't work correctly on Python 2.7. (Enable using `KSCONF_DEBUG=1`.)
- Extended the output of `ksconf --version` to show the names and version of external modules, when present.
- Improved some resource allocation in corner cases.
- Tested with Splunk 7.3 (numeric similarity in version numbers is purely coincidental)

Attention: API BREAKAGE

The `DiffOp` output values for `DIFF_OP_INSERT` and `DIFF_OP_DELETE` have been changed in a backwards-compatible breaking way. The values of `a` and `b` were previously reversed for these two operations, leading to some code confusion.

Release v0.7.2 (2019-03-22)

- Fixed bug where `filter` would crash when doing stanza matching if global entries were present. Global stanzas can be matched by searching for a stanza named `default`.

- Fixed broken `pre-commit` issue that occurred for the `v0.7.1` tag. This also kept `setup.py` from working if the `six` module wasn't already installed. Developers and pre-commit users were impacted.

Release v0.7.1 (2019-03-13)

- Additional fixes for UTF-8 BOM files which appear to happen more frequently with `local` files on Windows. This time some additional unit tests were added so hopefully there are few regressions in the future.
- Add the `ignore-missing` argument to *ksconf merge* to prevent errors when input files are absent. This allows bashisms `Some_App/{{default,local}}/savedsearches.conf` to work without errors if the local or default file is missing.
- Check for incorrect environment setup and suggest running sourcing `setSplunkEnv` to get a working environment. See #48 for more info.
- Minor improvements to some internal error handling, packaging, docs, and troubleshooting code.

Release v0.7.0 (2019-02-27)

Attention: For anyone who installed 0.6.x, we recommend a fresh install of the Splunk app due to packaging changes. This shouldn't be an issue in the future.

General changes:

- Added new *ksconf rest-publish* command that supersedes the use of `rest-export` for nearly every use case. Warning: No unit-testing has been created for this command yet, due to technical hurdles.
- Added *Cheat Sheet* to the docs.
- Massive doc cleanup of hundreds of typos and many expanded/clarified sections.
- Significant improvement to entrypoint handling and support for conditional inclusion of 3rd party libraries with sane behavior on import errors, and improved warnings. This information is conveniently viewable to the user via `ksconf --version`.
- Refactored internal diff logic and added additional safeties and unit tests. This includes improvements to TTY colorization which should avoid previous color leaks scenarios that were likely if unhandled exceptions occur.
- New support for metadata handling.
- CLI change for `rest-export`: The `--user` argument has been replaced with `--owner` to keep clean separation between the login account and object owners. (The old argument is still accept for now.)

Splunk app changes:

- Modified installation of python package installation. In previous releases, various `.dist-info` folders were created with version-specific names leading to a mismatch of package versions after upgrade. For this reason, we suggest that anyone who previously installed 0.6.x should do a fresh install.
- Changed Splunk app install script to `install.py` (it was `bootstrap_bin.py`). Hopefully this is more intuitive.
- Added Windows support to `install.py`.
- Now includes the Splunk Python SDK. Currently used for `rest-publish` but will eventually be used for additional functionally unique to the Splunk app.

3.12.2 Ksconf 0.6.x

Add deployment as a Splunk app for simplicity and significant docs cleanup.

Release v0.6.2 (2019-02-09)

- Massive rewrite and restructuring of the docs. Highlights include:
  - Reference material has been moved out of the user manual into a different top-level section.
  - Many new topics were added, such as
    - *Ksconf as external difftool*
    - *How Splunk writes to conf files*
    - *Configuration layers*
    - *What's so important about minimizing files?*
  - A new approach for CLI documentation. We're moving away from the **WALL OF TEXT** thing. (Yeah, it was really just the output from `--help`). That was limiting formatting, linking, and making the CLI output way too long.
- Refreshed Splunk app icons. Add missing alt icon.
- Several minor internal cleanups. Specifically the output of `--version` had a face lift.

Release v0.6.1 (2019-02-07)

- (Trivial) Fixed some small issues with the Splunk App (online AppInspect)

Release v0.6.0 (2019-02-06)

- Add initial support for building ksconf into a Splunk app.
  - App contains a local copy of the docs, helpful for anyone who's working offline.

  - Credit to Sarah Larson for the ksconf logos.
  - No `ksconf` functionality exposed to the Splunk UI at the moment.
- Docs/Sphinx improvements (more coming)
  - Begin work on cleaning up API docs.
  - Started converting various document pages into reStructuredText for greatly improved docs.
  - Improved PDF fonts and fixed a bunch of sphinx errors/warnings.
- Refactored the install docs into 2 parts. With the new ability to install ksconf as a Splunk app it's quite likely that most of the wonky corner cases will be less frequently needed, hence all the more exotic content was moved into the "Advanced Install Guide", tidying things up.

3.12.3 Ksconf 0.5.x

Add Python 3 support, new commands, support for external command plugins, tox and vagrant for testing.

Release v0.5.6 (2019-02-04)

- Fixes and improvements to the `filter` command. Found issue with processing from stdin, inconsistency in some CLI arguments, and finished implementation for various output modes.
- Add logo (fist attempt).

Release v0.5.5 (2019-01-28)

- New *ksconf filter* command added for slicing up a conf file into smaller pieces. Think of this as GREP that's stanza-aware. Can also allow or block attributes, if desirable.
- Expanded `rest-export` CLI capabilities to include a new `--delete` option, pretty-printing, and now supports stdin by allowing the user to explicitly set the file type using `--conf`.
- Refactored all CLI unittests for increased readability and long-term maintenance. Unit tests now can also be run individually as scripts from the command line.
- Minor tweaks to the `snapshot` output format, v0.2. This feature is still highly experimental.

Release v0.5.4 (2019-01-04)

- New commands added:
  - *ksconf snapshot* will dump a set of configuration files to a JSON formatted file. This can be used used for incremental "snapshotting" of running Splunk apps to track changes overtime.

  - *ksconf rest-export* builds a series of custom `curl` commands that can be used to publish or update stanzas on a remote instance without file system access. This can be helpful when pushing configs to Splunk Cloud when all you have is REST (splunkd) access. This command is indented for interactive admin not batch operations.

- Added the concept of command maturity. A listing is available by running `ksconf --version`
- Fix typo in `KSCONF_DEBUG`.
- Resolving some build issues.
- Improved support for development/testing environments using Vagrant (fixes) and Docker (new). Thanks to Lars Jonsson for these enhancements.

Release v0.5.3 (2018-11-02)

- Fixed bug where `ksconf combine` could incorrectly order directories on certain file systems (like ext4), effectively ignoring priorities. Repeated runs may resulted in undefined behavior. Solved by explicitly sorting input paths forcing processing to be done in lexicographical order.
- Fixed more issues with handling files with BOM encodings. BOMs and encodings in general are NOT preserved by ksconf. If this is an issue for you, please add an enhancement issue.
- Add Python 3.7 support
- Expand install docs specifically for offline mode and some OS-specific notes.
- Enable additional tracebacks for CLI debugging by setting `KSCONF_DEBUG=1` in the environment.

Release v0.5.2 (2018-08-13)

- Expand CLI output for `--help` and `--version`
- Internal cleanup of CLI entry point module name. Now the ksconf CLI can be invoked as `python -m ksconf`, you know, for anyone who's into that sort of thing.
- Minor docs and CI/testing improvements.

Release v0.5.1 (2018-06-28)

- Support external ksconf command plugins through custom *entry_points*, allowing for others to develop their own custom extensions as needed.
- Many internal changes: Refactoring of all CLI commands to use new entry_points as well as pave the way for future CLI unittest improvements.
- Docs cleanup / improvements.

Release v0.5.0 (2018-06-26)

- Python 3 support.
- Many bug fixes and improvements resulting from wider testing.

3.12.4 Ksconf 0.4.x

Ksconf 0.4.x switched to a modular code base, added build/release automation, PyPI package registration (installation via `pip install` and, online docs.

Release v0.4.10 (2018-06-26)

- Improve file handling to avoid "unclosed file" warnings. Impacted `parse_conf()`, `write_conf()`, and many unittest helpers.
- Update badges to report on the master branch only. (No need to highlight failures on feature or bug-fix branches.)

Release v0.4.9 (2018-06-05)

- Add some missing docs files

Release v0.4.8 (2018-06-05)

- Massive cleanup of docs: revamped install guide, added 'standalone' install procedure and developer-focused docs. Updated license handling.
- Updated docs configuration to dynamically pull in the ksconf version number.
- Using the classic 'read-the-docs' Sphinx theme.
- Added additional PyPi badges to README (GitHub home page).

Release v0.4.4-v0.4.1 (2018-06-04)

- Deployment and install fixes (It's difficult to troubleshoot/test without making a new release!)

Release v0.4.3 (2018-06-04)

- Rename PyPI package `kintyre-splunk-conf`
- Add support for building a standalone executable (zipapp).
- Revamp install docs and location
- Add GitHub release for the standalone executable.

Release v0.4.2 (2018-06-04)

- Add readthedocs.io support

Release v0.4.1 (2018-06-04)

- Enable PyPI production package building

Release v0.4.0 (2018-05-19)

- Refactor entire code base. Switched from monolithic all-in-one file to clean-cut modules.
- Versioning is now discoverable via `ksconf --version`, and controlled via git tags (via `git describe --tags`).

Module layout

- `ksconf.conf.*` - Configuration file parsing, writing, comparing, and so on
- `ksconf.util.*` - Various helper functions
- `ksconf.archive` - Support for uncompressing Splunk apps (tgz/zip files)
- `ksconf.vc.git` - Version control support. Git is the only VC tool supported for now. (Possibly ever)
- `ksconf.commands.<CMD>` - Modules for specific CLI functions. I may make this extendable, eventually.

3.12.5 Ksconf 0.3.x

First public releases.

Release v0.3.2 (2018-04-24)

- Add AppVeyor for Windows platform testing
- Add codecov integration
- Created ConfFileProxy.dump()

Release v0.3.1 (2018-04-21)

- Setup automation via Travis CI
- Add code coverage

Release v0.3.0 (2018-04-21)

- Switched to semantic versioning.
- 0.3.0 feels representative of the code maturity.

3.12.6 Ksconf legacy releases

Ksconf started in a private Kintyre repo. There are no official releases; all git history has been rewritten.

Release legacy-v1.0.1 (2018-04-20)

- Fixes to blocklist support and many enhancements to `ksconf unarchive`.
- Introduces parsing profiles.
- Lots of bug fixes to various subcommands.
- Added automatic detection of 'subcommands' for CLI documentation helper script.

Release legacy-v1.0.0 (2018-04-16)

- This is the first public release. First work began Nov 2017 (as a simple conf 'sort' tool, which was imported from yet another repo.) Version history was extracted/rewritten/preserved as much as possible.
- Mostly stable features.
- Unit test coverage over 85%
- Includes pre-commit hook configuration (so that other repos can use this to run `ksconf sort` and `ksconf check` against their conf files.

3.13 Known issues

3.13.1 General

- File encoding issues: Byte order markers and specific encodings are NOT preserved. All files are encoding using UTF-8 upon update, which is Splunk's expected encoding.

3.13.2 Splunk app

- File cleanup issues after *KSCONF app for Splunk* upgrades (impacts versions prior to 0.7.0). Old `.dist-info` folders or other stale files may be left around after upgrades. If you encounter this issue, either uninstall and delete the ksconf directory or manually remove the old 'bin'

folder and (re)upgrade to the latest version. The fix in 0.7.0 is to remove the version-specific portion of the folder name. (GH issue #37)

See more confirmed bugs in the issue tracker.

3.14 Advanced Installation Guide

The content in this document is a subsidiary to the *Installation Guide* because it became disorganized and the number of possible Python installation combinations and snags intensified. However, that culminated in the collection of excellent information that is provided here. Please remember, the Splunk app install approach was introduced to alleviate several of these issues.

A portion of this document is targeted at those who can't install packages as Admin or are forced to use Splunk's embedded Python. For everyone else, please start with the one-liner!

Tip: Do any of these words for phrases strike fear in your heart?

- `pip`
- `pipenv`
- `virtualenv`
- `wheel`
- `pyenv` (not the same as `pyvenv`)
- `python2.7` vs `python27` vs `py -27`
- `PYTHONPATH`
- `LD_LIBARY`
- RedHat Software Collections

If this list seems daunting, head over to *Install Splunk App*. There's no shame in it.

Contents

- *Advanced Installation Guide*
  - *Flowchart*
  - *Installation*
    - *Install from PyPI with PIP*
      - *Install ksconf into a virtual environment*
      - *Install ksconf system-wide*
    - *CentOS (RedHat derived) distros*
      - *RedHat Software Collections*
  - *Use the standalone executable*
    - *Install the Wheel manually (offline mode)*

        * *Install with Splunk's Python*
            - *On Linux or Mac*
            - *On Windows*
    - *Offline installation*
        * *Offline installation steps*
        * *Offline installation of pip*
            - *Use pip without installing it*
    - *Frequent gotchas*
        * *PIP Install TLS Error*
        * *No module named 'command.install'*
    - *Troubleshooting*
        * *Check Python version*
        * *Check PIP Version*
        * *Validate the install*
    - *Resources*

3.14.1 Flowchart

(Unfinished; more of a brainstorm at this point. . .)

- Is Python installed? (OS level)
  - Is the version greater than 2.7? (Some early 2.7 version have quarks, but typically this is okay)
  - If Python 3.x, is it greater than 3.4? (I'd like to drop 3.4, but lots of old distros still have it.)
- Do you have admin access? (root/Administrator; or can you get it? How hard? Will you need it each time you upgrade the ksconf?)
- Do you already have a large Python deployment or dependency? (If so, you'll probably be fine. Use virtualenv)
- Do you have any prior Python packaging or administration experience?
- Are you dealing with some vendor-specific solution?
  - Example: RedHat Software Collections – where they realize their software is way too old, so they try to make it possible to install newer version of things like Python, but since they aren't native or the default, you still end up jumping through a bunch of wonky hoops)
- Do you have Internet connectivity? (air gap or blocked outbound traffic, or proxy)

- Do you want to build/deploy your own ksconf extensions? If so, the Python package is a better option. (But at that point, you can probably already handle any packaging issues yourself.)

3.14.2 Installation

There are several ways to install ksconf. Technically, all standard Python packaging approaches should work just fine as there's no compiled code or external run-time dependencies so installation is fairly easy. However, for non-Python developers, there are some snags. Installation options are listed from the most easy and recommended, to more obscure and difficult:

Install from PyPI with PIP

The preferred installation method is to install via the standard Python package tool **pip**. Ksconf can be installed via the registered kintyre-splunk-conf package using the standard Python process.

There are 2 popular variations, depending on whether or not you would like to install for all users or test it locally.

Install ksconf into a virtual environment

Use this option if you don't have admin access

Installing `ksconf` with virtualenv is a great way to test the tool without requiring admin privileges and has many advantages for a production install. Here are the basic steps to get started.

Please change `venv` to a suitable path for your environment.

```
# Install Python virtualenv package (if not already installed)
pip install virtualenv

# Create and activte new 'venv' virtual environment
virtualenv venv
source venv/bin/activate

pip install kintyre-splunk-conf
```

Note: Windows users

The above virtual environment activation should be run as `venv\Scripts\activate.bat`.

Install ksconf system-wide

Important: This requires admin access.

This is the absolute easiest install method where 'ksconf' is available to all users on the system but it requires root access and `pip` must be installed and up-to-date.

On Mac or Linux, run:

```
sudo pip install kintyre-splunk-conf
```

On Windows, run this command from an Administrator console.

```
pip install kintyre-splunk-conf
```

CentOS (RedHat derived) distros

```
# Enable the EPEL repo so that `pip` can be installed.
sudo yum install -y epel-release

# Install pip
sudo yum install -y python-pip

# Install ksconf (globally, for all users)
sudo pip install kintyre-splunk-conf
```

RedHat Software Collections

The following assumes the `python27` software collection, but other version of Python are supported too. The initial setup and deployment of Software Collections is beyond the scope of this doc.

```
sudo scl enable python27 python -m pip install kintyre-splunk-conf
```

Hint: Missing pip?

If pip is missing from a RHSC, then install the following rpm.

```
yum install python27-python-pip
```

Unfortunately, the `ksconf` entrypoint script (in the `bin` folder) will not work correctly on it's own because it doesn't know about the scl environment, nor is it in the default PATH. To solve this, run the following:

```
sudo cat > /usr/local/bin/ksconf <<HERE
#!/bin/sh
source scl_source enable python27
exec /opt/rh/python27/root/usr/bin/ksconf "$@"
HERE
chmod +x /usr/local/bin/ksconf
```

3.14.3 Use the standalone executable

Deprecated since version 0.6.0: This option remains for historical reference and will likely be disabled in the future. If this seems like the best option to you, then please consider installing the KSCONF App for Splunk instead.

Ksconf can be installed as a standalone executable zip app. This approach still requires a Python interpreter to be present either from the OS or the one embedded with Splunk Enterprise. This works well for testing or when all other options fail.

From the GitHub releases page, grab the file name `ksconf-*.pyz`, download it, copy it to a `bin` folder in your PATH and rename it `ksconf`. The default shebang looks for 'python' in the PATH, but this can be adjusted as needed. Since installing with Splunk is a common use case, a second file named `ksconf-*-splunk.pyz` already has the shebang set for the standard `/opt/splunk` install path.

Typical embedded Splunk install example:

```
VER=0.5.0
curl https://github.com/Kintyre/ksconf/releases/download/v${VER}/ksconf-${VER}-splunk.pyz
mv ksconf-${VER}-splunk.pyz /opt/splunk/bin/
cd /opt/splunk/bin
ln -sf ksconf-${VER}-splunk.pyz ksconf
chmod +x ksconf
ksconf --version
```

Reasons why this is a non-ideal install approach:

- Lower performance since all Python files live in a zip file, and pre-compiled version's can be cached.
- No standard install pathway (doesn't use pip); user must manually copy the executable into place.
- Uses a non-standard build process. (May not be a big deal, but could cause things to break in the future.)

Install the Wheel manually (offline mode)

Download the latest "Wheel" file file from PyPI, copy it to the destination server and install with pip.

Offline pip install:

```
pip install ~/Downloads/kintyre-splunk-conf-0.4.2-py2.py3-none-any.whl
```

Install with Splunk's Python

Deprecated since version 0.6.0: Don't do this anymore. Please use the KSCONF App for Splunk instead.

Splunk Enterprise 6.x and later installs an embedded Python 2.7 environment. However, Splunk does not provide packing tools (such as `pip` or the `distutils` standard library which is required to bootstrap install `pip`). For these reasons, it's typically easier and cleaner to install `ksconf` with the system provided Python. However, sometimes the system-provided Python environment is the wrong version, is missing (like on Windows), or security restrictions prevent the installation of additional packages. In such cases, Splunk's embedded Python becomes a beacon of hope.

On Linux or Mac

Download the latest "Wheel" file file from PyPI. The path to this download will be set in the `pkg` variable as shown below.

Setup the shell:

```
export SPLUNK_HOME=/opt/splunk
export pkg=~/Downloads/kintyre_splunk_conf-0.4.9-py2.py3-none-any.whl
```

Run the following:

```
cd $SPLUNK_HOME
mkdir Kintyre
cd Kintyre
# Unzip the 'kconf' folder into SPLUNK_HOME/Kintyre
unzip "$pkg"

cat > $SPLUNK_HOME/bin/ksconf <<HERE
#!/bin/sh
export PYTHONPATH=$PYTHONPATH:$SPLUNK_HOME/Kintyre
exec $SPLUNK_HOME/bin/python -m ksconf \$*
HERE
chmod +x $SPLUNK_HOME/bin/ksconf
```

Test the install:

```
ksconf --version
```

On Windows

1. Open a browser and download the latest "Wheel" file file from PyPI.
2. Rename the `.whl` extension to `.zip`. (This may require showing file extensions in Explorer.)
3. Extract the zip file to a temporary folder. (This should create a folder named "ksconf")
4. Create a new folder called "Kintyre" under the Splunk installation path (aka `SPLUNK_HOME`) By default, this is `C:\Program Files\Splunk`.
5. Copy the "ksconf" folder to `%SPLUNK_HOME%\Kintyre`.
6. Create a new batch file called `ksconf.bat` and paste in the following. Be sure to adjust for a non-standard `%SPLUNK_HOME%` value, if necessary.

```
@echo off
SET SPLUNK_HOME=C:\Program Files\Splunk
SET PYTHONPATH=%SPLUNK_HOME%\bin;%SPLUNK_HOME%\Python-2.7\Lib\site-packages\win32;
↪%SPLUNK_HOME%\Python-2.7\Lib\site-packages;%SPLUNK_HOME%\Python-2.7\Lib
SET PYTHONPATH=%PYTHONPATH%;%SPLUNK_HOME%\Kintyre
CALL "%SPLUNK_HOME%\bin\python.exe" -m ksconf %*
```

7. Move `ksconf.bat` to the `Splunk\bin` folder. (This assumes that `%SPLUNK_HOME%/bin` is part of your `%PATH%`. If not, add it, or find an appropriate install location.)
8. Test this by running `ksconf --version` from the command line.

3.14.4 Offline installation

Installing ksconf to an offline or network restricted computer requires three steps: (1) download the latest packages from the Internet to a staging location, (2) transfer the staged content (often as a zip file) to the restricted host, and (3) use pip to install packages from the staged copy. Fortunately, pip makes offline workflows quite easy to achieve. Pip can download a Python package with all dependencies stored as wheels files into a single directory, and pip can be told to install from that directory instead of attempting to talk to the Internet.

The process of transferring these files is very organization-specific. The example below shows the creation of a tarball (since `tar` is universally available on Unix systems), but any acceptable method is fine. If security is a high concern, this step is frequently where safety checks are implemented: such as, antivirus scans, static code analysis, manual inspection, and/or comparison of cryptographic file hashes.

One additional use-case for this workflow, is to ensure the exact same version of all packages are deployed consistently across all servers and environments. Often, building a `requirements.txt` file with `pip freeze`, is a more appropriate solution. Alternatively, consider using `pipenv lock` for even more security benefits.

Offline installation steps

Important: Pip must be installed on the destination server for this process to work. If pip is NOT installed, see the *Offline installation of pip* section below.

Step 1: Use pip to download the latest package and their dependencies. Be sure to use the same version of Python that is running on destination machine.

```
# download packages
python2.7 -m pip download -d ksconf-packages kintyre-splunk-conf
```

A new directory named 'ksconf-packages' will be created and will contain the necessary `*.whl` files.

Step 2: Transfer the directory or archive to the remote computer. Insert whatever security and file copy procedures necessary for your organization.

```
# Compress directory (on staging computer)
tar -czvf ksconf-packages.tgz ksconf-packages

# Copy file using whatever means
scp ksconf-packages.tgz user@server:/tmp/ksconf-packages.tgz

# Extract the archive (on destination server)
tar -xzvf ksconf-packages.tgz
```

Step 3:

```
# Install ksconf package with pip
pip install --no-index --find-links=ksconf-packages kntyre-splunk-conf

# Test the installation
ksconf --version
```

The `ksconf-packages` folder can now be safely removed.

Offline installation of pip

Use the recommended `pip` install procedures listed elsewhere if possible. But if a remote bootstrap of pip is your only option, then here are the steps. (This process mirrors the steps above and can be combined, if needed.)

Step 1: Fetch bootstrap script and necessary wheels

```
mkdir ksconf-packages
curl https://bootstrap.pypa.io/get-pip.py -o ksconf-packages/get-pip.py
python2.7 -m pip download -d /tmp/my_packages pip setuptools wheel
```

The `ksconf-pacakges` folder should contain 1 script, and 3 wheel (`*.whl`) files.

Step 2: Archive and/or copy to offline server

Step 3: Bootstrap pip

```
sudo python get-pip.py --no-index --find-links=ksconf-packages/

# Test with
pip --version
```

Use pip without installing it

If you have a copy of the `pip*.whl` (wheel) file, then it can be executed directly by Python. This can be used to run `pip` without actually installing it, or for installing pip initially (bypassing the `get-pip.py` script step noted above.)

Here's an example of how this could work:

Step 1: Download the pip wheel on a machine where `pip` works, by running:

```
pip download pip -d .
```

This will create a file like `pip-19.0.1-py2.py3-none-any.whl` in the current working directory.

Step 2: Copy the pip wheel to another machine (likely where pip isn't installed.)

Step 3: Execute the wheel by running:

```
python pip-19.0.1-py2.py3-none-any.whl/pip list
```

Substitute the `list` command with whatever action you need (like `install` or whatever).

3.14.5 Frequent gotchas

PIP Install TLS Error

If `pip` throws an error message like the following:

```
There was a problem confirming the ssl certificate: [SSL: TLSV1_ALERT_PROTOCOL_VERSION]␣
↪tlsv1 alert protocol version
...
No matching distribution found for setuptools
```

The problem is likely caused by changes to PyPI website in April 2018 when support for TLS v1.0 and 1.1 were removed. Downloading new packages requires upgrading to a new version of pip. Like so:

Upgrade pip as follows:

```
curl https://bootstrap.pypa.io/get-pip.py | python
```

Note: Use `sudo python` above if not in a virtual environment.

Helpful links:

- Not able to install Python packages [SSL: TLSV1_ALERT_PROTOCOL_VERSION]
- 'pip install' fails for every package ("Could not find a version that satisfies the requirement")

No module named 'command.install'

If, while trying to install `pip` or run a `pip` command you see the following error:

```
ImportError: No module named command.install
```

Likely this is because you are using a crippled version of Python; like the one that ships with Splunk. This won't work. Either get a pre-packaged version (the `.pyz` file) or install using the OS-level Python.

3.14.6 Troubleshooting

Here are a few fact gathering type commands that may help you begin to track down problems.

Check Python version

Check your installed Python version by running:

```
python --version
```

Note that Linux distributions and Mac OS X that ship with multiple versions of Python may have renamed this to `python2`, `python2.7` or similar.

Check PIP Version

```
pip --version
```

If you are running a different Python interpreter version, you can instead run this as:

```
python2.7 -m pip --version
```

Validate the install

Confirm installation with the following command:

```
ksconf --version
```

If this works, it means that `ksconf` installed and is part of your `PATH` and should be useable everywhere in your system. Go forth and conquer!

If this doesn't work, here are a few things to try:

1. Check that your `PATH` is set correctly.
2. Try running ksconf as a "module" (sometimes works around a PATH issue). Run `python -m ksconf`
3. If you're running the Splunk app, try running the following:

   ```
   cd $SPLUNK_HOME/etc/apps/ksconf/bin/lib
   python -m ksconf --version
   ```

   If this works, then the issue has something to do with your path.

It may be helpful to uninstall (remove) the Splunk app and reinstall from scratch.

3.14.7 Resources

- Python packaging docs provide a general overview on installing Python packages, how to install per-user vs install system-wide.
- Install PIP docs explain how to bootstrap or upgrade `pip` the Python packaging tool. Recent versions of Python come with this by default, but releases before Python 2.7.9 do not.

3.15 License

```
                                 Apache License
                           Version 2.0, January 2004
                        http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

   "License" shall mean the terms and conditions for use, reproduction,
   and distribution as defined by Sections 1 through 9 of this document.

   "Licensor" shall mean the copyright owner or entity authorized by
   the copyright owner that is granting the License.

   "Legal Entity" shall mean the union of the acting entity and all
   other entities that control, are controlled by, or are under common
   control with that entity. For the purposes of this definition,
   "control" means (i) the power, direct or indirect, to cause the
   direction or management of such entity, whether by contract or
   otherwise, or (ii) ownership of fifty percent (50%) or more of the
   outstanding shares, or (iii) beneficial ownership of such entity.

   "You" (or "Your") shall mean an individual or Legal Entity
   exercising permissions granted by this License.

   "Source" form shall mean the preferred form for making modifications,
   including but not limited to software source code, documentation
   source, and configuration files.

   "Object" form shall mean any form resulting from mechanical
   transformation or translation of a Source form, including but
   not limited to compiled object code, generated documentation,
   and conversions to other media types.

   "Work" shall mean the work of authorship, whether in Source or
   Object form, made available under the License, as indicated by a
   copyright notice that is included in or attached to the work
   (an example is provided in the Appendix below).

   "Derivative Works" shall mean any work, whether in Source or Object
```

(continues on next page)

(continued from previous page)

```
   form, that is based on (or derived from) the Work and for which the
   editorial revisions, annotations, elaborations, or other modifications
   represent, as a whole, an original work of authorship. For the purposes
   of this License, Derivative Works shall not include works that remain
   separable from, or merely link (or bind by name) to the interfaces of,
   the Work and Derivative Works thereof.

   "Contribution" shall mean any work of authorship, including
   the original version of the Work and any modifications or additions
   to that Work or Derivative Works thereof, that is intentionally
   submitted to Licensor for inclusion in the Work by the copyright owner
   or by an individual or Legal Entity authorized to submit on behalf of
   the copyright owner. For the purposes of this definition, "submitted"
   means any form of electronic, verbal, or written communication sent
   to the Licensor or its representatives, including but not limited to
   communication on electronic mailing lists, source code control systems,
   and issue tracking systems that are managed by, or on behalf of, the
   Licensor for the purpose of discussing and improving the Work, but
   excluding communication that is conspicuously marked or otherwise
   designated in writing by the copyright owner as "Not a Contribution."

   "Contributor" shall mean Licensor and any individual or Legal Entity
   on behalf of whom a Contribution has been received by Licensor and
   subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of
   this License, each Contributor hereby grants to You a perpetual,
   worldwide, non-exclusive, no-charge, royalty-free, irrevocable
   copyright license to reproduce, prepare Derivative Works of,
   publicly display, publicly perform, sublicense, and distribute the
   Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of
   this License, each Contributor hereby grants to You a perpetual,
   worldwide, non-exclusive, no-charge, royalty-free, irrevocable
   (except as stated in this section) patent license to make, have made,
   use, offer to sell, sell, import, and otherwise transfer the Work,
   where such license applies only to those patent claims licensable
   by such Contributor that are necessarily infringed by their
   Contribution(s) alone or by combination of their Contribution(s)
   with the Work to which such Contribution(s) was submitted. If You
   institute patent litigation against any entity (including a
   cross-claim or counterclaim in a lawsuit) alleging that the Work
   or a Contribution incorporated within the Work constitutes direct
   or contributory patent infringement, then any patent licenses
   granted to You under this License for that Work shall terminate
   as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the
   Work or Derivative Works thereof in any medium, with or without
   modifications, and in Source or Object form, provided that You
   meet the following conditions:
```

(continues on next page)

(continued from previous page)

```
    (a) You must give any other recipients of the Work or
        Derivative Works a copy of this License; and

    (b) You must cause any modified files to carry prominent notices
        stating that You changed the files; and

    (c) You must retain, in the Source form of any Derivative Works
        that You distribute, all copyright, patent, trademark, and
        attribution notices from the Source form of the Work,
        excluding those notices that do not pertain to any part of
        the Derivative Works; and

    (d) If the Work includes a "NOTICE" text file as part of its
        distribution, then any Derivative Works that You distribute must
        include a readable copy of the attribution notices contained
        within such NOTICE file, excluding those notices that do not
        pertain to any part of the Derivative Works, in at least one
        of the following places: within a NOTICE text file distributed
        as part of the Derivative Works; within the Source form or
        documentation, if provided along with the Derivative Works; or,
        within a display generated by the Derivative Works, if and
        wherever such third-party notices normally appear. The contents
        of the NOTICE file are for informational purposes only and
        do not modify the License. You may add Your own attribution
        notices within Derivative Works that You distribute, alongside
        or as an addendum to the NOTICE text from the Work, provided
        that such additional attribution notices cannot be construed
        as modifying the License.

    You may add Your own copyright statement to Your modifications and
    may provide additional or different license terms and conditions
    for use, reproduction, or distribution of Your modifications, or
    for any such Derivative Works as a whole, provided Your use,
    reproduction, and distribution of the Work otherwise complies with
    the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise,
   any Contribution intentionally submitted for inclusion in the Work
   by You to the Licensor shall be under the terms and conditions of
   this License, without any additional terms or conditions.
   Notwithstanding the above, nothing herein shall supersede or modify
   the terms of any separate license agreement you may have executed
   with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade
   names, trademarks, service marks, or product names of the Licensor,
   except as required for reasonable and customary use in describing the
   origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or
   agreed to in writing, Licensor provides the Work (and each
```

(continues on next page)

(continued from previous page)

```
   Contributor provides its Contributions) on an "AS IS" BASIS,
   WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
   implied, including, without limitation, any warranties or conditions
   of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A
   PARTICULAR PURPOSE. You are solely responsible for determining the
   appropriateness of using or redistributing the Work and assume any
   risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory,
   whether in tort (including negligence), contract, or otherwise,
   unless required by applicable law (such as deliberate and grossly
   negligent acts) or agreed to in writing, shall any Contributor be
   liable to You for damages, including any direct, indirect, special,
   incidental, or consequential damages of any character arising as a
   result of this License or out of the use or inability to use the
   Work (including but not limited to damages for loss of goodwill,
   work stoppage, computer failure or malfunction, or any and all
   other commercial damages or losses), even if such Contributor
   has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing
   the Work or Derivative Works thereof, You may choose to offer,
   and charge a fee for, acceptance of support, warranty, indemnity,
   or other liability obligations and/or rights consistent with this
   License. However, in accepting such obligations, You may act only
   on Your own behalf and on Your sole responsibility, not on behalf
   of any other Contributor, and only if You agree to indemnify,
   defend, and hold each Contributor harmless for any liability
   incurred by, or claims asserted against, such Contributor by reason
   of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

Copyright 2019 Kintyre

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

    http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.
```

3.16 API Reference

Note: As of now, no assumptions should be made about APIs remaining stable

KSCONF is first and foremost a CLI tool, so backwards incompatible changes are more of a concern for CLI breakage than for API breakage. That being said, there are a number of helpful features in the core `ksconf` Python module. So if anyone is interested in using the API, please feel free to do so, but let us know *how* you are using it and we'll find a way to keep the the important bits stable. We'd love to make it more clear what APIs are stable and which are likely to change.

As of right now, the general rule of thumb is this: Anything well-covered by the unit tests should be be fairly safe to build on top of, but again, *ping us*.

3.16.1 ksconf

Subpackages

ksconf.commands package

Module contents

class ksconf.commands.**KsconfCmd**(*name*)

Bases: `object`

Ksconf command specification base class.

add_parser(*subparser*)

description = None

format = 'default'

help = None

launch(*args*)

Handle flow control between pre_run() / run() / post_run()

maturity = 'alpha'

parse_conf(*path*, *mode='r'*, *profile=None*, *raw_exec=False*)

post_run(*args*, *exec_info=None*)

Any custom clean up work that needs done. Always called if run() was. Presence of exc_info indicates failure.

pre_run(*args*)

Pre-run hook. Any exceptions here prevent run() from being called.

register_args(*parser*)

This function in passed the

run(*args*)

Actual works happens here. Return code should be an EXIT_CODE_* from consts.

version_extra = None

class ksconf.commands.**ConfDirProxy**(*name*, *mode*, *parse_profile=None*)

Bases: object

get_file(*relpath*)

class ksconf.commands.**ConfFileProxy**(*name*, *mode*, *stream=None*, *parse_profile=None*, *is_file=None*)

Bases: object

close()

data

dump(*data*, ***kwargs*)

exists()

is_file()

load(*profile=None*)

readable()

reset()

set_parser_option(***kwargs*)

Setting a key to None will remove that setting.

stream

unlink()

writable()

class ksconf.commands.**ConfFileType**(*mode='r'*, *action='open'*, *parse_profile=None*, *accept_dir=False*)

Bases: object

Factory for creating conf file object types; returns a lazy-loader ConfFile proxy class

Started from argparse.FileType() and then changed everything. With our use case, it's often necessary to delay writing, or read before writing to a conf file (depending on whether or not –dry-run mode is enabled, for example.)

Instances of FileType are typically passed as type= arguments to the ArgumentParser add_argument() method.

Parameters

- **mode** (str) – How the file is to be opened. Accepts "r", "w", and "r+".
- **action** (str) – Determine how much work should be handled during argument parsing vs handed off to the caller. Supports 'none', 'open', 'load'. Full descriptions below.
- **parse_profile** – parsing configuration settings passed along to the parser

- **accept_dir** (bool) – Should the CLI accept a directory of config files instead of an individual file. Defaults to *False*.

Values for action

Action	Description
none	No preparation or testing is done on the filename.
open	Ensure the file exists and can be opened.
load	Ensure the file can be opened and parsed successfully.

Once invoked, instances of this class will return a ConfFileProxy object, or a ConfDirProxy object if a directory is passed in via the CLI.

ksconf.commands.**dedent**(*text*)
Remove any common leading whitespace from every line in *text*.

This can be used to make triple-quoted strings line up with the left edge of the display, while still presenting them in the source code in indented form.

Note that tabs and spaces are both treated as whitespace, but they are not equal: the lines " hello" and "thello" are considered to have no common leading whitespace. (This behaviour is new in Python 2.5; older versions of this module incorrectly expanded tabs before searching for common leading whitespace.)

ksconf.commands.**get_all_ksconf_cmds**(*on_error='warn'*)

ksconf.commands.**get_entrypoints**

ksconf.commands.**add_splunkd_access_args**(*parser*)

ksconf.commands.**add_splunkd_namespace**(*parser*)

ksconf.conf package

Submodules

ksconf.conf.delta module

class ksconf.conf.delta.**DiffGlobal**(*type*)
Bases: tuple

type
Alias for field number 0

class ksconf.conf.delta.**DiffHeader**(*name*, *mtime=None*)
Bases: object

detect_mtime()

class ksconf.conf.delta.**DiffOp**(*tag*, *location*, *a*, *b*)
Bases: tuple

a
Alias for field number 2

b
Alias for field number 3

location
Alias for field number 1

tag
Alias for field number 0

class ksconf.conf.delta.**DiffStanza**(*type*, *stanza*)
Bases: tuple

stanza
Alias for field number 1

type
Alias for field number 0

class ksconf.conf.delta.**DiffStzKey**(*type*, *stanza*, *key*)
Bases: tuple

key
Alias for field number 2

stanza
Alias for field number 1

type
Alias for field number 0

ksconf.conf.delta.**compare_cfgs**(*a*, *b*, *allow_level0=True*)
Return list of 5-tuples describing how to turn a into b.

Note: The *Opcode* tags borrowed from SequenceMatcher class in the difflib standard Python module.

Each tuple takes the form:

(tag, location, a, b)

tag:

Value	Meaning
'replace'	same element in both, but different values.
'delete'	remove value b
'insert'	insert value a
'equal'	same values in both

location is a tuple that can take the following forms:

Tuple form	Description
(0)	Global file level context (e.g., both files are the same)
(1, stanza)	Stanzas are the same, or completely different (no shared keys)
(2, stanza, key)	Key level, indicating

Possible alternatives:

https://dictdiffer.readthedocs.io/en/latest/#dictdiffer.patch

ksconf.conf.delta.**compare_stanzas**(*a*, *b*, *stanza_name*)

ksconf.conf.delta.**is_equal**(*delta*)
Is the delta output show that the compared objects are identical

ksconf.conf.delta.**reduce_stanza**(*stanza*, *keep_attrs*)
Pre-process a stanzas so that only a common set of keys will be compared. :param stanza: Stanzas containing attributes and values :type stanza: dict :param keep_attrs: Listing of :type keep_attrs: (list, set, tuple, dict) :return: a reduced copy of `stanza`.

ksconf.conf.delta.**show_diff**(*stream*, *diffs*, *headers=None*)

ksconf.conf.delta.**show_text_diff**(*stream*, *a*, *b*)

ksconf.conf.delta.**summarize_cfg_diffs**(*delta*, *stream*)
Summarize a delta into a human readable format. The input *delta* is in the format produced by the compare_cfgs() function.

ksconf.conf.merge module

ksconf.conf.merge.**merge_conf_dicts**(**dicts*)

ksconf.conf.merge.**merge_conf_files**(*dest*, *configs*, *dry_run=False*, *banner_comment=None*)

ksconf.conf.parser module

Parse and write Splunk's .conf files

According to this doc:

https://docs.splunk.com/Documentation/Splunk/7.2.3/Admin/Howtoeditaconfigurationfile

1. Comments must start at the beginning of a line (#)
2. Comments may not be after a stanza name or on an attribute's value
3. Supporting encoding is UTF-8 (and therefore ASCII too)

exception ksconf.conf.parser.**ConfParserException**
Bases: `Exception`

exception ksconf.conf.parser.**DuplicateKeyException**
Bases: ksconf.conf.parser.ConfParserException

exception ksconf.conf.parser.**DuplicateStanzaException**
Bases: ksconf.conf.parser.ConfParserException

class ksconf.conf.parser.**Token**
Bases: object

Immutable token object. deepcopy returns the same object

ksconf.conf.parser.**conf_attr_boolean**(*value*)

ksconf.conf.parser.**cont_handler**(*iterable*, *continue_re=re.compile('^(.*)\\\\$')*, *breaker='\n'*)
Look for trailing backslashes ("\") which indicate a value for an attribute is split across multiple lines. This function will group such lines together, and pass all other lines through as-is. Note that the continuation character must be the very last character on the line, trailing whitespace is not allowed.

Parameters

- **iterable** (iter) – lines from a configuration file
- **continue_re** – regular expression to detect the continuation character
- **breaker** – joining string when combining continued lines into a single string. Default '\n'

Returns lines of text

Return type str

ksconf.conf.parser.**detect_by_bom**(*path*)

ksconf.conf.parser.**inject_section_comments**(*section*, *prepend=None*, *append=None*)

ksconf.conf.parser.**parse_conf**(*stream*, *profile={'dup_key': 'overwrite', 'dup_stanza': 'exception', 'keep_comments': True, 'strict': True}*, *encoding=None*)
Parse a .conf file. This is a wrapper around parse_conf_stream() that allows filenames or stream to be passed in.

Parameters

- **stream** (str, file) – the path to a configuration file or open file-like object to be parsed
- **profile** – parsing configuration settings
- **encoding** – Defaults to the system default, "uft-8"

Returns a mapping of the stanza and attributes. The resulting output is accessible as [stanaza][attribute] -> value

Return type dict

ksconf.conf.parser.**parse_conf_stream**(*stream*, *keys_lower=False*, *handle_conts=True*, *keep_comments=False*, *dup_stanza='exception'*, *dup_key='overwrite'*, *strict=False*)

ksconf.conf.parser.**section_reader**(*stream*, *section_re=re.compile('^ [\\s\\t]*\\[(.*)\\]\\s*$')*)

This generator break a configuration file stream into sections. Each section contains a name and a list of text lines held within that section.

Sections that have no entries must be preserved. Any lines before the first section are send back with the section name of None.

Parameters

- **stream** (file) – configuration file input stream
- **section_re** – regular expression for detecting stanza headers

Returns sections in the form of *(section_name, lines_of_text)*

Return type tuple

ksconf.conf.parser.**smart_write_conf**(*filename*, *conf*, *stanza_delim='\n'*, *sort=True*, *temp_suffix='.tmp'*)

ksconf.conf.parser.**splitup_kvpairs**(*lines*, *comments_re=re.compile('^ \\s*[#;]')*, *keep_comments=False*, *strict=False*)

Break up 'attribute=value' entries in a configuration file.

Parameters

- **lines** (iter) – the body of a stanza containing associated attributes and values
- **comments_re** – Regular expression used to detect comments.
- **keep_comments** (bool, optional) – Should comments be preserved in the output. Defaults to *False*.
- **strict** (bool, optional) – Should unknown content in the stanza stop processing. Defaults to *False* allowing "junk" to be silently ignored allowing for a best-effort parse.

Returns iterable of (attribute,value) tuples

ksconf.conf.parser.**write_conf**(*stream*, *conf*, *stanza_delim='\n'*, *sort=True*)

ksconf.conf.parser.**write_conf_stream**(*stream*, *conf*, *stanza_delim='\n'*, *sort=True*)

Module contents

ksconf.util package

Submodules

ksconf.util.compare module

ksconf.util.compare.**file_compare**(*fn1*, *fn2*)

ksconf.util.compare.**fileobj_compare**(*f1*, *f2*)

ksconf.util.completers module

ksconf.util.completers.**DirectoriesCompleter**(**args*, ***kwargs*)

ksconf.util.completers.**FilesCompleter**(**args*, ***kwargs*)

ksconf.util.completers.**autocomplete**(**args*, ***kwargs*)

ksconf.util.file module

class ksconf.util.file.**ReluctantWriter**(*path*, **args*, ***kwargs*)
Bases: object

Context manager to intelligently handle updates to an existing file. New content is written to a temp file, and then compared to the current file's content. The file file will be overwritten only if the contents changed.

ksconf.util.file.**dir_exists**(*directory*)
Ensure that the directory exists

ksconf.util.file.**expand_glob_list**(*iterable*)

ksconf.util.file.**file_fingerprint**(*path*, *compare_to=None*)

ksconf.util.file.**file_hash**(*path*, *algorithm='sha256'*)

ksconf.util.file.**match_bwlist**(*value*, *bwlist*, *escape=True*)

ksconf.util.file.**relwalk**(*top*, *topdown=True*, *onerror=None*, *followlinks=False*)
Relative path walker Like os.walk() except that it doesn't include the "top" prefix in the resulting 'dirpath'.

ksconf.util.file.**smart_copy**(*src*, *dest*)
Copy (overwrite) file only if the contents have changed.

ksconf.util.file.**which**(*cmd*, *mode=1*, *path=None*)
Given a command, mode, and a PATH string, return the path which conforms to the given mode on the PATH, or None if there is no such file. *mode* defaults to os.F_OK | os.X_OK. *path* defaults to the result of os.environ.get("PATH"), or can be overridden with a custom search path.

ksconf.util.rest module

ksconf.util.rest.**build_rest_namespace**(*base*, *owner=None*, *app=None*)

ksconf.util.rest.**build_rest_url**(*base*, *service*, *owner=None*, *app=None*)

ksconf.util.terminal module

class ksconf.util.terminal.**TermColor**(*stream*)
Bases: object

Simple color setting helper class that's a context manager wrapper around a stream. This ensure that the color is always reset at the end of a session.

color(**codes*)

reset()

write(*content*)

ksconf.util.terminal.**tty_color**(*stream*, **codes*)

Module contents

ksconf.util.**debug_traceback**()
If the 'KSCONF_DEBUG' environmental variable is set, then show a stack trace.

ksconf.vc package

Submodules

ksconf.vc.git module

class ksconf.vc.git.**GitCmdOutput**(*cmd*, *returncode*, *stdout*, *stderr*, *lines*)
Bases: tuple

cmd
Alias for field number 0

lines
Alias for field number 4

returncode
Alias for field number 1

stderr
Alias for field number 3

stdout
Alias for field number 2

exception ksconf.vc.git.**GitNotAvailable**
Bases: Exception

ksconf.vc.git.**git_cmd**(*args*, *shell=False*, *cwd=None*, *capture_std=True*, *encoding='utf-8'*)

ksconf.vc.git.**git_cmd_iterable**(*args*, *iterable*, *cwd=None*, *cmd_len=1024*)

ksconf.vc.git.**git_is_clean**(*path=None*, *check_untracked=True*, *check_ignored=False*)

ksconf.vc.git.**git_is_working_tree**(*path=None*)

ksconf.vc.git.**git_ls_files**(*path*, **modifiers*)

ksconf.vc.git.**git_status_summary**(*path*)

ksconf.vc.git.**git_status_ui**(*path*, **args*)

ksconf.vc.git.**git_version**()

Module contents

Submodules

ksconf.archive module

ksconf.archive.**GenArchFile**
alias of ksconf.archive.GenericArchiveEntry

ksconf.archive.**extract_archive**(*archive_name*, *extract_filter=None*)

ksconf.archive.**gaf_filter_name_like**(*pattern*)

ksconf.archive.**gen_arch_file_remapper**(*iterable*, *mapping*)

ksconf.archive.**sanity_checker**(*iterable*)

ksconf.consts module

ksconf.setup_entrypoints module

Defines all command prompt entry points for CLI actions

This is a silly hack that serves 2 purposes:

(1) It works around an apparent Python 3.4/3.5 bug on Windows where [options.entry_point] in setup.cfg is ignored hence 'ksconf' isn't installed as a console script and custom ksconf_* entry points are not available. (So no CLI commands are available)

(2) **It allows for fallback mechanism when**

  (a) running unit tests (can happen before install)

  (b) if entrypoints or pkg_resources are not available at run time (Splunk's embedded python)

class ksconf.setup_entrypoints.**Ep**(*name*, *module_name*, *object_name*)
Bases: tuple

module_name
Alias for field number 1

name
Alias for field number 0

object_name
Alias for field number 2

class ksconf.setup_entrypoints.**LocalEntryPoint**(*data*)
Bases: object

Bare minimum stand-in for entrypoints.EntryPoint

load()

ksconf.setup_entrypoints.**debug**()

ksconf.setup_entrypoints.**get_entrypoints_fallback**(*group*)

ksconf.setup_entrypoints.**get_entrypoints_setup**()

Module contents

ksconf - Kintyre Splunk CONFig tool

Design goals:

- Multi-purpose go-to .conf tool.
- Dependability
- Simplicity
- No eternal dependencies (single source file, if possible; or packable as single file.)
- Stable CLI
- Good scripting interface for deployment scripts and/or git hooks

exception ksconf.**KsconfPluginWarning**
Bases: Warning

CHAPTER 4

Indices and tables

- genindex
- modindex
- search

Bibliography

[SPLKDOC1] https://docs.splunk.com/Documentation/Splunk/7.2.3/Admin/Configurationfiledirectories

Python Module Index

k

ksconf, 100
ksconf.archive, 99
ksconf.commands, 90
ksconf.conf, 96
ksconf.conf.delta, 92
ksconf.conf.merge, 94
ksconf.conf.parser, 94
ksconf.consts, 99
ksconf.setup_entrypoints, 99
ksconf.util, 98
ksconf.util.compare, 97
ksconf.util.completers, 97
ksconf.util.file, 97
ksconf.util.rest, 97
ksconf.util.terminal, 98
ksconf.vc, 99
ksconf.vc.git, 98

Index

A

a (*ksconf.conf.delta.DiffOp attribute*), 92
add_parser() (*ksconf.commands.KsconfCmd method*), 90
add_splunkd_access_args() (*in module ksconf.commands*), 92
add_splunkd_namespace() (*in module ksconf.commands*), 92
autocomplete() (*in module ksconf.util.completers*), 97

B

b (*ksconf.conf.delta.DiffOp attribute*), 93
build_rest_namespace() (*in module ksconf.util.rest*), 97
build_rest_url() (*in module ksconf.util.rest*), 97

C

close() (*ksconf.commands.ConfFileProxy method*), 91
cmd (*ksconf.vc.git.GitCmdOutput attribute*), 98
color() (*ksconf.util.terminal.TermColor method*), 98
compare_cfgs() (*in module ksconf.conf.delta*), 93
compare_stanzas() (*in module ksconf.conf.delta*), 94
conf_attr_boolean() (*in module ksconf.conf.parser*), 95
ConfDirProxy (*class in ksconf.commands*), 91
ConfFileProxy (*class in ksconf.commands*), 91
ConfFileType (*class in ksconf.commands*), 91
ConfParserException, 94
cont_handler() (*in module ksconf.conf.parser*), 95

D

data (*ksconf.commands.ConfFileProxy attribute*), 91
debug() (*in module ksconf.setup_entrypoints*), 100
debug_traceback() (*in module ksconf.util*), 98
dedent() (*in module ksconf.commands*), 92
description (*ksconf.commands.KsconfCmd attribute*), 90
detect_by_bom() (*in module ksconf.conf.parser*), 95
detect_mtime() (*ksconf.conf.delta.DiffHeader method*), 92
DiffGlobal (*class in ksconf.conf.delta*), 92
DiffHeader (*class in ksconf.conf.delta*), 92
DiffOp (*class in ksconf.conf.delta*), 92
DiffStanza (*class in ksconf.conf.delta*), 93
DiffStzKey (*class in ksconf.conf.delta*), 93
dir_exists() (*in module ksconf.util.file*), 97
DirectoriesCompleter() (*in module ksconf.util.completers*), 97
dump() (*ksconf.commands.ConfFileProxy method*), 91
DuplicateKeyException, 94
DuplicateStanzaException, 95

E

Ep (*class in ksconf.setup_entrypoints*), 99
exists() (*ksconf.commands.ConfFileProxy method*), 91
expand_glob_list() (*in module ksconf.util.file*), 97
extract_archive() (*in module ksconf.archive*), 99

F

file_compare() (*in module ksconf.util.compare*), 97
file_fingerprint() (*in module ksconf.util.file*), 97
file_hash() (*in module ksconf.util.file*), 97
fileobj_compare() (*in module ksconf.util.compare*), 97
FilesCompleter() (*in module ksconf.util.completers*), 97
format (*ksconf.commands.KsconfCmd attribute*), 90

G

gaf_filter_name_like() (*in module ksconf.archive*), 99
gen_arch_file_remapper() (*in module ksconf.archive*), 99
GenArchFile (*in module ksconf.archive*), 99
get_all_ksconf_cmds() (*in module ksconf.commands*), 92
get_entrypoints (*in module ksconf.commands*), 92
get_entrypoints_fallback() (*in module ksconf.setup_entrypoints*), 100
get_entrypoints_setup() (*in module ksconf.setup_entrypoints*), 100
get_file() (*ksconf.commands.ConfDirProxy method*), 91
git_cmd() (*in module ksconf.vc.git*), 98
git_cmd_iterable() (*in module ksconf.vc.git*), 99
git_is_clean() (*in module ksconf.vc.git*), 99
git_is_working_tree() (*in module ksconf.vc.git*), 99
git_ls_files() (*in module ksconf.vc.git*), 99
git_status_summary() (*in module ksconf.vc.git*), 99
git_status_ui() (*in module ksconf.vc.git*), 99
git_version() (*in module ksconf.vc.git*), 99
GitCmdOutput (*class in ksconf.vc.git*), 98
GitNotAvailable, 98

H

help (*ksconf.commands.KsconfCmd attribute*), 90

I

inject_section_comments() (*in module ksconf.conf.parser*), 95
is_equal() (*in module ksconf.conf.delta*), 94
is_file() (*ksconf.commands.ConfFileProxy method*), 91

K

key (*ksconf.conf.delta.DiffStzKey attribute*), 93
ksconf (*module*), 100
ksconf.archive (*module*), 99
ksconf.commands (*module*), 90
ksconf.conf (*module*), 96
ksconf.conf.delta (*module*), 92
ksconf.conf.merge (*module*), 94
ksconf.conf.parser (*module*), 94
ksconf.consts (*module*), 99
ksconf.setup_entrypoints (*module*), 99
ksconf.util (*module*), 98
ksconf.util.compare (*module*), 97
ksconf.util.completers (*module*), 97
ksconf.util.file (*module*), 97
ksconf.util.rest (*module*), 97
ksconf.util.terminal (*module*), 98
ksconf.vc (*module*), 99
ksconf.vc.git (*module*), 98
KsconfCmd (*class in ksconf.commands*), 90
KsconfPluginWarning, 100

L

launch() (*ksconf.commands.KsconfCmd method*), 90
lines (*ksconf.vc.git.GitCmdOutput attribute*), 98
load() (*ksconf.commands.ConfFileProxy method*), 91
load() (*ksconf.setup_entrypoints.LocalEntryPoint method*), 100
LocalEntryPoint (*class in ksconf.setup_entrypoints*), 100
location (*ksconf.conf.delta.DiffOp attribute*), 93

M

match_bwlist() (*in module ksconf.util.file*), 97
maturity (*ksconf.commands.KsconfCmd attribute*), 90
merge_conf_dicts() (*in module ksconf.conf.merge*), 94
merge_conf_files() (*in module ksconf.conf.merge*), 94

module_name (*ksconf.setup_entrypoints.Ep attribute*), 99

N

name (*ksconf.setup_entrypoints.Ep attribute*), 100

O

object_name (*ksconf.setup_entrypoints.Ep attribute*), 100

P

parse_conf() (*in module ksconf.conf.parser*), 95
parse_conf() (*ksconf.commands.KsconfCmd method*), 90
parse_conf_stream() (*in module ksconf.conf.parser*), 95
post_run() (*ksconf.commands.KsconfCmd method*), 90
pre_run() (*ksconf.commands.KsconfCmd method*), 90

R

readable() (*ksconf.commands.ConfFileProxy method*), 91
reduce_stanza() (*in module ksconf.conf.delta*), 94
register_args() (*ksconf.commands.KsconfCmd method*), 90
ReluctantWriter (*class in ksconf.util.file*), 97
relwalk() (*in module ksconf.util.file*), 97
reset() (*ksconf.commands.ConfFileProxy method*), 91
reset() (*ksconf.util.terminal.TermColor method*), 98
returncode (*ksconf.vc.git.GitCmdOutput attribute*), 98
run() (*ksconf.commands.KsconfCmd method*), 90

S

sanity_checker() (*in module ksconf.archive*), 99
section_reader() (*in module ksconf.conf.parser*), 96
set_parser_option() (*ksconf.commands.ConfFileProxy method*), 91
show_diff() (*in module ksconf.conf.delta*), 94
show_text_diff() (*in module ksconf.conf.delta*), 94
smart_copy() (*in module ksconf.util.file*), 97
smart_write_conf() (*in module ksconf.conf.parser*), 96
splitup_kvpairs() (*in module ksconf.conf.parser*), 96
stanza (*ksconf.conf.delta.DiffStanza attribute*), 93
stanza (*ksconf.conf.delta.DiffStzKey attribute*), 93
stderr (*ksconf.vc.git.GitCmdOutput attribute*), 98
stdout (*ksconf.vc.git.GitCmdOutput attribute*), 98
stream (*ksconf.commands.ConfFileProxy attribute*), 91
summarize_cfg_diffs() (*in module ksconf.conf.delta*), 94

T

tag (*ksconf.conf.delta.DiffOp attribute*), 93
TermColor (*class in ksconf.util.terminal*), 98
Token (*class in ksconf.conf.parser*), 95
tty_color() (*in module ksconf.util.terminal*), 98
type (*ksconf.conf.delta.DiffGlobal attribute*), 92
type (*ksconf.conf.delta.DiffStanza attribute*), 93
type (*ksconf.conf.delta.DiffStzKey attribute*), 93

U

unlink() (*ksconf.commands.ConfFileProxy method*), 91

V

version_extra (*ksconf.commands.KsconfCmd attribute*), 90

W

which() (*in module ksconf.util.file*), 97
writable() (*ksconf.commands.ConfFileProxy method*), 91
write() (*ksconf.util.terminal.TermColor method*), 98
write_conf() (*in module ksconf.conf.parser*), 96
write_conf_stream() (*in module ksconf.conf.parser*), 96