

Live Kernel Patching Using kGraft

SUSE Linux Enterprise Server 12 SP5

This document describes the basic principles of the kGraft live patching technology and provides usage guidelines for the SLE Live Patching service.

kGraft is a live patching technology for runtime patching of the Linux kernel, without stopping the kernel. This maximizes system uptime, and thus system availability, which is important for mission-critical systems. By allowing dynamic patching of the kernel, the technology also encourages users to install critical security updates without deferring them to a scheduled downtime.

A kGraft patch is a kernel module, intended for replacing whole functions in the kernel. kGraft primarily offers in-kernel infrastructure for integration of the patched code with base kernel code at runtime.

SLE Live Patching is a service provided on top of regular SUSE Linux Enterprise Server maintenance. kGraft patches distributed through SLE Live Patching supplement regular SLES maintenance updates. Common update stack and procedures can be used for SLE Live Patching deployment.

Publication Date: 07 Nov 2024

Contents

- 1 Advantages of kGraft 3
- 2 Low-level Function of kGraft 3

3	Installing kGraft Patches	4
4	Patch Life Cycle	6
5	Removing a kGraft Patch	6
6	Stuck Kernel Execution Threads	6
7	The kgr Tool	7
8	Scope of kGraft Technology	7
9	Scope of SLE Live Patching	8
10	Interaction with the Support Processes	8
11	GNU Free Documentation License	8

1 Advantages of kGraft

Live kernel patching using kGraft is especially useful for quick response in emergencies (when serious vulnerabilities are known and should be fixed when possible or there are serious system stability issues with a known fix). It is not used for scheduled updates where time is not critical. Typical use cases for kGraft include systems like memory databases with huge amounts of RAM, where boot-up times of 15 minutes or more are not uncommon, large simulations that need weeks or months without a restart, or infrastructure building blocks providing continuous service to many consumers.

The main advantage of kGraft is that it never requires stopping the kernel, not even for a short time period.

A kGraft patch is a `.ko` kernel module in a RPM package. It is inserted into the kernel using the `insmod` command when the package is installed or updated. kGraft replaces whole functions in the kernel, even if they are being executed. An updated kGraft module can replace an existing patch if necessary.

kGraft is also lean—it contains only a small amount of code, because it leverages other standard Linux technologies.

2 Low-level Function of kGraft

kGraft uses the `ftrace` infrastructure to perform patching. The following describes the implementation on the AMD64/Intel 64 architecture.

To patch a kernel function, kGraft needs some space at the start of the function to insert a jump to a new function. This space is allocated during kernel compilation by GCC with function profiling turned on. In particular, a 5-byte call instruction is injected to the start of kernel functions. When such instrumented kernel is booting, profiling calls are replaced by 5-byte NOP (no operation) instructions.

After patching starts, the first byte is replaced by the `INT3` (breakpoint) instruction. This ensures atomicity of the 5-byte instruction replacement. The other four bytes are replaced by the address to the new function. Finally, the first byte is replaced by the `JMP` (long jump) opcode.

Inter-processor non-maskable interrupts (IPI NMI) are used throughout the process to flush speculative decoding queues of other CPUs in the system. This allows switching to the new function without ever stopping the kernel, not even for a very short moment. The interruptions by IPI NMIs can be measured in microseconds and are not considered service interruptions as they happen while the kernel is running in any case.

Callers are never patched. Instead, the caller's NOPs are replaced by a JMP to the new function. JMP instructions remain forever. This takes care of function pointers, including in structures, and does not require saving any old data for the possibility of un-patching.

However, these steps alone would not be good enough: since the functions would be replaced non-atomically, a new fixed function in one part of the kernel could still be calling an old function elsewhere or vice versa. If the semantics of the function interfaces changed in the patch, chaos would ensue.

Thus, until all functions are replaced, kGraft uses an approach based on trampolines and similar to RCU (read-copy-update), to ensure a consistent view of the world to each user space thread, kernel thread and kernel interrupt. A per-thread flag is set on each kernel entry and exit. This way, an old function would always call another old function and a new function always a new one. Once all processes have the "new universe" flag set, patching is complete, trampolines can be removed and the code can operate at full speed without performance impact other than an extra-long jump for each patched function.

3 Installing kGraft Patches

This section describes the activation of the SUSE Linux Enterprise Live Patching extension and the installation of kGraft patches.

3.1 Activation of SLE Live Patching

To activate SLE Live Patching on your system, follow these steps:

1. If your SLES system is not yet registered, register it. Registration can be done during the system installation or later using the YaST *Product Registration* module (**yast2 registration**). After registration, click *Yes* to see the list of available online updates.

If your SLES system is already registered, but SLE Live Patching is not yet activated, open the YaST *Product Registration* module (**yast2 registration**) and click *Select Extensions*.

2. Select *SUSE Linux Enterprise Live Patching 12* in the list of available extensions and click *Next*.
3. Confirm the license terms and click *Next*.
4. Enter the SLE Live Patching registration code and click *Next*.
5. Check the *Installation Summary* and selected *Patterns*. The pattern Live Patching should be selected for installation.
6. Click *Accept* to complete the installation. This will install the base kGraft components on your system together with the initial live patch.

3.2 Updating System

1. SLE Live Patching updates are distributed in a form that allows using standard SLE update stack for patch application. The initial live patch can be updated using zypper patch, YaST Online Update or equivalent method.
2. The kernel is patched automatically during the package installation. However, invocations of the old kernel functions are not completely eliminated until all sleeping processes wake up and get out of the way. This can take a considerable amount of time. Despite this, sleeping processes that use the old kernel functions are not considered a security issue. Nevertheless, in the current version of kGraft, it is not possible to apply another kGraft patch until all processes cross the kernel-user space boundary to stop using patched functions from the previous patch.

To see the global status of patching, check the flag in /sys/kernel/kgraft/in_progress. The value 1 signifies the existence of sleeping processes that still need to be woken (the patching is still in progress). The value 0 signifies that all processes are using solely the patched functions and patching has finished already. Alternatively, use the kgr status command to obtain the same information.

The flag can be checked on a per-process basis too. Check the number in /proc/PROCESS_NUMBER/kgf_in_progress for each process individually. Again, the value 1 signifies sleeping process that still needs to be woken. Alternatively, use the kgr blocking command to output the list of sleeping processes.

4 Patch Life Cycle

Expiration dates of live patches can be accessed with `zypper lifecycle`. Make sure that the package `lifecycle-data-sle-live-patching` is installed.

```
tux > zypper lifecycle

Product end of support
Codestream: SUSE Linux Enterprise Server 12                2024-10-31
SUSE Linux Enterprise Server 12 SP2                       n/a*

Extension end of support
SUSE Linux Enterprise Live Patching                       2017-10-31

Package end of support if different from product:
SUSEConnect                Now, installed 0.2.41-18.1, update available
  0.2.42-19.3.1
apache2-utils              Now

*) See https://www.suse.com/lifecycle for latest information
```

When the expiration date of a patch is reached, no further live patches for this kernel version will be supplied. Plan an update of your kernel before the end of the live patch life cycle period.

5 Removing a kGraft Patch

To remove a kGraft patch, use the following procedure:

1. First remove the patch itself using Zypper:

```
zypper rm kgraft-patch-3_12_32-25-default
```

2. Then reboot the machine.

6 Stuck Kernel Execution Threads

Kernel threads need to be prepared to handle kGraft. Third-party software may not be ready for kGraft adoption and its kernel modules may spawn kernel execution threads. These threads will block the patching process indefinitely. As an emergency measure kGraft offers the possibility

to force finishing of the patching process without waiting for all execution threads to cross the safety checkpoint. This can be achieved by writing `0` into `/sys/kernel/kgraft/in_progress`. Consult SUSE Support before performing this procedure.

7 The **kg**r Tool

Several kGraft management tasks can be simplified with the **kg**r tool. The available commands are:

kgr status

Displays the overall status of kGraft patching (`ready` or `in_progress`).

kgr patches

Displays the list of loaded kGraft patches.

kgr blocking

Lists processes that are preventing kGraft patching from finishing. By default only the PIDs are listed. Specifying `-v` prints command lines if available. Another `-v` displays also stack traces.

For detailed information, see `man kg`r.

8 Scope of kGraft Technology

kGraft is based on replacing functions. Data structure alteration can be accomplished only indirectly with kGraft. As a result, changes to kernel data structure require special care and, if the change is too large, rebooting might be required. kGraft also might not be able to handle situations where one compiler is used to compile the old kernel and another compiler is used for compiling the patch.

Because of the way kGraft works, support for third-party modules that are spawning kernel threads is limited.

9 Scope of SLE Live Patching

Fixes for SUSE Common Vulnerability Scoring System (CVSS) level 7+ vulnerabilities and bug fixes related to system stability or data corruption will be shipped in the scope of SLE Live Patching. It might not be possible to produce a live patch for all kinds of fixes fulfilling the above criteria. SUSE reserves the right to skip fixes where production of a kernel live patch is unviable because of technical reasons. For more information on CVSS 3.0, which is the base for the SUSE CVSS rating, see <https://www.first.org/cvss/>.

10 Interaction with the Support Processes

While resolving a technical difficulty with SUSE Support, you may receive a so-called Program Temporary Fix (PTF). PTFs may be issued for various packages including those forming the base of SLE Live Patching.

kGraft PTFs complying with the conditions described in the previous section can be installed as usual and SUSE will ensure that the system in question does not need to be rebooted and that future live updates are applied cleanly.

PTFs issued for the base kernel disrupt the live patching process. First, installing the PTF kernel means a reboot as the kernel cannot be replaced as a whole at runtime. Second, another reboot is needed to replace the PTF with any regular maintenance updates for which the live patches are issued.

PTFs for other packages in SLE Live Patching can be treated like regular PTFs with the usual guarantees.

GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non-commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects. If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.

- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles. You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.  
Permission is granted to copy, distribute and/or modify this document  
under the terms of the GNU Free Documentation License, Version 1.2  
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.