Table of Contents

Software update policy for SLC (for desktops and auto-updating machines)

Software updates get triggered by the following events:

- notification from T.U.V. (mail from enterprise-watch-list) - standard workflow, SRPMs available
- availability of updates from SL (no notification?) - standard workflow, RPMs available
- direct security notification from TAM, CERN security team, other mailing list or (occasionally) users - SRPMs typically not available.
- direct errata notification or update request from any add-on software provider

Assessment

Severity will be defined similar to Red Hat (see their classification) or MS:

- **Critical**: roughly: local root exploit, remote user exploit and similar
- **Important**: local user/daemon exploit (automatic or user-assisted); system-wide DoS (kernel panic etc); remote service DoS
- **Moderate**: anything else
- **Low**: technically a security exploit but low/no impact at CERN assumed (i.e. DoS on an service that typically isn't used at CERN)

Impact on CERN for all security updates needs to be assessed. If available, input from TUV (or secunia advisories etc) can be taken as a base, and gets modified:

- local root exploits in a "CERN recommened setup" install have **critical** priority (even if TUV only labels these as "important" since local access is required first) - provide advance warning to FIO SMOD in these cases.
- user-assisted attacks (user needs to perform some action - visit "evil" web page, read "evil" attachment) on the user account are **not critical** but **important**, unless there is reason to believe there is an ongoing attack (which is working against Linux machines). This includes browsers, email clients, multimedia and other plugins/applications such as PDF viewers.

Timeline/delay goals

(these eventually should be measured automatically via IncomingUpdateWorkflow tools)

- *standard workflow* updates get pushed maximum once per week, announce beforehand at CCSR if reboot will be required to activate
- non-Critical updates should be in the *testing* repository for at least 2 days (i.e. full day after the machines subscribed to that repository did an automatic update)
- **Important** and **Critical** standard updates should not get delayed for more than 7 days after release by the vendor, unless testing has shown them to fail.

Implementation

Procedure is at SoftwareUpdatesOnSLCOld.

Responsible

Person on duty for Linux 3rd-level REMEDY support needs to add all updates during their week to "testing" repository and performs weekly push to production. Anything not pushed during the week needs to be passed on to next support rota person, with an assessment whether the software in question should (or not) get pushed during the next week.

Software update policy for RHEL

Updates from RHEL are pushed whenever they come out of RHN, last step on our side is the announcement to project-elfms (after which the FIO-FS "monthly intervention" should kick in). Some short (= days) delay on the Linux Support side is OK, in order to synchronize with matching SLC updates. See SoftwareUpdatesOnRHE2 for the procedure.

Other policies

Updates on the FIO-centrally managed machines are governed by a different policy (roughly: monthly update after one week lock-off period, see FsScheduledLinuxUpgradeTemplate). Emergency interventions can still be rushed.

This topic: LinuxSupport > SoftwareUpdatesPolicy
Topic revision: r5 - 2013-01-07 - ThomasOulevey