

June 11, 2009
NTT

NILFS merged into Linux as the first File System from Japan

- Fine-grain data protection with continuous snapshotting -

NILFS, a Log-structured File System ^{([*1](#))} developed at NTT (headquartered in Chiyoda-ku, Tokyo with Satoshi Miura as President) is now merged into the 2.6.30 release of the mainline Linux ^{([*2](#))} kernel. It is the first-ever file system, one of the core components of operating systems (OS), developed in the Japanese community to get merged into Linux.

With NILFS, a snapshot ^{([*3](#))}, a historical view of the file system, gets continuously taken without stopping services, and the system's state or user's data can be restored from any point in the past. As a result of this mainline merge, NILFS is now distributed as part of Linux. The combination of NILFS and Linux would increase the maintainability of the system.

1. Background and significance of mainline merge

NTT Cyber Space Laboratories (referred to as simply NTT Labs below) is utilizing open source software ^{([*4](#))}, aiming to support enterprise core systems, and concentrating to improve reliability and functionality of OS and middleware such as database management systems. NTT Labs has developed NILFS as part of this activity, and released the first version in September 2005, the second version (NILFS2) in June 2007 as open source software mainly for open source developers. NILFS2 has been packaged in Linux distributions ^{([*5](#))} such as Debian GNU/Linux ^{([*6](#))}. NTT Labs has been proposing the mainline merge of NILFS to achieve stable and lasting use in the actively developed Linux kernel. As a result, NILFS is now introduced as one of the file systems in Linux.

With the mainline merge of NILFS, fine-grain data protection with continuous snapshotting is now part of Linux. Furthermore, by providing a file system with distinct features, NTT Labs contributes to the active discussion and the development of the next generation file systems that is the core component for data reliability, and to the improvement of future Linux in functionality and reliability.

2. Technical features

NILFS is a file system for Linux that saves a snapshot, a historical view of the file system, automatically and continuously. This feature enables the system not only to quickly recover from system failures, but also to recover from human operation errors because the system's state can be restored from any point in the past.

- (1) All additions and modifications to data are appended to the storage with checksums ^{([*7](#))} and existing data won't be overwritten until the protection period gets expired. This reduces the probability of data corruption dramatically.

- (2) A time-stamped snapshot is taken when a series of data writes (transaction) is completed, which enables recovering from failures without consistency checks. Furthermore, as long as there is enough storage space, the system's state and user's data at any point in the past can be restored.
- (3) The snapshots are taken automatically and continuously without stopping services. Furthermore, because the snapshots can be read even when the system is running, the whole system backup can be taken without stopping services.
- (4) B-tree (*8) is exploited for managing the file system to achieve efficient processing of large files and large numbers of files.
- (5) A cleaner mechanism (introduced from NILFS2) which deletes snapshots automatically and reclaims unnecessary disk space is implemented. To keep more important snapshots longer, it can be configured to prevent them from deletion. NILFS is the first log-structured file system to achieve both multiple snapshots and the cleaner mechanism.
- (6) NILFS can be used on commodity PC without specific hardware therefore, it can be applied to various Linux systems ranging from personal desktop PC to enterprise servers.

3.Future plans

NTT Labs will continue collaborating with the Linux development community that consists of developers around the world, and will contribute to the active discussion and the development of the next generation file systems, and to the improvement of future Linux in functionality and reliability. Furthermore, NTT Labs will pursue the technology based on NILFS, and contribute to achieve a safe and secure information sharing society.

[Notes]

"TOMOYO", the security enhancement mechanism developed at NTT Data Corporation, is also merged into the 2.6.30 release of the mainline Linux kernel. [Please refer to the news release at the NTT Data Corporation website.](#) "TOMOYO" is a registered trademark of NTT Data Corporation.

Terminology

*1: File System

A file system is a mechanism to manage data recorded to storages such as hard drives or USB flash drives. It provides functionalities of the OS related to data storing, such as creating files or directories (folders), reading or writing data of a file, and moving or deleting files or directories.

*2: Linux

Linux is an operating system kernel developed by Linus Torvalds in 1991 while he was studying at Helsinki University in Finland. The code was released as open source software under the GNU GPL and under continuous improvement by volunteer developers around the world. Linux is a registered trademark of Linus Torvalds in the U.S. and other countries.

*3: Snapshot

Snapshot is a whole image of the file system at the specific point in time.

*4: Open source software

This is software for which the source code is published on the Internet or other media so that anyone can improve or redistribute the code. The term "open source" attracted much attention and its meaning had come to depend on who was using it, so the Open Source Initiative (OSI) has proposed "The Open Source Definition" (OSD).

***5: Linux distribution**

Linux distribution is a package of the Linux kernel, device drivers, window systems and collection of software applications.

***6: Debian GNU/Linux**

Debian GNU/Linux is a Linux distribution that the Debian Project, an association of individuals who have made common cause to create a free operating system, is developing. Debian is a registered trademark of Software in the Public Interest, Inc.

***7: Checksum**

This is one technique for detecting errors in data that employs simple computation such as summing up all of the data. The value of the checksum is stored along with the data. When the data is read, the same computation is performed and the result is compared with the stored checksum value. The difference of these values means corruption in the data.

***8: B-tree**

This is a method for rapid search of dynamically changing data. It has been used widely in database systems since it was reported in 1972, and its use in the recent file systems is increasing.

Send inquiries to:
Public Relations, Planning Division
NTT Cyber Communication Laboratories
TEL: 046-859-2032
E-mail: randd@lab.ntt.co.jp

NTT NEWS RELEASE 