




IBM Systems Group

# Network Related Improvements in the 2.6 Linux Kernel

By Robbie Williamson

[robbiew@us.ibm.com](mailto:robbiew@us.ibm.com)

## Background Information on the LTP



You Could Learn A Lot From The LTP.

[www.linuxtestproject.org](http://www.linuxtestproject.org)

## Introduction and Overview

- ★ The Linux Test Project is a joint project started by SGI™, with the goal of delivering test suites to the open source community that validate the reliability, robustness, and stability of Linux.
- ★ Currently, it is **the only** comprehensive opensource test suite available for the Linux Operating System.
- ★ The number one result on Yahoo and Google when “linux test” is used as search parameters.

## Current Collaborating Organizations

- ★ SGI™
- ★ IBM®
- ★ Open Source Development Lab (OSDL™)
- ★ Group Bull®
- ★ Wipro Technologies
- ★ Intel® (Open POSIX Test Suite)
- ★ SUSE
- ★ Debian
- ★ IBM® OZLabs
- ★ IBM® China
- ★ IBM® India

## Examples of LTP Tested Architectures

- ★ i386/i586/i686
- ★ ia64
- ★ ppc
- ★ ppc64
- ★ s390
- ★ s390x
- ★ x86\_64 (AMD's Opteron™)
- ★ Embedded Architectures:
  - ▶ ARM
  - ▶ the IXIA® 10/100 Mbps Ethernet Load Module
  - ▶ IBM's FipS (Power 5 Firmware)

## LTP Tested Areas (2100+ Tests)

- ★ Filesystems
- ★ Logical Volume Manager (LVM)
- ★ Memory
- ★ Kernel Module Functions
- ★ Timers
- ★ Hyperthreading
- ★ Scheduler
- ★ System Calls
- ★ Signals
- ★ Floating Point Operations
- ★ Application Development Environment Commands
- ★ System Administration Commands
- ★ Large Page Support
- ★ Asynchronous I/O
- ★ Disk I/O
- ★ Threading
- ★ Shared Memory
- ★ PTY
- ★ Semaphores
- ★ Message Queues
- ★ IPv4/IPv6
- ★ Sockets
- ★ Network File System (NFS)
- ★ Remote Procedure Calls (RPC)
- ★ Stream Control Transmission Protocol (SCTP)
- ★ Networking Related Commands
- ★ Virtual Devices
- ★ Open Hardware Platform Interface (HPI)



## Some Uses of the LTP Test Suite

- **Pre-GA Hardware Validation** - such as SUSE's Andi Kleen, who uses the test suite for testing his x86\_64 (AMD Opteron) port of the kernel. IBM uses the LTP for testing the 2.6 Linux kernel with their latest Power 5 architecture machines.
- **Application Development** – Application developers use the suite to verify correct system call behavior when porting their UNIX applications to Linux.
- **Distribution Regression Testing** - Debian includes the LTP as an installable package for their release. SUSE includes the LTP in their distro release test cycle.
- **Kernel Development** – Kernel maintainer, Matthew Wilcox, uses the suite for testing file locking. The 2.6 Linux kernel maintainer, Andrew Morton, uses the LTP for regression testing the release candidates.
- **Reliability Testing** – The IBM Linux Technology Center uses the LTP testsuite in its ongoing effort to test the long-term reliability and stability of the Linux operating system. IBM's developerWorks website published their work, "Putting Linux Reliability to the Test".

**LTP**

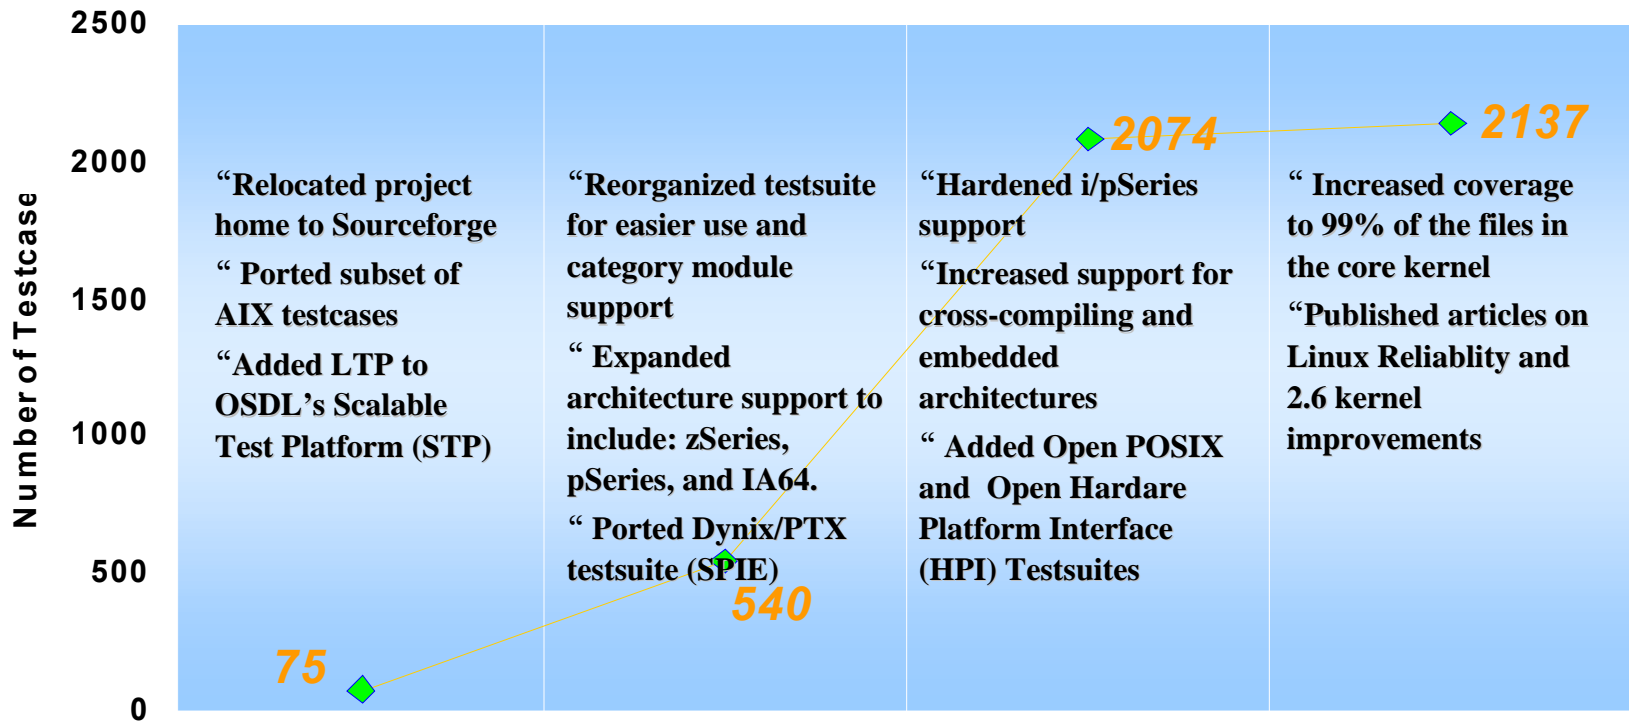
Testcase Coverage

Examples

- Asynch I/O
- Commands
- Device Drivers
- Filesystems
- I/O
- Inter-Process Control
- Large Page Support
- Memory
- Modules
- Pty
- Scheduler
- SCSI
- Syscalls
- Timers
- Math
- Networking
- Open POSIX
- Open HPI
- Virtual Devices



www.linuxtestproject.org



“Relocated project home to Sourceforge  
 “ Ported subset of AIX testcases  
 “Added LTP to OSDL’s Scalable Test Platform (STP)

“Reorganized testsuite for easier use and category module support  
 “ Expanded architecture support to include: zSeries, pSeries, and IA64.  
 “ Ported Dynix/PTX testsuite (SPIE)

“Hardened i/pSeries support  
 “Increased support for cross-compiling and embedded architectures  
 “ Added Open POSIX and Open Hardware Platform Interface (HPI) Testsuites

“ Increased coverage to 99% of the files in the core kernel  
 “Published articles on Linux Reliability and 2.6 kernel improvements

2001	2002	2003	2004 (1st Half)
<ul style="list-style-type: none"> <li>•IBM joins the Linux Test Project</li> <li>•Open Source Development Labs (OSDL) joins project</li> </ul>	<ul style="list-style-type: none"> <li>•Group Bull joins project</li> <li>•Used as 2.4 Kernel Release Candidate Regression Testsuite</li> </ul>	<ul style="list-style-type: none"> <li>•Wipro Technologies joins project</li> <li>•Project booth at Linux World Expo</li> </ul>	<ul style="list-style-type: none"> <li>•IBM China becomes major test contributor</li> <li>•LTP used across all IBM server brands for testing Linux</li> </ul>
<p>Page Views: 28,530</p> <p>Downloads: 2,463</p>	<p>Page Views: 104,786</p> <p>Downloads: 7,102</p>	<p>Page Views: 162,818</p> <p>Downloads: 8,679</p>	<p>Page Views: 196,660*</p> <p>Downloads: 13,192*</p> <p>*projected total for the year</p>



# Overview of 2.6 Network Improvements

There are many improvements in the new Linux 2.6 kernel, when compared to the current 2.4 version. One area of technical advancement is in the kernel networking options. The following are some of the major areas of improvement:

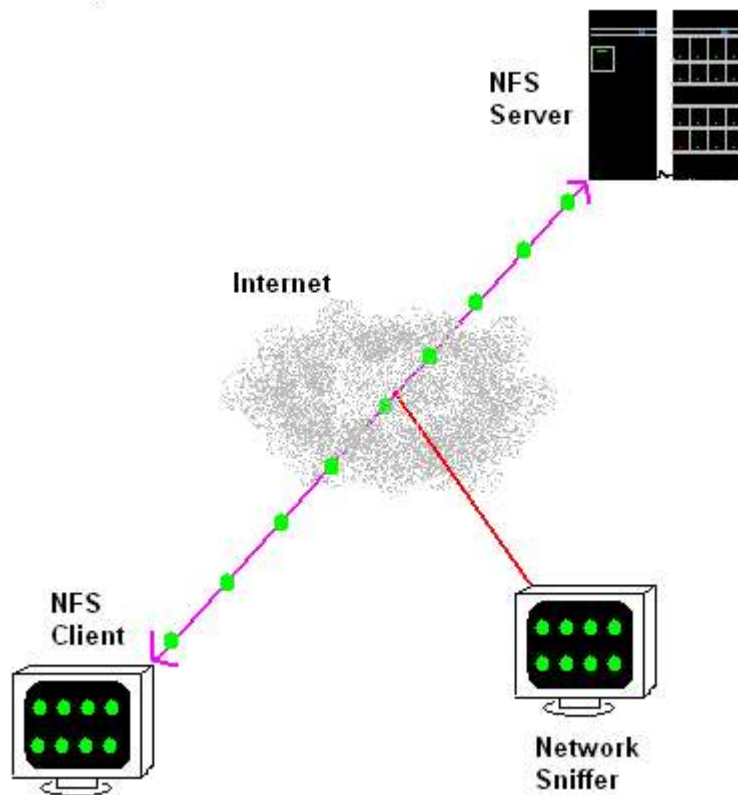
- **NFSv4: Networking File System version 4**
  - **Cryptographic Authentication**
  - **Compound Procedures**
  - **Server Daemon Reduction**
  - **Support for Filesystem Migration and Replication**
- **SCTP: Stream Control Transmission Protocol**
  - **Multi-Streaming**
  - **Multi-Homing**
- **IPSec: Internet Protocol Security**
  - **Authentication Header (AH)**
  - **Encapsulated Security Payload (ESP)**
- **IPComp: IP Payload Compression**
- **IPv6 Improvements**
  - **IPSec & IPComp with Tunneling Support**
  - **Privacy Extensions**

# Network File System version 4


# NFSv4

**“Users of NFS may now conduct secure transactions utilizing RPCSEC\_GSS, which is an implementation of the General Security Service (GSS) API with Remote Procedure Calls (RPC).**”

### NFSv2/3



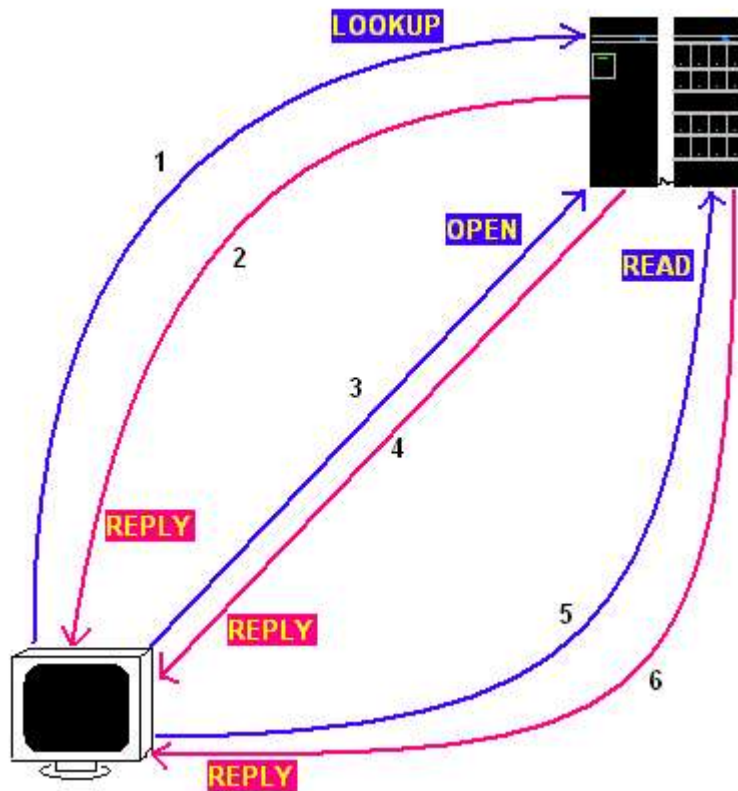
### NFSv4




# NFSv4

“Designers introduced the compound procedure, which allows combination of multiple requests into one.

## NFSv2/3

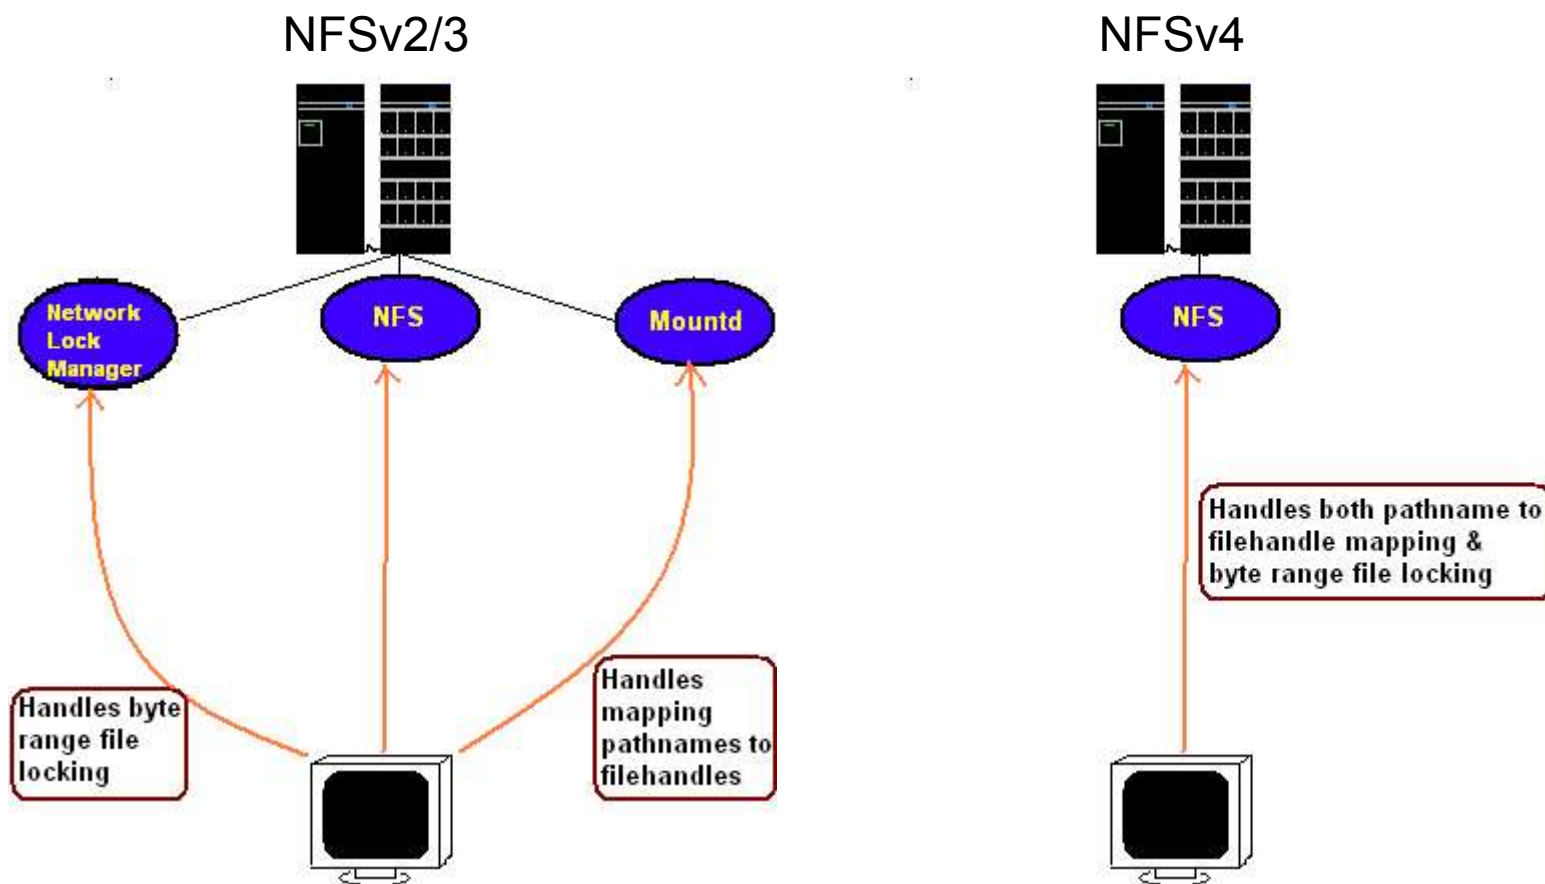


## NFSv4



# NFSv4

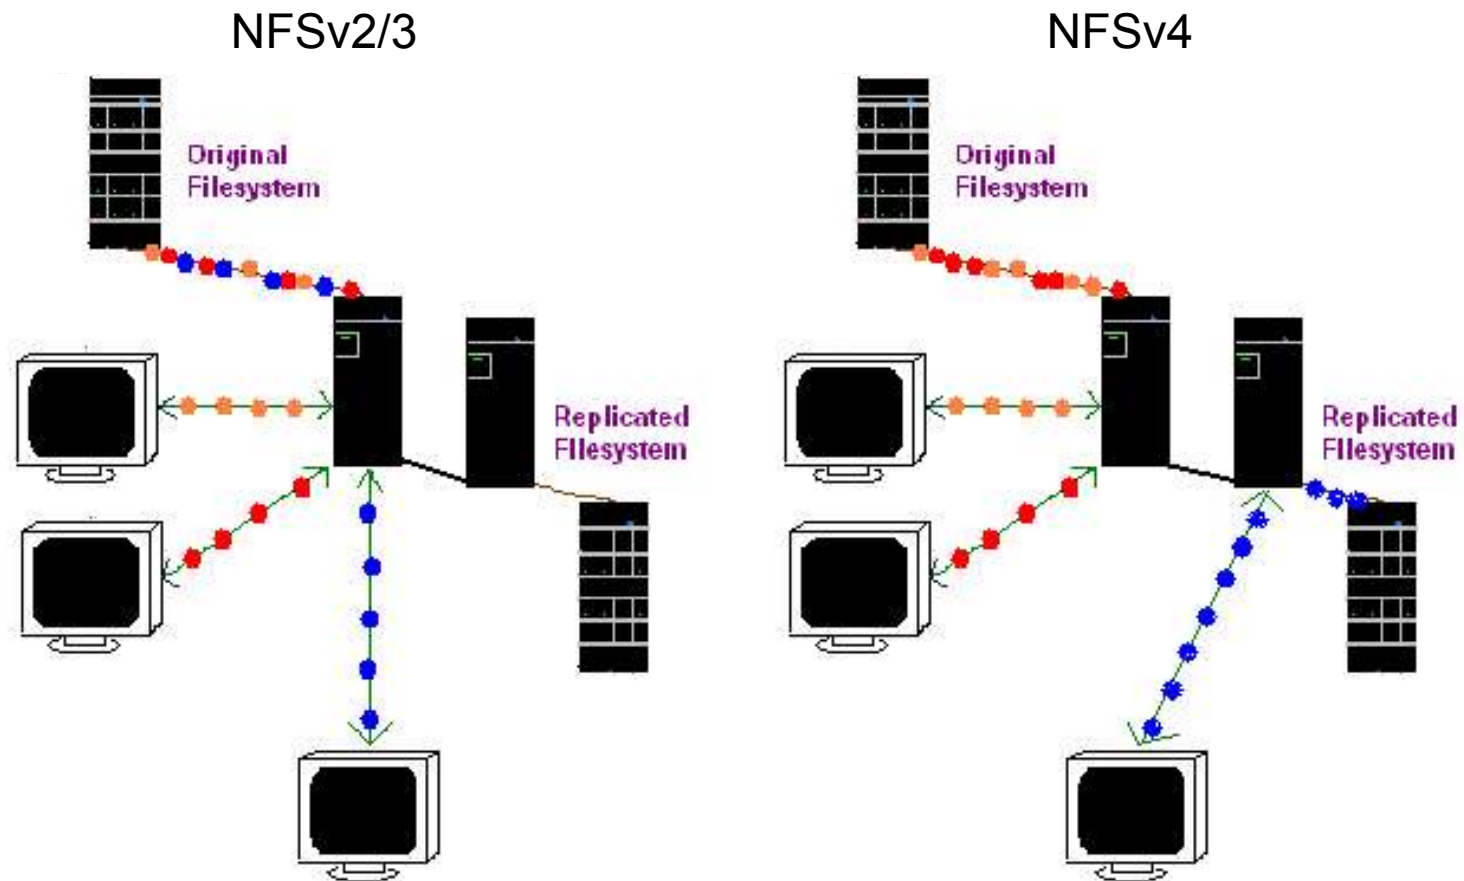
“NFS now handles filehandle-to-path name mapping and byte range file locking, which reduces the number of server-side support daemons required.





# NFSv4

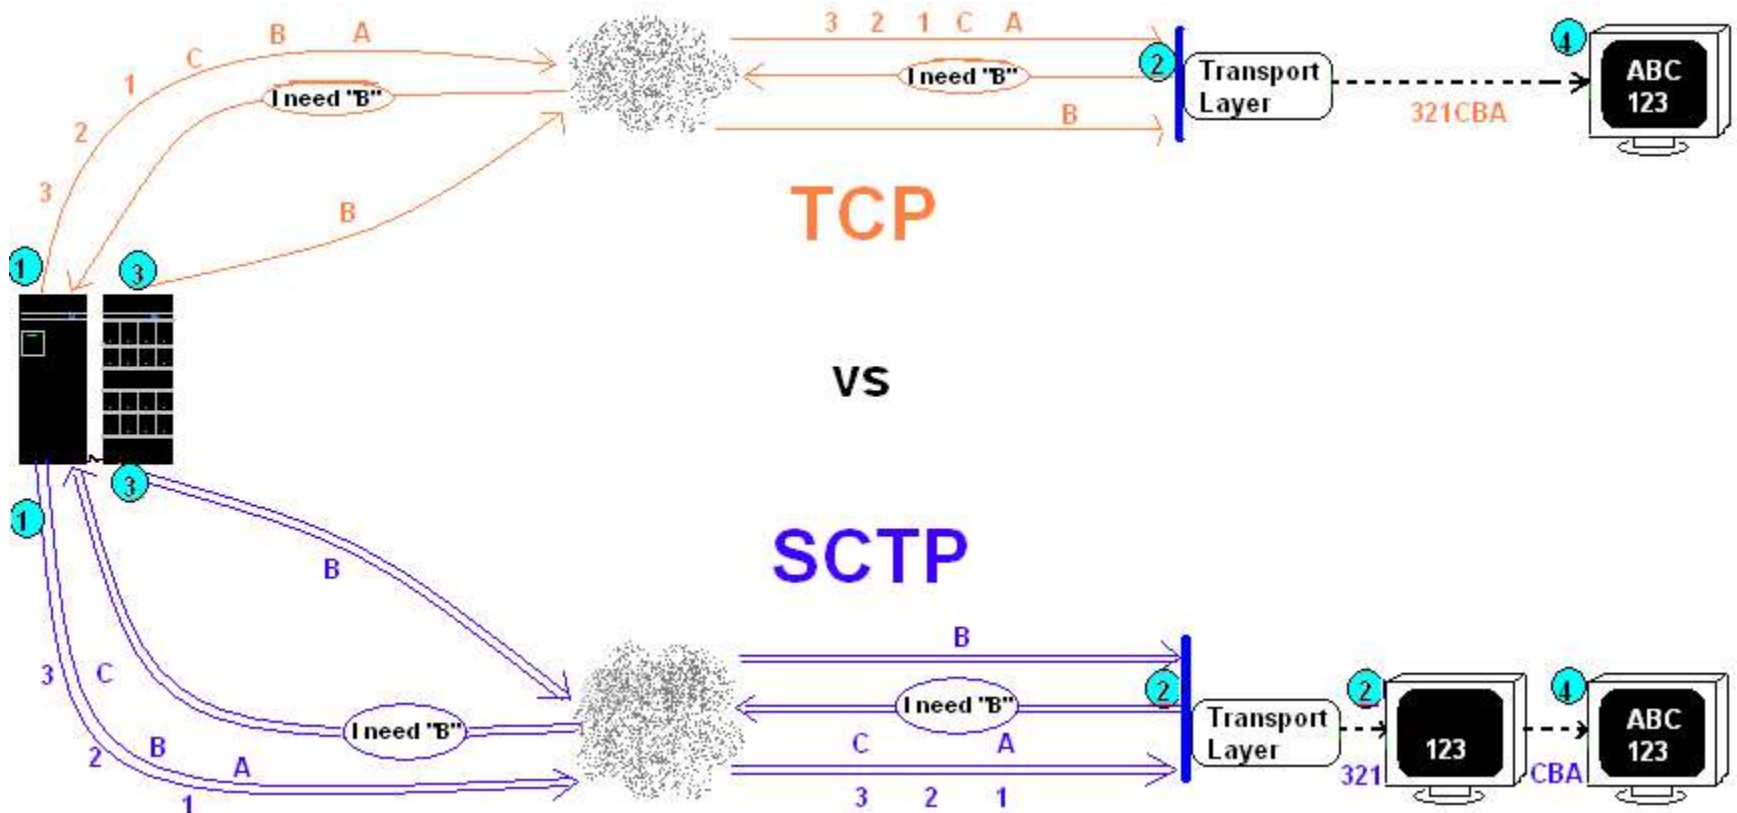
“Provides support for server migration and replication to enable clients to seamlessly change filesystems when either event occurs.



# Stream Control Transmission Protocol

# SCTP

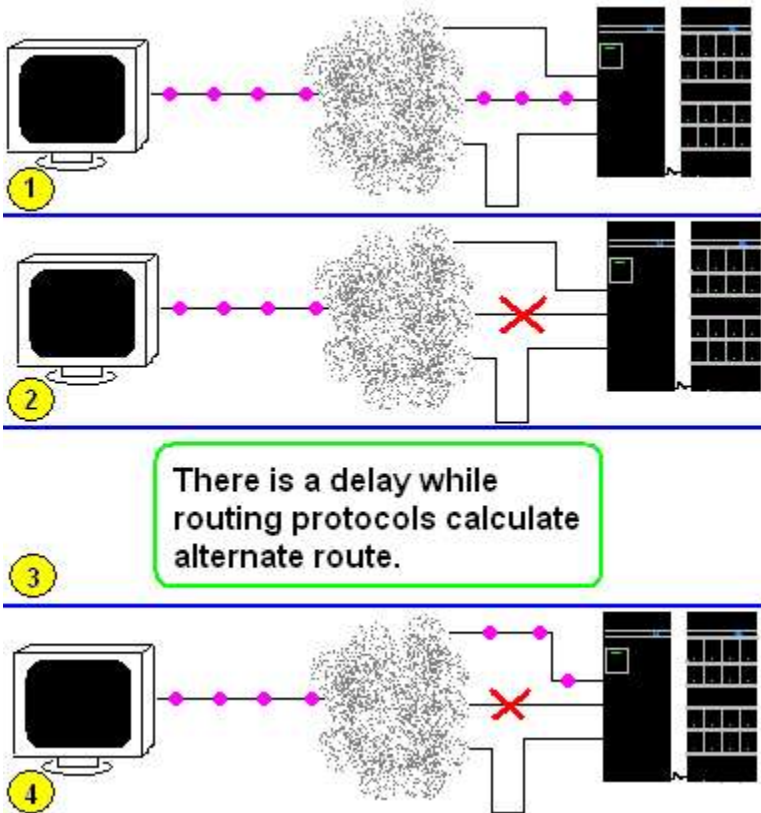
“Multi-Streaming allows data to be partitioned into multiple independently sequenced streams.



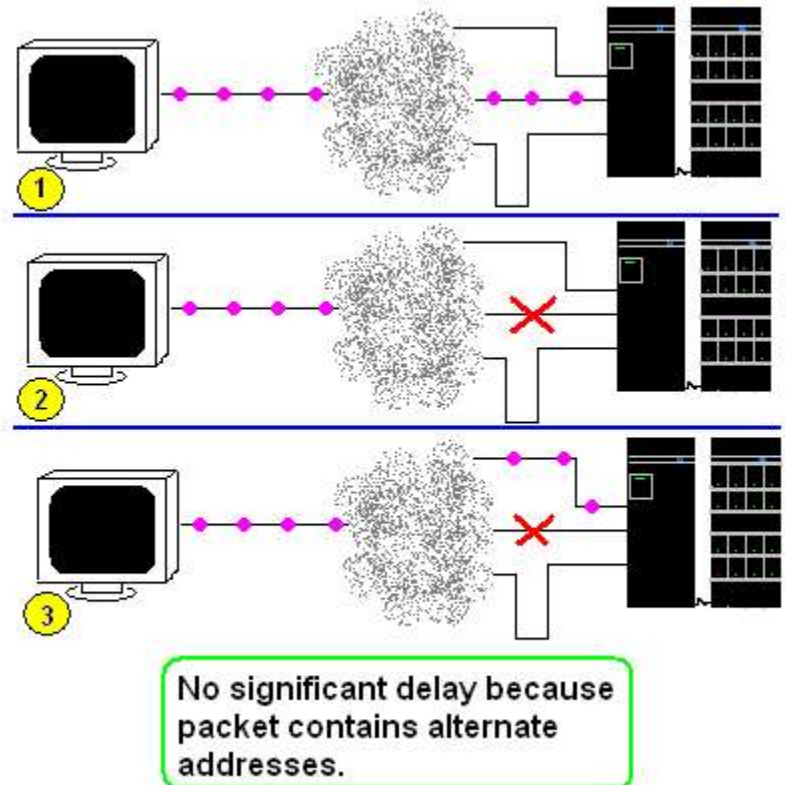
# SCTP

“Multi-Homing provides the ability for a single SCTP endpoint to support multiple IP address.

## TCP



## SCTP




# Internet Protocol Security

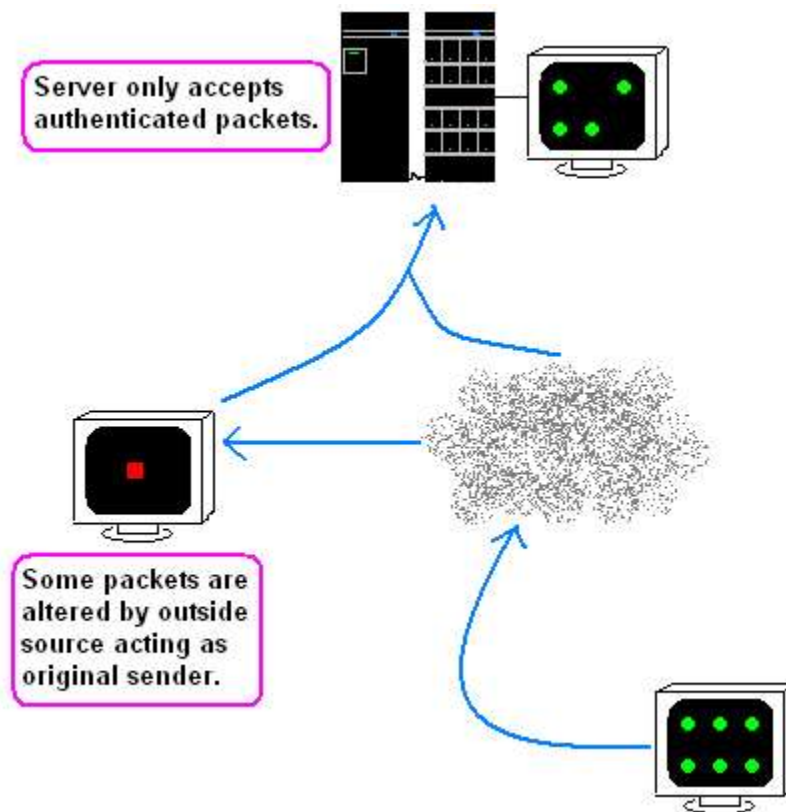
# IPSec

**“The Authentication Header (AH) allows users to be sure that the packet received came from a particular machine, and it's contents were not altered somewhere along the way.**

### Without AH



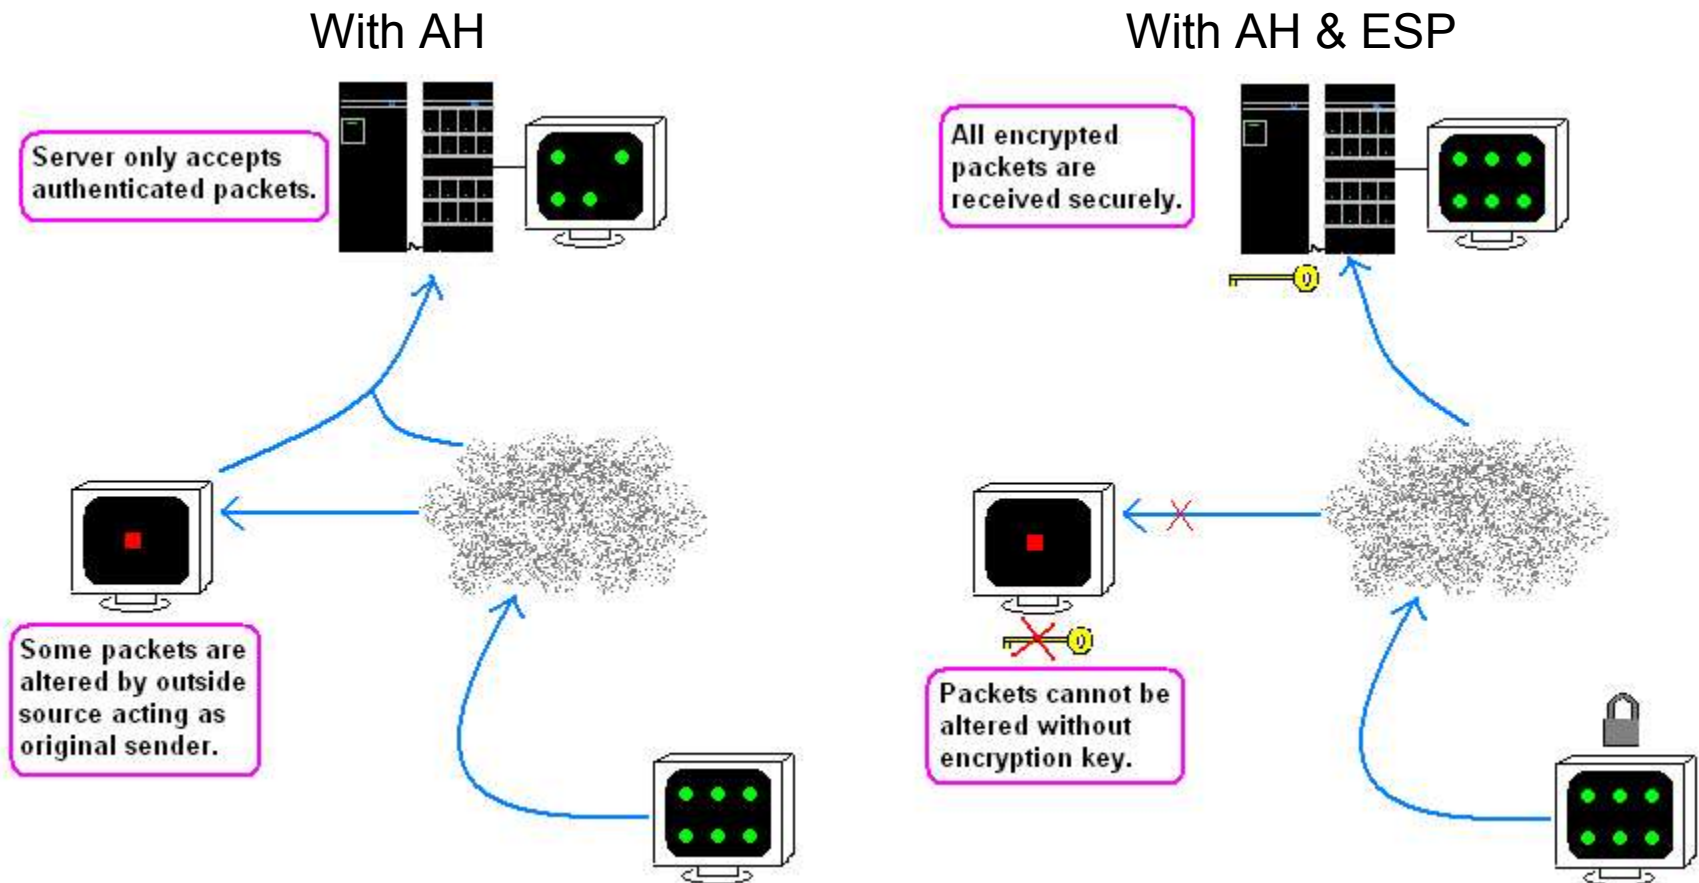
### With AH





# IPSec

**“With the Encapsulated Security Payload (ESP) header users have the benefit of encryption, as well as packet authentication.”**




# IP Payload Compression


# IPComp

**“IP Payload Compression (IPComp) reduces the size of IP datagrams, which can improve the performance of data transfer between two endpoints.”**

Without IPComp



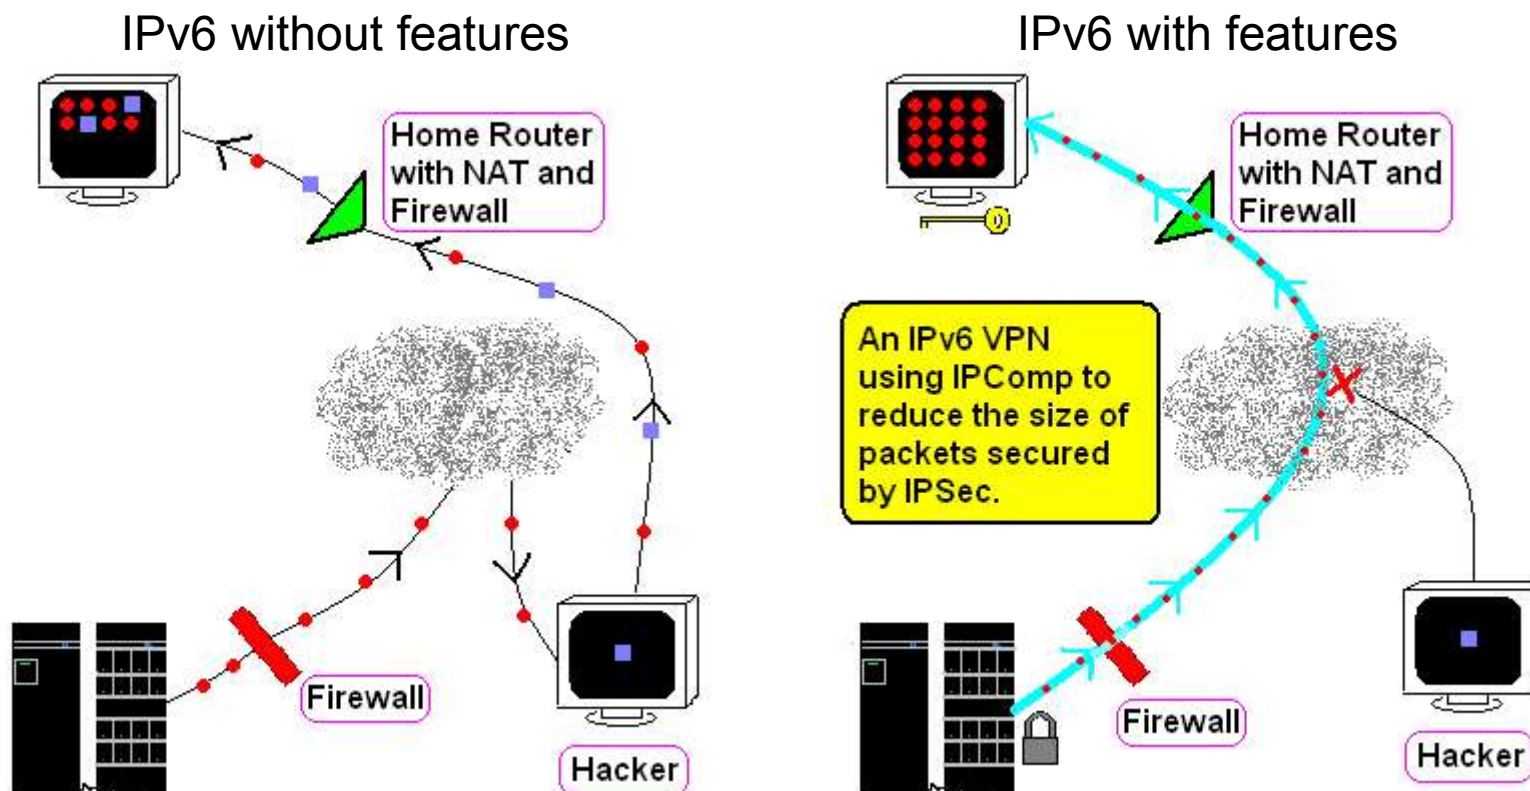
With IPComp



# IPv6 Improvements

# IPv6

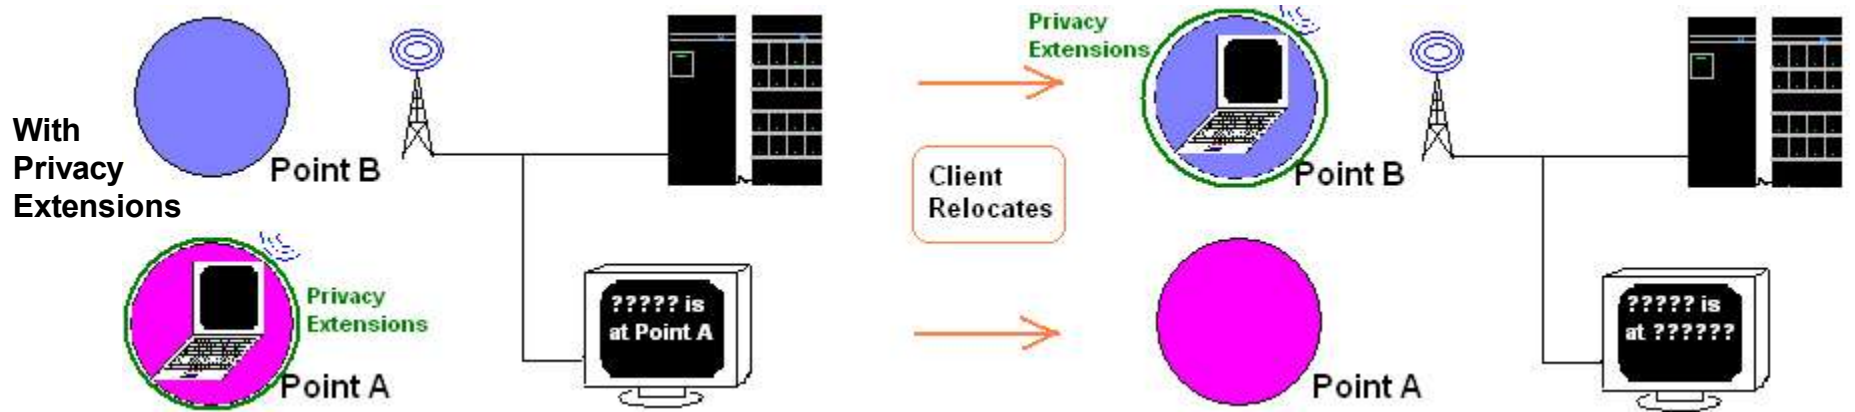

**“IPSec for IPv6 provides the same level of authentication and security as it does for IPv4. The inclusion of support for IPv6-to-IPv6 tunneling allows for secure seamless communication to occur between two end-points, i.e. Virtual Private Networks.**





# IPv6

**“IPv6 Privacy Extensions give users the option to protect their identity with regard to their IPv6 address.”**





## Conclusion

**Not all Linux users will benefit from these advances in networking technology provided by the 2.6 kernel. However, most users may find that one or more of these features can improve the way they utilize Linux in their system environment:**

- ▶ **Current users of NFS, looking for improved performance or security, will gain in both areas by migrating to version 4.**
- ▶ **Developers of carrier-grade and telephony applications can utilize the features provided by SCTP to ensure better, more reliable service for customers.**
- ▶ **IPSec provides solutions for people and businesses who need a method of transmitting secure data across insecure networks.**
- ▶ **IPComp allows these same groups the ability to improve data communication across the Internet by using smaller packet sizes during transmission.**
- ▶ **The enhancements to IPv6 provide better security and privacy to those using this next generation internet protocol, and allow for more IPv4 application developers to make the transition to using this improved version of IP, without compromising security.**

# References

- ❖ **Conta, A. and Deering, S., "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998**
- ❖ **Kent, S. and Atkinson, R., "IP Authentication Header", RFC 2402, November 1998**
- ❖ **Kent, S. and Atkinson, R., "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998**
- ❖ **Narten, T. and Draves, R., "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001**
- ❖ **Ong, L. and Yoakum, J., "An Introduction to the Stream Control Transmission Protocol (SCTP)", RFC 3286, May 2002**
- ❖ **Shacham, A., Monsour, B., Pereira, R., and Thomas, M., "IP Payload Compression Protocol (IPComp)", RFC 3173, September 2001**
- ❖ **Shepler, S., Callaghan, B., Robinson, D., Thurlow, R., Beame, C., Eisler, M., and Noveck, D., "Network File System (NFS) version 4 Protocol", RFC 3530, April 2003**

## For More Information

### “ IBM Linux Technology Center

- Internal - <http://ltc.linux.ibm.com>
- External - <http://www.ibm.com/linux/ltc>

### “IBM developerWorks - Linux

- <http://www.ibm.com/developerworks/linux/>

### “The Linux Test Project

- <http://ltp.sourceforge.net>
- or
- <http://www.linuxtestproject.org>