

Studies in Linux data storage

SAVE IT

This month we look at filesystems for SSDs and show you how to get connected with a Windows Active Directory file server. **BY JOE CASAD AND RAINER LUKAS**

Life was so easy when all the data for a standalone computer stayed on a little local hard drive. If the hard disk died, you were out of luck (unless you had the habit of performing regular backups to a tape drive or a bevy of floppy disks), but as long as it was working, you never had to worry about connectivity, network authentication, and the array of hardware and filesystem compatibility issues facing today's IT professionals. The good news is, it is much easier to back up your data now. The bad news is, the amount of data you have to back up is astronomically larger, and the tools for accessing data storage systems are vastly more complicated.

This month we study Linux filesystem options for Solid-State Drives (SSDs), and we examine some important techniques for accessing Windows file servers in Active Directory environments. You'll learn how to set up your Linux system as an Active-Directory-ready Kerberos client, and we'll even show you a shortcut for easy AD access using Likewise Open.

So Much Data

By the end of the five-year period between 2006 and 2011, the volume of data stored globally – our digital universe – will have exploded to 10 times its current size. Toward the end of this period, 1,800 exabytes of new data are being added each year. This unimaginable mass of information spans a variety

COVER STORY

Optimized Filesystem for SSD.....22

Linux with Active Directory26

of formats and containers, with the number of containers growing one and a half times as fast as the volume of data itself. The forecast figure for 2011 is 20 quadrillion – 20 million billion – containers (data files, images, tags, and so on).

This explosion in data storage has caused (or been caused by) an explosion in the data storage hardware. The world's bits now reside on a rich collection of exotic devices. Some of the terms that turn up frequently in storage discussions include:

- DAS (Direct Attached Storage), or Server Attached Storage – a disk attached to an individual host. The typical interface here is SCSI, and increasingly SAS. However, any block-oriented data transfer protocol, such as ATA/ATAPI, Fibre Channel, iSCSI, or FICON/ESCON is possible.
- SAN (Storage Area Network) – a network between servers and the storage resources they use. The data traffic on a SAN mainly comprises block-based data transfers, wherein individual blocks are transferred instead of whole files. The transport protocols commonly seen here are SCSI, Fibre Channel, or iSCSI.
- NAS (Network Attached Storage) – basically, an easy-to-manage file server. NAS is typically used to add storage capacity to an existing computer network without too much administrative overhead. In contrast to Direct Attached Storage (DAS), a NAS is always a separate host with its own operating system, although the OS will be heavily customized for the file server role. The link to the NAS uses Ethernet/IP. The overhead associated with supporting standard Ethernet-based network communication precludes the possibility of high-speed mass storage. A Storage Area Network (SAN) avoids these drawbacks.

Many devices come with built-in (or easily configurable) fault tolerance features, and remote backup to a data center is always an option for an extra measure of safety.

Migrating from one storage device to another presents a range of complications, depending on the devices and the means of migration. Possibly the easiest situation for data migration occurs when data are stored on local disks managed by a RAID controller. If the migration target has a more recent RAID controller, however, you won't be able to just reconnect the disks. If the data reside on a software RAID, a physical move might be possible in some circumstances, but chances are you would rather upgrade to newer, larger, and faster disks. If you use an external RAID array, you will typically have the option of replicating the data on a second identical (or similar) system. This will normally mean purchasing the new system from the same manufacturer. If you can't attach the old memory subsystem to the new server, you have to use the IP network to transfer the data.

If the data are on a Direct Attached Storage system (Figure 1), such as a disk array connected directly to the server, the situation is very similar to the case of an internal hard disk, the only difference being that a professional storage subsystem often offers the ability to connect to multiple servers at the same time. Thanks to this option, you can connect your new server to the existing array, mount the old and new volumes for the migration, and simply create a local copy of the data.

With a Storage Area Network (SAN), you can simply mount the old server's volumes on the new server, assuming you can configure the new server to interoperate with the SAN.

If you already have Network Area Storage (NAS) systems running at your data center, operating system migration is fairly simple. Because filesystems are managed by the storage hardware and not the servers, migration is typically not even necessary.

If you purchase a new NAS system from the same vendor as its predecessor, you will typically have some options for replicating the data online on the new system.

Read On

Keep reading for more on storing data on solid-state drives and configuring your Linux clients to access data on Windows Active Directory servers. ■

Giants Have Their Own Rules

Web 2.0 platforms like MySpace, YouTube, and Second Life, as well as Internet giants like eBay, Google, or Yahoo, generate unimaginably large volumes of data. Multiple petabytes of data are the norm. For example, the Kodak EasyShare Gallery stores about 8 petabytes (8,000,000 gigabytes) of data. Dailymotion, a European MySpace competitor, has added an average of 1 terabyte a day since it started up in 2005.

In cloud computing, the data volumes are even larger. The tier 1 data for the Large Hadron Collider (LHC) at CERN is sent to Karlsruhe, Germany, for processing. The cluster computer that handles this has 16 petabytes of storage at its disposal.



Figure 1: A typical DAS system, the RDL-AS42S3, with 42 disk bays in a four-height-unit-tall case.