STORAGETEK®

# APPLICATION NOTE

# Mainframe Linux with Shared Virtual Array® disk system and SnapVantage™ software

## Commanding the Linux army of servers in the z/VM environment

APRIL 2004

1 Introduction . . . . 4

2 Virtual Linux . . . . 4

3 Virtual Linux and the SVA disk system . . . . 5

4 Creating an army of virtual servers . . . . 5

5 The SVA system and SnapShot software . . . . 6

6 SnapVantage software . . . . 6

7 Server deployment . . . . 7

8 Server management . . . . 7

9 Summary . . . . 7

## 1 Introduction — Why Linux? Why Linux on a mainframe?

In May of 2000, the first commercial Linux distributions were announced for the IBM zSeries (formerly called S/390) mainframe architecture. In just two short years since those announcements, approximately 15 percent of mainframe customers have implemented Linux on their mainframes and nearly half are considering such a move. In recognition of this demand for an open operating system on the traditionally closed zSeries platform, companies such as IBM, StorageTek®, Oracle, SAP, BMC and Computer Associates have been busy developing and porting applications for this exciting new environment. This paper describes the value of server consolidation onto the mainframe and expands that value to the total enterprise by offering increased savings and ease of infrastructure management. How? Match the virtual server consolidation to virtual disk capabilities.

The initial question asked by nearly everyone about the mainframe Linux environment is: Why? Why would anyone want to run an operating system, originally developed on a 386 PC by a college student, on a multi-million dollar mainframe? The answers are many and varied.

- **Linux is the fastest growing server operating system.** Since Linux emerged from Linus Torvald's PC in 1991, a legion of programmers have contributed millions of lines of code to the platform under continuous public review. The result is that today, Linux is an extremely robust and scalable operating system. Linux has been ported to everything from wristwatches to the world's largest mainframes and biggest clusters. The research firm of IDC reports that as of 2000, Linux had 26.9 percent of the new server installations, and projects that Linux will be the fastest growing server operating system (OS) through 2004.
- **Escalating cost of discrete servers.** As open and PC-based servers propagate through the data center, the costs of supporting discrete servers and their necessary infrastructure have increased at a faster rate than the growth in applications themselves. Even though the cost of server hardware has fallen, the overall operational costs of data centers full of discrete servers has risen, due to maintenance, power, cooling, floor space, network and personnel requirements of such an installation.
- **Difficulty of management of discrete servers.** As the number of servers grows, the tasks of managing the software, hardware and network configuration of an installation also grow progressively more difficult. Even with automated software tools, soft switches and rack-mounted hardware, managing operating system, middleware and hardware upgrades for legions of discrete servers is a daunting task.
- **Disaster recovery of discrete servers.** Disaster recovery is a long-practiced and finely tuned process in mainframe environments. Disaster recovery of a room full of discrete servers can be nearly impossible. Just the amount of time necessary to physically restore the processor and network complex for hundreds or thousands of discrete servers virtually rules out the traditional strategy of rapidly restoring a data center at a remote recovery site.
- **The zSeries hardware platform is the most reliable in the world.** The IBM zSeries hardware platform is one of the world's leaders in reliability and failover capability. This is due to its long development history and the capability and reliability of the processors and the fact that it has reliable hypervisor software.

## 2 Virtual Linux

The key to the adoption of Linux on mainframes is the z/VM operating system. Developed by IBM in the late 1960s, VM has the capability of completely emulating the zSeries hardware environment in software. An operating system can run under VM as a "guest" and VM will manage that operating system's interaction with the CPU, memory and peripheral devices such that the guest OS runs almost exactly as it would on the bare hardware. CPU, memory and network resources are dynamically allocated to the guest by z/VM based on availability and tuning requirements. z/VM can manage thousands or even tens of thousands of operating system guests on one zSeries processor complex. In fact, in a test environment, thousands of individual Linux images have been successfully booted under one z/VM on one zSeries processor. This "virtual Linux" environment leverages the strengths of the zSeries hardware to provide a platform for thousands of production Linux images to run simultaneously. A "virtual army" of Linux servers can be deployed on one zSeries processor running z/VM to replace an entire data center of discrete servers.

Some of the major benefits of running an army of virtual Linux servers over discrete servers are:

- **A virtual server is more scalable than a discrete server.** When a server is built for a specific application or group of applications, it is often difficult to project the exact load that server will bear. Sometimes administrators will "over spec" the server hardware to prevent the need to move the application to new hardware shortly after implementation. Or, even worse, the server hardware is under-provisioned to show a low cost up front, necessitating an emergency move to new hardware later. With a virtual environment, z/VM automatically load balances between the virtual servers. If a particular server sees a spike in demand, z/VM can allocate the necessary resources to accommodate the server. When demand drops, those resources can be reallocated to another server. z/VM does this automatically, without any administrator intervention.
- **All networking between individual servers on the same z/VM is done internally in memory, at memory speeds.** The only physical network cabling required is between the zSeries processor and the outside world. This drastically increases the security, availability and performance of the network. Internal networking can also be established between virtual servers and other mainframe operating systems (such as z/OS or VSE) running on the same processor complex.
- **The storage for virtual servers is provided on ECKD (mainframe protocol) DASD (disk) and tape.** Just as mainframe CPU complexes are the undisputed champions of reliability and redundancy, mainframe storage is the gold standard of reliability, failover and fast replication.
- **Virtual servers can be deployed in seconds, rather than days.** Using mainframe fast data replication features like StorageTek's SnapShot software, a virtual server can be created, booted and serving users in less than 30 seconds.
- **Virtual servers are easier to manage than dispersed discrete servers.** Hardware capacity upgrades can be done non-disruptively, one time. Software upgrades can be tested in virtual images identical to the production images and propagated into production in seconds. If a software upgrade causes a problem, the upgrade can be reversed in seconds.

## 3 Virtual Linux and the SVA disk system

The power of mainframe Linux is unleashed by the virtual technology of z/VM. Without z/VM, many of the mainframe advantages of scalability, reliability, resource sharing and rapid centralized server deployment would be impractical or impossible to supply to the Linux environment. The answer to the "impossibility" is found in StorageTek's Shared Virtual Array® (SVA™) family of products.

StorageTek, recognizing the power of virtual technology in the early 1990s, brought one of the world's first virtual disk devices to market as the Iceberg® 9200 library. The Iceberg library evolved into the IBM 9393 RAMAC Virtual Array (RVA), the StorageTek 9500 and V960 Shared Virtual Array systems and, most recently, the re-architected StorageTek V2X family. Today, the SVA is one of the few remaining disk storage devices that still innovate and extend the virtual architecture, bringing benefits of management to storage customers. The same scalability, reliability, sharing capabilities and rapid replication afforded by the virtual architecture of z/VM are applied to storage when the disk system is the SVA.

The V2X combines multiple RAID-6 disk arrays with eight microprocessors, front-end data compression engines and a large cache to provide up to 4,096 virtual disk images. The images can be either standard mainframe ECKD 3380/3390 DASD images, or Fibre Channel FBA architecture images. For mainframe Linux today, DASD images are the only choice. Virtual disk, virtual memory and virtual machine concepts are similar. The goals of virtualization are the same in all three cases: the separation of the user's logical view of a collection of data from the physical management of that data on a device.

## 4 Creating an army of virtual servers

An administrator building a "Linux army" will generally build several types of server images, one for each specific task or application to be served. Once these images are built, the required number of each server will be "cloned" to create unique server images. This cloning consists of creating the z/VM guest framework, copying the base Linux image to a new disk location and customizing the new copy to make it a unique server, capable of responding to a specific TCP/IP address. As business needs change and software evolves, the mix of server types and software levels will also change. This requires the administrator to periodically re-deploy the various servers. Redeployment requires backing up the current server, destroying or adjusting the z/VM guest framework, allocating fresh resources and re-cloning the image. The work of the administrator requires much resource management and uses precious system resources to accomplish the cloning operations. The virtual architecture of the SVA disk system provides a platform for minimizing the administrator and system resources necessary to provision and command a Linux army.

## 5 The SVA disk system and SnapShot software — The key to provisioning the Linux army

As we have seen, z/VM leverages its virtual architecture to provide hundreds or thousands of individual Linux servers the same view of processor resources that they would have if they were running alone on the processor. z/VM virtualizes the CPU, the memory and the network — and those are necessary elements in the building of the Linux army. But there is one piece missing: the virtualization of disk storage.

**The SVA disk system provides the missing element of true virtualization for Linux disk storage.**

The cloning operations necessary to build the Linux army are resource intensive. The copy operation takes CPU, control unit and channel interface resources. Once the copy is complete, you have used twice as much disk space. Creating an army of 1,000 servers requires 1,000 times the disk space of the original image, and hours or even days of CPU, channel and/or control unit processing resources. The unique virtual architecture of StorageTek's SVA disk system allows it to provide a powerful data replication function called SnapShot. SnapShot is a combined software/hardware feature that uses the characteristics of the virtual architecture to replicate disk storage almost instantaneously, while using no space on the disk arrays in the SVA and almost no control unit processor resources.

With the SVA SnapShot feature, a z/VM Linux disk partition can be replicated by simply copying its pointers in control unit memory. The physical data corresponding to the source of the SnapShot is not moved, copied or altered in any way. The target of the copy can be an unused portion of the same disk or another virtual disk, anywhere else in the SVA disk system.

The result of the SnapShot is that a complete replication of the source disk partition appears in another location on the SVA disk system. This copy is usable for both read and write operations. That the SnapShot image is a virtual copy is completely unknown to the user of the copy, just as virtual memory or virtual machines are completely transparent to the guest operating system in z/VM. At time zero of the replication, absolutely no additional space is needed on the disk arrays. As data is written to the copy, each updated track is written to a different location in the disk arrays, preserving the original track image for the original user (and other SnapShot copies).

**This replication takes place in milliseconds using absolutely no space on the disk arrays.**

Now imagine an administrator creating a 1,000-member Linux army. Instead of using 1,000 times the space of the original image and days of CPU/disk controller/channel resources on a traditional disk array, with the SVA disk system the army can be created in minutes, with the only additional disk space being those tracks modified to describe the unique network configuration of the servers.

## 6 SnapVantage software — Deploying and commanding the Linux Army

Even with an amazing replication feature like SnapShot, deploying thousands of servers can be a huge burden on system administrators. Once those servers are deployed, the ongoing control and administration functions are nearly impossible without automated tools. To assist administrators managing the virtual server environment, StorageTek has created a virtual server deployment and administration package called SnapVantage™ software. SnapVantage software leverages the strengths of the SVA disk system, SnapShot software and z/VM operating system to deploy and manage virtual Linux servers.

## 7 Server deployment

Deployment of virtual servers is usually done by cloning a pre-built image that has the software configuration and customization necessary for the server's role. Each server role requires a server with a different "personality." The tasks involved in cloning include:

- Creating the new z/VM guest user
- Allocating disk space to the guest
- Copying a pre-configured Linux system to the new guest's disk
- Customizing the network configuration

The basic value of the SVA SnapShot capability on the V2X in an environment of 4,000 virtual servers is obvious. However, if only a few servers are being created, it may not seem like much of an issue to do the cloning operation manually. But using the SVA SnapShot feature and SnapVantage software to automate the process gives the administrator the flexibility to use the cloning process to perform software upgrades and make personality changes to a virtual server at will. If a server has a role today that requires a SuSE Web server personality, that personality can be given to the server in under a minute. If next week it would be better for that server to have a Red Hat MySQL database server personality, the virtual server can be shut down, re-cloned and restarted in under a minute. This gives a data center unparalleled flexibility for providing the services demanded by their customers. Combined with z/VM and the SVA disk system's intrinsic load balancing capability, SnapVantage software allows a data center to be truly responsive to its customers' needs with a minimum of administrator and data center resources.

## 8 Server management

Ongoing administration tasks are performed on individual virtual Linux servers exactly as they would be performed on discrete servers. Boot/shutdown, user management, software maintenance and upgrade, log checking and disk management are all still required for the virtual Linux army. SnapVantage software is designed to assist in the management of Linux virtual servers under z/VM. From the Web-based SnapVantage software console, the necessary z/VM and Linux commands can be issued to control each virtual server. End-of-life or redeployment of a server is a matter of just a few clicks on the SnapVantage software user interface, or running a shell script that invokes the SnapVantage software command line interface. All z/VM and Linux command interfaces necessary to control the virtual Linux army are included in the SnapVantage software Web interface. Because the Web interface is so powerful, it is password-secured by z/VM. Extensive self-diagnostic and logging capability is built into SnapVantage software so that the administrator can concentrate on serving customers rather than on maintaining the administration tool. SnapVantage software is the only Linux administration tool that exploits the virtual power of z/VM, SnapShot software and the SVA disk system.

## 9 Summary — The armies march forward

The speed of adoption of virtual Linux into traditionally conservative mainframe data centers is unprecedented. The low-cost, open source nature of Linux encourages the leveraging of virtual technology to run a large number of servers with distinct personalities, rather than a small number of servers with multiple personalities. z/VM is perfectly suited to the task of virtualizing the CPU resources necessary to support server armies. Similarly, the SVA disk system is perfectly suited to the task of virtualizing the storage resources necessary to support server armies. SnapVantage software, enabled by the virtual technology of the SVA disk system and the power of SnapShot software, makes it easier to command these new armies. Together, there is realized a synergy that enables rapid deployment and centralized management of hundreds or thousands of servers with the lowest cost, smallest footprint and highest reliability in the world.

STORAGETEK®

## ABOUT STORAGETEK®

Storage Technology Corporation (NYSE: STK), a $2 billion worldwide company with headquarters in Louisville, CO, has been delivering a broad range of storage management solutions designed for IT professionals for over 30 years. StorageTek offers solutions that are easy to manage, integrate well with existing infrastructures and allow universal access to data across servers, media types and storage networks. StorageTek's practical and safe storage solutions for tape automation, disk storage systems and storage integration, coupled with a global services network, provide IT professionals with confidence and know-how to manage their entire storage management ecosystem today and in the future.

StorageTek products are available through a worldwide network. For more information, visit www.storagetek.com, or call 1.800.275.4785 or 01.303.673.2800.

## WORLD HEADQUARTERS

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
1.800.877.9220 or 01.303.673.5151



MV 9113 C  04/04