



RED HAT ENTERPRISE LINUX 5 BUSINESS CRITICAL DATA MANAGEMENT

DATA SERVICES TO RELY ON

Business continuity depends upon the ready availability of data, even in the face of equipment failures and other disaster scenarios. While the dollar cost of data recovery is high, it is dwarfed by the real cost of lost business. And in today's regulatory environment, failure to produce key information can threaten an organization's existence.

Are you ready when a server fails? While Storage Area Network (SAN) storage provides greater flexibility and performance, it can be subject to storage network outages from broken fibre channel cables, failed switches, and storage controllers. If a single mishap takes out an entire lab, can you access key information? Are you taking advantage of a SAN's full ability for disaster recovery protection?

And while unplanned events are often the focus, planned outages represent the majority of downtime. When can you schedule taking down that key array? You're forced to expect the unexpected while facing the common dilemma: How can I do more with less?

Add to this mix the promise of processor virtualization. You need a storage environment that fully supports and, even leverages the capability to strengthen your operational efficiencies and improve uptime.

Red Hat Enterprise Linux provides the solid foundation every enterprise solution can rely on to ensure high data availability and stellar performance. From great SAN support, to resilient volume management, to cluster file systems, Red Hat covers your business up and

down the storage stack.

MIRRORING FOR CONTINUOUS OPERATION WITH NO DATA LOSS

"We lost nearly two hours of orders when our storage system went offline. Those new storage arrays guaranteed continuous uptime. But they don't protect against every failure imaginable.

We can ride through single disk failures, even controller issues, but when the lab goes out, we need to replicate or mirror the data somewhere else. That's expensive and only works with arrays from the same vendor."

Whether your IT environment experiences a lost storage array or an entire lab, data mirroring provides continuous availability of your critical data. Unlike the recovery techniques of traditional backup, mirroring keeps your data available with no downtime and no lost information.

Using the host-based synchronous mirroring (RAID 1) feature of the Logical Volume Manager (LVM), each write is simultaneously and completely written to two or more local or SAN disks. In the event of an array failure, LVM will automatically detect the failure and use the identical, mirrored disk or LUN.




Illustration 1

Once a replacement LUN or disk has been brought on line, the recovery process proceeds in the background copying the data set over to the new disk unit. If the disruption to the original disk was temporary (a switch or cable failure) the recovery process will use a transaction log to rapidly replay IOs missed during the outage. Either way, recovery occurs while your application stays online and operational.

ARRAY BASED REPLICATION

- Synchronous

Red Hat Enterprise Linux also works well with enterprise-class storage arrays that offer optional remote replication capability. In this instance there is a direct connection between two or more storage arrays. For local and campus-based configurations this offers the equivalent of host based synchronous mirroring.




Illustration 2

- Asynchronous for disaster recovery

Federal or regulatory requirements may dictate that your critical business information be replicated at least 200 miles away. Given terrorist threats and unpredictable weather patterns, this may not be an unreasonable goal.




Illustration 3

Asynchronous replication can provide a reasonable

Unfortunately, this distance is too far for synchronous mirroring without introducing an unacceptable performance penalty, due to the high latency.

balance of protection and performance. This technique works like synchronous mirroring, except that the write to disk is considered complete once the local disk has been written to. The replicated copy will hit the remote volume shortly afterward. This key trade-off leaves a small window of vulnerability which can lead to loss of information.

CLOSING THE BACKUP WINDOW WITH VOLUME SNAPSHOTS

Though preservation technologies have become increasingly sophisticated, backup is still a key capability for the protection of data. Red Hat Enterprise Linux supports all major backup applications. Still, one key problem is finding a large enough time window to perform the process. Also, nightly backups simply aren't frequent enough. The volume snapshot support of the Logical Volume Manager (LVM) is able to close the backup window and make frequent backups easy and non-disruptive.

LVM volume snapshots create a point in time view of a volume, allowing your applications to continue while you make use of this older view. Creation of the snapshot volume is fast because LVM actively maintains a set of

"Our new data mining app is really proving valuable. We'd like to expand it to our retail projects. But the run will need about 3 hours a day and we're already up against the backup window. We've got maybe an hour of open time.

There must be a way to shorten that backup time."

block deltas between the active volume and the snapshot. As blocks are overwritten, copies of the original are maintained for the use by the snapshot. Once the snapshot is created, simply mount and use it as a normal file system.

Illustration 4 shows how snapshots can be used for backing up a database. First, the database is quiesced to ensure that a consistent state exists on the file system. The snapshot is created and the database is allowed to resume operation. Creation of the snapshot is nearly immediate. Next, the snapshot is mounted to a separate mount point



and the backup process is run. Once complete, the snapshot is simply removed.




Illustration 4

Beyond backups, snapshots can also be used for testing new applications against real data. A snapshot can be mounted read/write. The application can then be run against the actual data without risking the production data set.

DATA SHARING IN A VIRTUALIZED WORLD

Processor virtualization represents a paradigm shift in the configuration and management of IT resources. Finer grain isolation of applications within virtual machines requires a shared, high speed infrastructure. In a computing world where virtual instances are moved from one physical server to another, storage must be a common resource equally available to all physical systems hosting these mobile operating systems.

"The new virtualized CPU environment provides the isolation and protection I need, plus the ability to dial performance up and down. But I'm worried that we're creating silos of information and it won't be easy to share applications and common data sets. Isn't that going to lead to replicated information and synchronization issues?"

Not with Red Hat Global File System. Share any data you need to. It doesn't matter where the virtual OS instances are running. With your Storage Area Network, you have a completely shared file system infrastructure. Virtual machines and virtual storage.

For the Red Hat Enterprise Linux Advanced Platform we've included a powerful capability with our clustered file system, Red Hat Global File System. Based on over five years of commercial cluster file system deployments, the latest version of Global File System raises the performance bar and makes virtualized environments even easier and more flexible.




Illustration 5

Global File System provides a high speed, shared file system across all of your physical and virtual machines. This allows easy sharing of information and eliminates the need to store redundant copies of data. Simplified data management. Without a performance penalty.

In illustration 5 above, note how all virtualized instances have equal read/write access to the file systems across the SAN. Even as instances migrate from one server to another, the data stays fully available and fully shared.

Perhaps you are not sure whether you plan on utilizing the power of Global File System at initial setup. That's not a problem. Just create a file system for a single machine in the SAN and turn on the cluster sharing later. With Red Hat Enterprise Linux 5 we've greatly enhanced the performance of Global File System so there is no compromise of functionality versus performance.



MULTIPATHING SAN PROVIDES SPEED AND REDUNDANCY

“Who re-configured that storage network again? We lost 15 minutes of order processing while we tracked that down. SAN storage is great, but...”

With Red Hat Enterprise Linux storage management, there are ample redundant connections. With the Red Hat Device Mapper MultiPath IO, the traffic will just be rerouted. Plus, with active/active arrays, you can use all the paths for even higher throughput. And that's a standard part of the Red Hat solution for no additional cost.

Storage Area Network (SAN) storage provides a rich set of capabilities. It's the right solution for the enterprise delivering RAID storage for speed and protection against disk and controller failure. Still, SAN introduces additional complexity that must be accounted for to create a robust solution. A centerpiece of Red Hat Enterprise Linux storage management is Device Mapper MultiPath IO (MPIO). MPIO has the ability to connect and manage multiple paths through the SAN to the storage array. In the event of any single component failure, MPIO will automatically redirect data traffic via a redundant path.




Illustration 6

In illustration 6, the left hand picture shows two servers connected to storage via the correctly functioning red path. In the event of a failure of the Host Bus Adaptor in the server, cable, switch or array controller, MPIO will automatically detect and reroute IO traffic via the second path.

Further, many arrays systems support dual active controllers. This allows for IO traffic to travel down multiple paths, improving the performance of the system while protecting against component failure. In fact, you can configure more than two paths to further improve both performance and protection. Red Hat Enterprise Linux even supports multipathing across the boot partition to further ensure continued operation.




Illustration 7

THE RIGHT STORAGE SOLUTIONS, JUST IN TIME

More with less. You're asked to do it every day. You've got enough headaches. Red Hat Enterprise Linux 5 provides a complete storage solution for enterprise-grade applications. Red Hat Enterprise Linux Advanced Platform delivers a solid storage foundation through a storage stack that combines the performance, flexibility, and data integrity that you need.