RED HAT ENTERPRISE LINUX 5 CONTINUOUS AVAILABILITY

Businesses continuity needs to be at the heart of any enterprise IT deployment. Even a modest disruption in service is costly in terms of lost revenue and customer perception. Red Hat Enterprise Linux provides full support for fault tolerant hardware. Still, business critical systems demand high availability; outages must be rare and short.

Traditional techniques of redundancy through the full replication of equipment are expensive and limit their deployment to only the most critical applications. As a result, many important business applications are left vulnerable to long recovery times and lost information. Plus, not all downtime is unexpected. Planned downtime for maintenance operations can be equally disruptive.

The Red Hat Enterprise Linux Advanced Platform, powered by processor virtualization and High Availability Clustering, marshals reclaimed computational resources to greatly enhance availability at a price point affordable to the entire suite of applications.

Following is an overview for technical advisors and decision makers who need a response to the common dilemma, How can I do more with less? Red Hat Enterprise Linux 5 solutions are reviewed that improve SLAs through decreased failure rates, faster recovery times, and ways to minimize planned downtime.

UNDERSTANDING AVAILABILITY

A key goal for all IT operations is to keep system downtime rare and recovery fast. Keep in mind that maintenance downtime, that is planned downtime, must also be accounted for. Availability can then be described as the ratio of uptime, divided by the total time, namely uptime plus unplanned recovery time plus unplanned downtime.

$$\text{AVAILABILITY} = \frac{\text{UPTIME}}{\text{UPTIME} + \text{UNPLANNED RECOVERY TIME} + \text{PLANNED DOWNTIME}}$$

Achieving great availability requires that uptime be large relative to the remaining factors. Increasing uptime requires increasingly expensive systems with redundant components (power, failover memory, fans etc.) As an example, assume that a system takes four hours to fully recover by replacing hardware and recovering data. Notice that the uptime needs to improve tenfold for each additional '9' of availability; an expensive proposition.

UPTIME HOURS	RECOVERY MINS	AVAILABILITY
400	240	99.010%
4,000	240	99.900%
40,000	240	99.990%
400,000	240	99.999%

A much more cost effective mechanism for improving availability is to use the failover recovery mechanism of clustering utilizing software rather than hardware. Revisiting the above example, with a modest 400 hours uptime, observe the effect of lowering the recovery time.

UPTIME HOURS	RECOVERY MINS	AVAILABILITY
400	240.00	99.010%
400	24.00	99.900%
400	2.40	99.990%

The key here is the use of software plus the hardware you already have to dramatically lower the recovery time of an unplanned failure.

ACHIEVING FAST RECOVERY FOR HIGH AVAILABILITY

"I just got those estimates on the new fault tolerant equipment. Ouch! We need to keep things running, but I just can't get that kind of expenditure through the budget cycle."

Red Hat combines CPU virtualization with cluster technology, so you can meet your availability metrics and stay well within budget.

Virtualization enables Red Hat Enterprise Linux to harvest the excess capacity in IT environments. Clustering then uses a portion of this reclaimed capability to greatly improve the recovery time of the system.

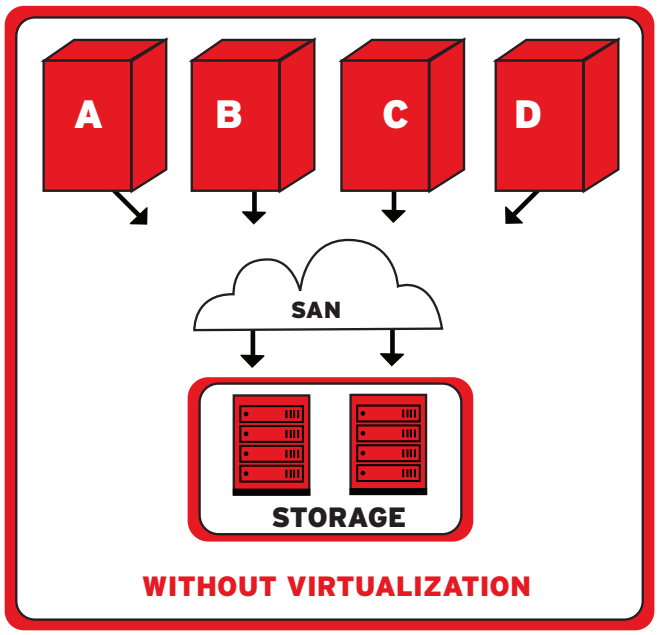

Illustration 1

Without virtualization, many organizations allocate unique, physical machines to each application. As a result, there is significant excess capacity on each of the machines. Industry data shows that utilization leaves 85% excess capacity (See Illustration 1). With virtualization, we can combine multiple applications onto a single machine while maintaining the desired software isolation.

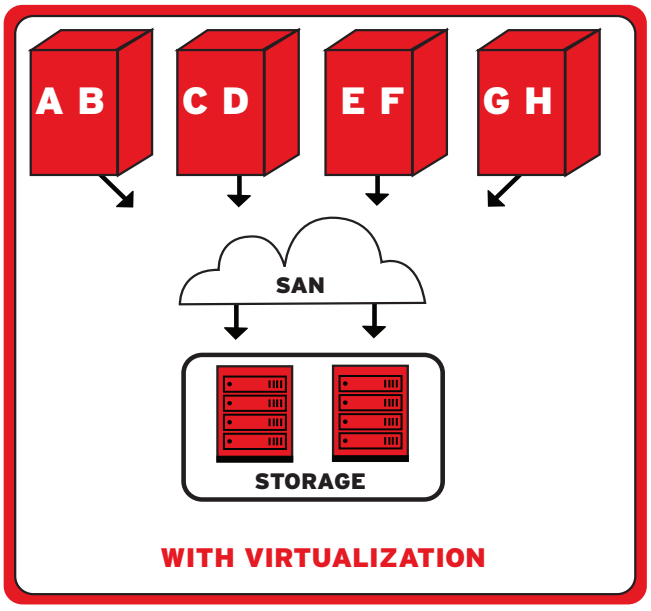

Illustration 2

Using virtualization, we can isolate each of the applications in their own virtual machine while doubling the number of running applications in with the previous configuration. Still, there is more we can do (See Illustration 2).

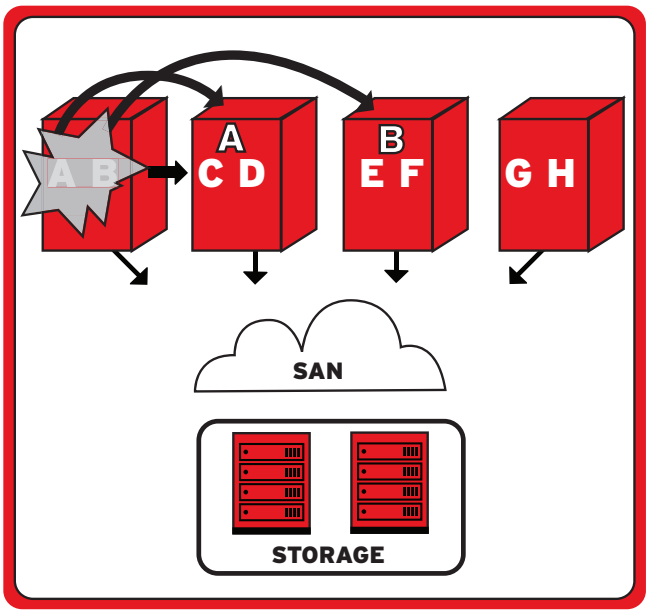

Illustration 3

With the clustering capability of the Red Hat Enterprise Linux Advanced Platform, we take a portion of the remaining capacity to power high availability. If one of the machines in the cluster fails, the virtual machines with the associated embedded applications are automatically restarted on a running machine in the cluster. The time to detect the machine failure and to restart on another machine is the recovery time described above and can generally occur in under a minute. Additionally, each of the virtual machines are started on separate nodes to better balance the load (See Illustration 3).

Once the failed machine is brought back on line, the system can then automatically utilize live migration to move the virtual machine and its hosted applications back to their original server. In this case there is no disruption in service operation (See Illustration 4).

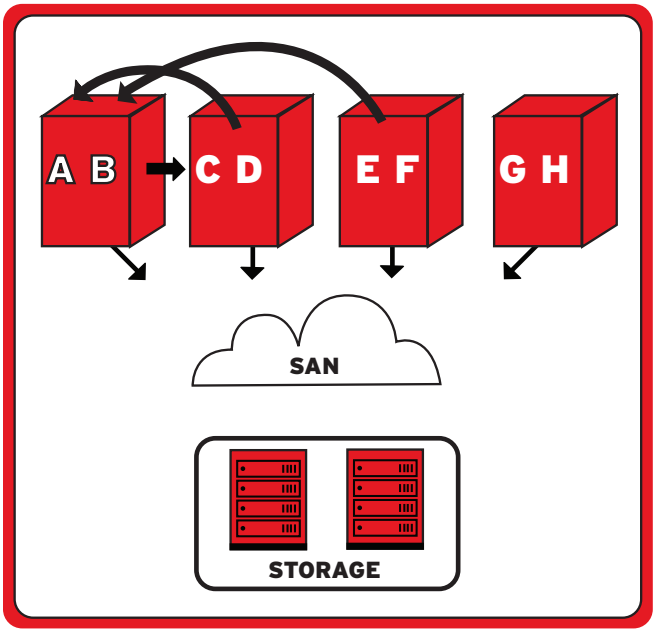

Illustration 4

TRADING PLANNED DOWNTIME FOR VIRTUALIZATION

"It's great that we're getting more memory for our data mining app. But, can't we wait a few weeks? I need that running 24x7."

It's never just the memory. Perhaps you have a fan that's out as well. Whatever the problem, you can't wait much longer or you'll be facing more than lost time.

While often overlooked, planned downtime can be as costly as the unplanned. Virtual machine migration provides the isolation your critical applications need to keep running even while swapping out failing components or upgrading equipment.

In this example (Illustration 5), applications A and B are migrated while active, from machine 1 to machines 2 and 3, balancing the load out. Note how GFS, our shared SAN filesystem, provides a fully shared data set that carries through the migration. Maintenance can then freely be applied. In fact an entire machine can be swapped out. This maintenance can include a significant upgrade to the equipment, including extra processors and more memory. No rush to finish is needed, because there has been no business interruption. You can take the time to run full diagnostics to ensure a successful operation. Plus this can all take place during normal business hours and at standard rates. Once complete, your applications can then be migrated back.

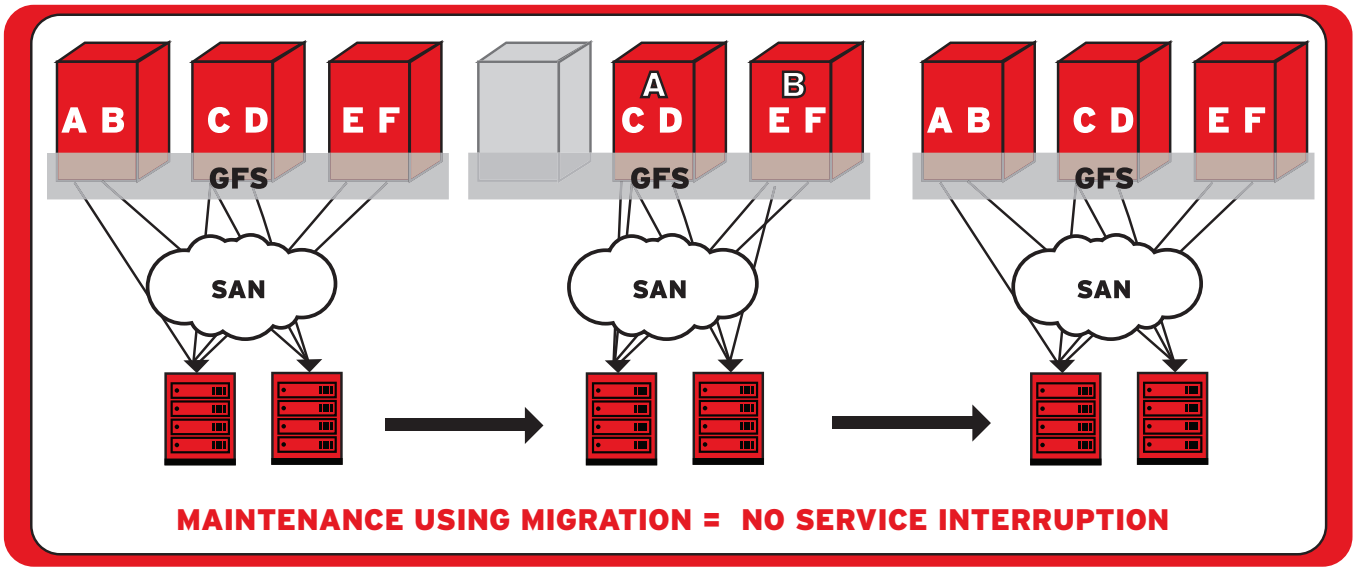

Illustration 5

MULTIPATHING KEEPS YOUR SAN RUNNING

"Who reconfigured 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 arrays, you can use all the paths for even higher throughput.

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.

In illustration 6, the left hand picture shows two servers connected to storage via the correctly functioning green 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 insure continued operation.

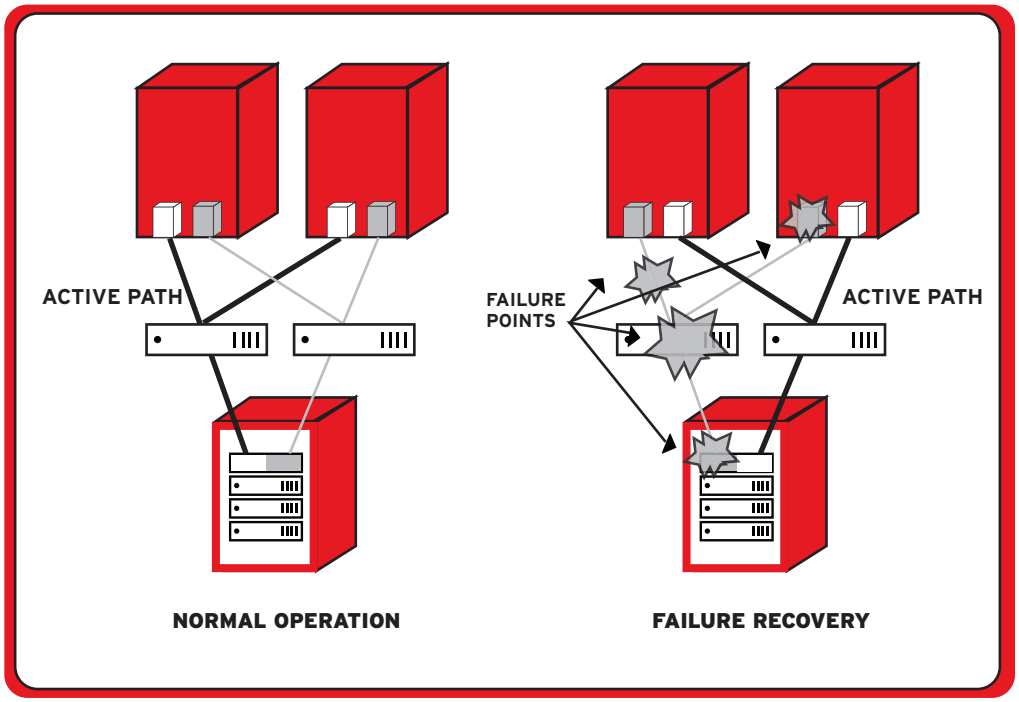

Illustration 6

STILL HIGHER PERFORMANCE AND RESILIENCY

"So, Red Hat has a really nice set of high availability services. But, we're building some custom apps. I'd love to plug into their clustering APIs. We could get even more resiliency with zero information loss if a machine goes out."

The Distributed Lock Manager is fast and has been used by other customers in prior releases. Now we've added the ability to share state and messages across the cluster.

In addition to the support of failover production of cluster unaware applications, Red Hat Enterprise Linux Advanced Platform provides a powerful set of libraries for the creation of cluster aware applications. Utilizing the Distributed Lock Manager and OpenAIS libraries, customers can create applications that span virtual and physical machines maintaining uninterrupted service through system faults.

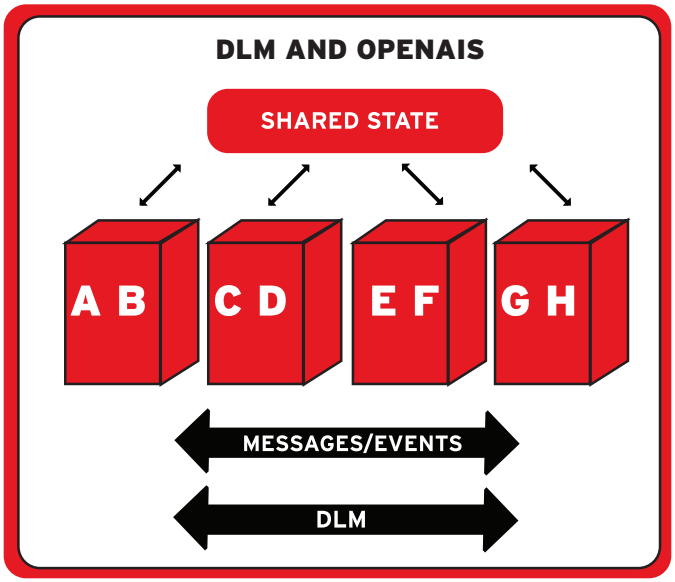

Illustration 7

DISTRIBUTED LOCK MANAGER

The Distributed Lock Manager is a high speed mechanism for resource synchronization across one or more machines. This allows applications to coordinate their activities and process information in parallel, using as many virtual and physical machines as needed to solve a problem.

OPENAIS

OpenAIS implements the SAForum specification that was born of the need for TELCO-grade applications to operate across a cluster of machines in demanding, zero downtime environments. With this powerful service, applications can share state using checkpointing, as well as coordinate activities through eventing and messaging APIs.

CONCLUSION

Red Hat Enterprise Linux Advanced Platform is the solution for enterprise grade applications. It delivers a solid foundation for hosting business-critical applications. The powerful combination of integrated CPU virtualization, service failover, cluster file system, and cluster application APIs provide an easy, powerful, and inexpensive way to achieve an always available environment. Red Hat enables you to do more with less.