

# Symantec Live Migration Transitions With State

## Quick Note

Thomas Deane  
16 August 2005

Copyright © 2005  
Illuminata, Inc. Licensed to  
Symantec Corporation for  
web posting. Do not  
reproduce.

It's hard to argue with a goal like uninterrupted, seamless service through any sort of failure or upgrade. Users in such an ideal environment would remain blithely unaware and uncaring about any outage or behind-the-scenes recovery. However, the search for such bliss has long seemed akin to the quest for the Holy Grail—a challenging journey with no real end in sight.

In the Unix world, clustering has been one common route for the Grail seekers. But even the most advanced clustering implementations—such as those originally developed by Digital Equipment for OpenVMS and Tru64—were complex, relatively specialized, and they never *completely* slew some real technical dragons.

Process migration is one such dragon. A process is the fundamental unit of work in many operating systems, including Unix. It

is a container of sorts that encapsulates program code along with current state, security permissions, and accessible resources. However, it's a very low-level container that tends to be deeply rooted in the system on which it was born.<sup>1</sup> This makes a thorny problem of moving a process to a different system to recover from a failure or to rebalance a workload.

Given that the on-the-fly migration of processes is a formidable quest, many vendors have pursued other, more attainable—but still beneficial—goals. One such is “Stateful Transition”<sup>2</sup> (ST), the movement of active applications and users from one up-and-running system to another without state or data loss. Meiosys, recently purchased by IBM, is one example of the ST approach.<sup>3</sup> Live Migration from Symantec, by way of Ejacent and VERITAS, is slated to be another.

<sup>1</sup> That's because it's closely tied to specific physical resources like shared memory.

<sup>2</sup> Some vendors use the term Stateful Migration instead to describe similar technology.

<sup>3</sup> See <http://www.illuminata.com/perspectives?p=45>



## From Eja-cent to VERITAS to Symantec

Live Migration's underlying technology comes from a product called UpScale developed by Eja-cent, a company that VERITAS purchased last year. This year, Symantec, in turn, acquired VERITAS. Currently, Symantec is in the process of retooling the technology to work with Version 5.0 of VERITAS Cluster Server (VCS).<sup>4</sup>

The release date of Live Migration has not been totally locked down, but it is likely to be around the time that VCS 5.0 is available—the middle of 2006. Unlike some other VCS options<sup>5</sup> which are chargeable, Symantec plans for Live Migration to be a bundled, user-enabled feature that will serve as a higher-level application state transition tool to complement the low-level failover capabilities that are inherent in its clustering software.

As currently described, Live Migration will move state between VCS cluster nodes by encapsulating the entire context of a running application into a container. The container will include the state of the application and its associated processes. It will also include any temporary files used by these processes and the network state of the application. By so doing, it will allow an administrator to move applications and users from one system to another with little or no interruption and with no loss of state or data—and without the need to quiesce or otherwise shutdown the application.

Symantec's plans for Live Migration call for it to support only Sun's Solaris in its first release and to focus on Oracle database environments specifically; this combination of Solaris and Oracle is a Symantec sweet spot today. Support for other OSs, including IBM's AIX 5L and Linux, is on the drawing board for the future, but with neither firm commitments nor schedules.

<sup>4</sup> Even though Symantec has bought VERITAS, this product's name will remain *VERITAS Cluster Server*.

<sup>5</sup> For example, "VERITAS Enterprise Agent for Oracle"

## Why ST?

"Stateful Transition" technology, such as that in Live Migration, can help an administrator handle impending system failures or planned downtime less disruptively. It can also simplify workload management. And, while doing these things, it preserves application network connections to databases. Let's discuss each of these benefits in more detail.

System failures don't always happen out of the blue. Many remote system management and monitoring tools, such as the HP Systems Insight Manager<sup>6</sup> and IBM Director, can proactively detect certain impending system failures, such as a failing memory DIMM or disk drive, through a variety of soft-error thresholding and other techniques. In these cases, ST would let an administrator move the applications and users from the troubled—but still breathing—system to another, stable system without losing data, state, or context.

The benefits are similar for an administrator who needs to take a system down for *planned* system maintenance. System maintenance can be handled during normal business hours because applications and users can be moved at-will and without disruption.

ST is also useful when an administrator wants to dynamically balance workloads across multiple systems. It lets system utilization be optimized in real-time by moving active users from one system to another server based on changing workload demand—again without disruption. Such flexibility is one (but just one) of the pieces needed to fulfill the promises of grids and utility computing environments in which workloads are moved dynamically from system to system as user demands change.

ST's complementary ability to maintain persistent connections from applications to back-end databases is particularly important and desirable when moving applications and users in the application server tier. If these connections are lost, nasty things often occur. For example, database transac-

<sup>6</sup> See Illuminata report "HP Systems Insight Manager: One Console to Rule Them All" (April 2005).

tions may have to be recovered by rolling back transactions and restarting. Even the application or database itself may need to be completely restarted. Data in flight could even be permanently lost. By preserving these connections, ST brings the movement of applications with persistent connections much closer to being a reality in a production environment.

### Limits of Live Migration

Of course, like all the other approaches to task, process, and job migration, ST isn't a panacea.

For example, it doesn't, in and of itself, provide failover if there's a catastrophic system failure. Thus, you wouldn't use Live Migration to protect a critical always-up database; a cluster or even a virtuoso fault tolerant system like HP's Integrity NonStop is the right solution there. It's better thought of as a tool for preventive maintenance—whether to deal with a system that might fail soon, a system that needs planned maintenance, or simply to redistribute application loads.

Like other forms of container and virtual machine technology, there's also the issue of whether an ISV will support software running within a container. If a problem occurs when running the application in a Live Migration container, it might not be clear if the problem was in the application, database, system, or Live Migration itself. And, in any case, the vendor who wrote the application may decide as a matter of standard practice to only provide support when running on a bare OS. However, VERITAS routinely dealt with such support issues, as have other vendors such as VMware.<sup>7</sup>

Live Migration also requires a common data store. Network Attached Storage (NAS) or a storage area network (SAN)—or some combination thereof—is needed for this to work, as is a supporting storage

network infrastructure. However, centralized storage architectures are necessary for all manner of data protection and application migration architectures; they're hardly a unique requirement for Live Migration, nor at all unusual in modern datacenters.

Finally, as with any new technology, users and admins need to accept the new technology and understand where it fits—and where it doesn't. They need to understand what tasks it simplifies and the sorts of failures it protects against. ST can clearly be useful, but primarily in the context of managing running systems rather than dealing with dead ones. At the very least, it must be complemented with other technologies that more directly address riding through unexpected system failures—as Symantec is doing with Live Migration by integrating it with VCS.

### Alternatives

ST is one type of virtualization. Other companies<sup>8</sup> have taken different virtualization approaches aimed at different system layers. Some of these are complementary. In fact, since Live Migration will run on Solaris, it could complement Solaris Containers (which virtualize the OS) by letting application containers be created within a container.<sup>9</sup>

Application containers could, in principle, also complement system virtualization at the hardware level, such as VMware's virtual machines (VM) provide. However, in practice, VMware is aggressively fielding its own take on stateful transitions. Its VMotion add-on to ESX Server and its VirtualCenter management tool allow active VMs to be moved "on the fly" between two systems running VMware's ESX without bringing down the system, the OS, or the application. What's more, an administrator can archive VM snapshots for later recovery from a failure or, in these days of increased

<sup>7</sup> For example, VMware has often had to demonstrate that application or OS problems that appeared within its virtual machines were likewise present running on bare iron. Virtualization (and VMware's flavor, in particular) is much more widely accepted today, but this support issue was once a significant headache for VMware.

<sup>8</sup> VMware, SWsoft, Meiosys, Trigence, SUN to name a few.

<sup>9</sup> Although, given how lightweight Solaris Containers are, it's perhaps unlikely that such mixing would be common occurrence. See Illuminata reports "Solaris Risks" (November 2004) and Illuminata report "New Containers for New Times" (January 2005).

governmental regulations, to satisfy a compliance request. VMware's approach of migrating and storing entire OS images is very different technologically from Live Migration's application containers—but it addresses many of the same user needs, and it does so today.

### **Conclusion**

Symantec hasn't yet released all the details of its still-in-development Live Migration. However, its cornerstone will clearly be its ability to perform stateful transitions, a technology approach that virtualizes the application by placing it and its state into a container.

VERITAS Cluster Server has the failover which Live Migration lacks. However, by adding the ability to move applications and users around without loss of state, Live Migration will make VCS better able to handle a number of sticky system management issues—including handling planned outages and load balancing workloads.

While Live Migration provides mastery over some technical dragons, it is not the Holy Grail. But the marriage of Live Migration and VCS is perhaps, for Symantec, a match made in heaven.



Through subscription research, advisory services, speaking engagements, strategic planning, product selection assistance, and custom research, Illuminata helps enterprises and service providers establish successful information technology.