Levanta Marries Provisioning with Data Virtualization

Quick Note

Thomas Deane
July 8, 2005

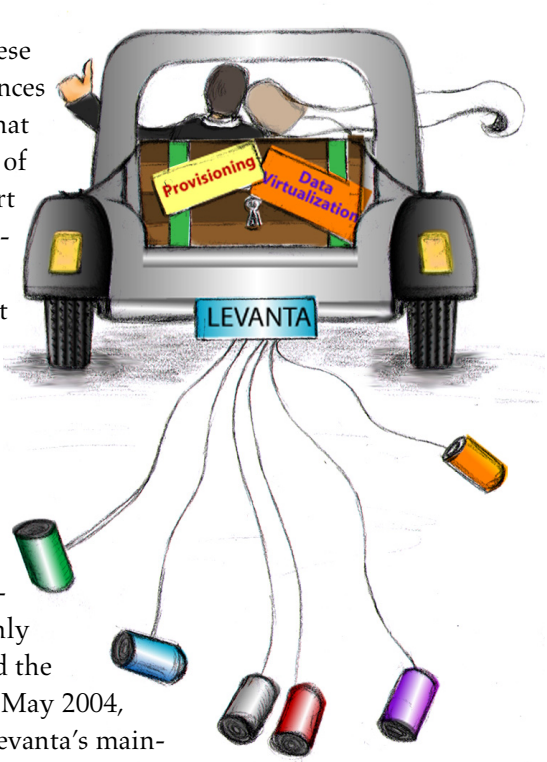

So many products do system provisioning these days that it's hard to keep track of the differences—both claimed and real—between them. If that weren't enough, there's another (smaller) set of products that do data virtualization of one sort or another. Levanta stands out because it marries these approaches in order to enable the rapid provisioning and on-going management of scale-out computing—from 10s or 100s to 1,000s of Linux systems.

Levanta began life[1] offering outsourced Linux development, deployment, and system management. When, amidst the dot-com meltdown, that business model didn't work, the company refocused from services to products. Levanta, its provisioning engine for highly virtualized IBM[2] zSeries mainframes, codified the company's Linux know-how into product. In May 2004, Levanta shipped Release 3, which extended Levanta's mainframe-inspired capabilities to scale-out x86 environments. The latest product addition is the start of a set of plug-and-play appliance servers that make it easier to adopt and deploy Levanta provisioning.

Let the Wedding Bells Ring

There are many script- and image-based provisioning systems in the market. They work by replicating operating system, middleware, and application code onto the direct-attached storage (DAS) of each and every managed server. This approach, taken by the likes of Altiris, HP's Remote Deployment Pack,[3] and VERITAS' OpForce, can dramatically simplify the installation and updating of software across many machines. It is, however, somewhat brute-force. It creates many distinct system instances, each of which, once installed, must be managed and updated quasi-independently.[4]

1 As LinuxCare

2 Levanta hovered very close to the IBM mother-ship. It is an IBM partner and was recommended by IBM to its customers.

3 Which is also Altiris based.

4 Too be sure, they provide tools for posting common updates to multiple servers. They

Levanta works differently. It uses a form of data virtualization to simplify provisioning. Levanta's MapFS virtual file system creates an illusion: each system *thinks* it has its own copy of software. In reality, however, all of the software is located on a shared data store that's accessed across the network. Because of this, Levanta can provision a system much faster, with much less bursty network traffic, than copy-everything-local methods. The software stack lives in a central location, so great gobs of data do not have to fly hither and thither when a system is provisioned.

All files are initially kept in a Read-Only Repository.[5] When a system attempts to modify one of these files, MapFS performs a "copy on write"[6] and replicates a virgin copy of the file into the Overlay —a Read/Write storage repository that contains a private, independent area for each system's modified files. If a system has not modified a file, MapFS will direct a file request from a system to the Read-Only Repository. On the other hand, if the file has been modified, MapFS will direct the system to the Overlay, as shown in Figure 1.

Although Levanta's core functions depend on MapFS, systems can also independently access direct-attached storage (DAS), other network-attached storage (NAS), and Fibre Channel storage area network (SAN) storage if needed or desired. Levanta recommends DAS for temporary files—swap space, for example—to improve performance. Other non-local (NAS and SAN) storage pools also often contain user data and databases.

MapFS maintains a Smart Mirror to facilitate system checkpoints/snapshots and rollbacks. When a snapshot is taken, only a recording of the file deltas between the Read Only version (in the Read-Only Repository) and the Read/Write version (in the Overlay) are made. Because of this, less data must be saved than if the files themselves needed to be recorded in their entirety.

Working hand-in-hand with the Smart Mirror, the Overlay also archives files changes. At any time—including a rollback—a file exists in either a modified or unmodified state. If unmodified, a MapFS rollback only has to read the file from Read-Only Repository. If modified, however, the Smart Mirror just instructs the system to load the Overlay's archived version.

Is there a cost to this virtualization? Some. What were local disk I/Os become network I/Os, and there's some overhead in the management of the "copy on write" operations. But networked storage is now widely deployed, and any incremental overhead associated with MapFS operation is trivial in comparison to many virtualization products.[7]

Figure 1 – The MapFS File System

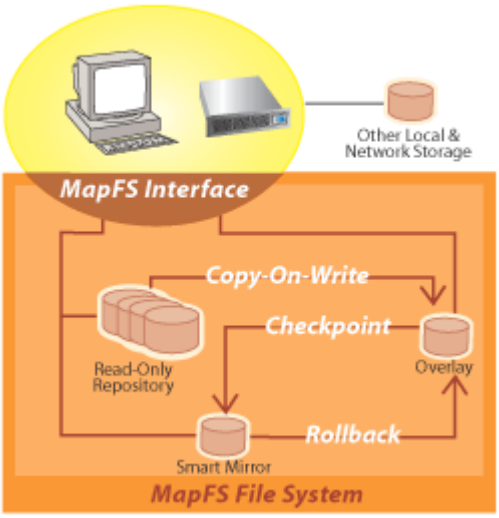

Templates and Recipes

Baking wedding cakes requires both ingredients and recipes. Think of the common read-only software repository as the pantry where all the ingredients are kept, and templates as a recipe book. Levanta's templates are metadata files describing a server's components—hardware, storage, and software. Hardware includes the processor, memory,

become much less effective, however, if servers cannot be updated in relative lock-step.

5 It is read-only from the point of view of managed systems. Administrators can, of course, update it as necessary.

6 A technique more commonly found in virtual memory systems.

7 Virtual machine overhead can easily run 20 percent. See Illuminata report "VMware on the March" (Aug 2004).

and network parts. Storage includes the storage space allocation and volume definitions. Software includes the software components needed for a particular variant of server or workstation.

Levanta provides a number of standard templates. As an example, one specifies how to turn Fedora Linux and the Apache HTTP server into an operational web server. Home-grown templates/recipes can also be created as needed.

A Happy Union

Levanta's capabilities go beyond the imaging, scripting, and control/monitoring found in many other provisioning systems. Whereas most competitors resort to make-a-copy-of-everything imaging methodology,[8] Levanta only has to record the changes made for a specific system or template. These changes are deltas from the Levanta-controlled, common, read-only, centrally-located software base. The vast majority of any OS and application stack is read-only; consequently, these deltas are invariably smaller than the entire software stack, so snapshots execute rapidly. Rollbacks are equally speedy because only the deltas need be restored.

The ability to quickly perform snapshots and rollbacks simplifies software versioning. New versions of software can be quickly installed and tested—and quickly rolled back if needed. Because Levanta-managed systems can share a single copy of the OS and application stack, and therefore run exactly the same versions of software[9] across the entire software stack, many software versioning (and version consistency) problems simply—*poof!*—disappear.

Wedding Day Jitters

To be sure, there are some trade-offs to keep in mind about automated provisioning systems. One is that they become effective only when the number of systems to be managed is fairly large: tens at least. For just a few systems, a sophisticated provisioning infrastructure is not worth the effort.

And while automated provisioning can substantially reduce the management costs and TCO for large data centers, installing the provisioning system itself can be a hill to climb. There is no magic pixie dust that configures "just want we want." Beyond product cost, it requires time and effort to learn the tool and create appropriate plans and templates. System administrators thus have to buy into the chosen tools in a big way.

Automated provisioning also dictates more commonality than many shops have been used to. Although differences among systems do not have to be *entirely* eliminated, the goal is minimizing one-offs—and more tightly managing them when they do occur. Levanta works best in an environment where many systems are quite similar, and ideally the same. For example, a data center filled with uniform racks or blade servers would be perfect. But Levanta also accommodates environments with dissimilar systems. Indeed, one of its biggest advantages over script- or image-based approaches is how well Levanta works in diverse, heterogeneous environments.[10]

Additionally, the Levanta software requires a supportive data center infrastructure. Although local disk can be used for temporary and personal file systems, operating systems and applications live across the network. This mandates fast, robust networks. For this reason, Levanta requires either SAN or Gigabit Ethernet (GbE) connectivity, or a combination, depending on the customer's preference and performance needs.

If SANs are used, each system needs an HBA, and the environment needs SAN switches. If GbE is

---

[8] Not to mention an "update every copy" patching technique.

[9] Levanta does not preclude someone from running a different version of any software—it handles these cases quite well. But in general usage, these cases should be the exception rather than the norm.

[10] Many Levanta customers exhibit considerable diversity. For example, one customer, a large financial institution, has 87 unique, non-cloned zVMs on a single mainframe. Electronic Arts has 6-10 different online gaming stacks as well as two different sets of blades, each requiring different underlying x86 Linux kernels. These all access the same binaries in the Read-Only Repository.

used, each system needs a GbE NIC and the environment needs GbE switches.[11] Existing environments lacking such equipment would need to be upgraded.

Because of Levanta's use of storage/data virtualization, a common data store is also required. In large enterprises, this requires separate SAN or iSCSI storage devices. In small/medium enterprises (SME), a network-accessible common data store is required.

Finally, Levanta is a Linux-only solution. Windows servers need not apply.

Like clusters and other sophisticated infrastructure, provisioning is often not necessarily easy to deploy—especially the first time. Provisioning engines need to be customized, either by trained, local people, or by bringing in a team of external experts through services. And Levanta has some specific infrastructure needs on top of the requirements for provisioning systems, in general.

Appliances: The Honeymoon Package

Realizing the various adoption barriers, Levanta now offers a packaged provisioning solution. Like honeymoon packages and other vacation plans that have many of the details already worked out, Levanta's Intrepid provisioning appliance makes it easier to figure out, buy, and use the company's provisioning technology. Levanta sees appliances at the core of its future growth.

The Intrepid M (M for manager) targets small and medium enterprises (SMEs), as well as large departments and workgroups. Two other appliances are slated to launch later this year. These are the Intrepid B (B for backup), and the Intrepid S (S for storage), a supporting NFS-based NAS appliance.

The Intrepid M is a low-cost entry point appliance that includes internal storage accessible by the other systems over GbE. Starting at 1.4 terabytes of storage space, it can manage up to 128 servers and workstations. It includes templates for various system types.

Although Intrepid M disks are internally RAID-protected, a single Intrepid M doesn't provide redundancy at the server level. Thus the Intrepid B. Intrepid M and B combine to form a highly available pair. This is crucial because Levanta holds all the eggs. If the data store supporting MapFS[12] goes down, so do all managed systems—*Gulp!*

For its SME target audience, Intrepid offers a known-price solution up front. Unlike most service-led engagements, no blank check is required.

Levanta is also readying a higher-end, enterprise-focused provisioning appliance. Based on high-performance Fibre Channel SANs and iSCSI shared storage already in the datacenter,[13] this device (code-named "LM3000") will scale to manage more servers than the Intrepid M.

To keep things simple, appliances are less flexible than a custom service-led engagement. Levanta's flagship Release 3 software, however, continues to be available for more complex, customized, and larger-scale deployments.

And They Lived Happily Ever After

Levanta does provisioning—and it does it in a unique way. It does so by expanding beyond the traditional "load all *this* software onto *those* servers!" approach. Its MapFS file system creates the illusion of separate data for each system, but in truth centralizes the storage and management of the software stack. This marriage of data virtualization with provisioning supports the rapid provisioning of large numbers of Linux systems.

---

[11] With GbE, Levanta recommends a dedicated storage network for best performance.

[12] MapFS isn't the data store itself. Specifically, it's the Linux kernel-loadable interface module that resides on the clients and brokers requests to the MapFS data store. If the MapFS interface goes belly up on a client system, only that client system will be affected. But if the data store supporting MapFS goes down, all clients are affected.

[13] Many of Levanta's enterprise customers have only a subset of their Linux systems hooked up to SANs. They would likely use the "LM3000" for those systems, while using the Intrepid M for the remainder.

Relatively rapid provisioning is not unique to Levanta. It can also be done with image- and script-based tools. Its ability to provision a system in under five minutes—*rapid* even among other "rapid" tools—is particularly desirable in a dynamic, on-demand environments. But initial speed is not its biggest win.

Levanta is particularly graceful at managing further updates to each system, including quick snapshots and rollbacks. Levanta manages software versions efficiently because most, if not all, of the software stack is centrally located on a network share, and managed (through MapFS) according to changes that occur. While relative homogeneity and simplicity is still the goal of provisioned environments, Levanta degrades gracefully when the reality of system diversity sets in. As a result, Levanta provides a more elegant and effective way of performing software installation, snapshots, rollbacks, and versioning than many competing solutions.

The adoption of any automated provisioning system, and no less Levanta, requires a robust infrastructure and some thoughtful planning. Levanta's new appliances, however, can make this hill a bit easier to climb. Overall, Levanta's wedding of system provisioning and data virtualization takes a significant step closer to the agility and ease promised by on-demand computing.