

Project Live*

> *Radically simplifying software deployment, installation, and configuration*

One of the great things about today's consumer electronic devices — cell phones, Blu-Ray players, wireless routers, plasma TVs, etc. — is that they provide very sophisticated functionality without requiring you to install any software. All of the software is built into the device's firmware. Everything just runs, right out of the box, and you can usually personalize specific features very easily.

That's not the case with computer systems. Whether you're setting up your new PC or bringing up new servers in an enterprise datacenter, software installation is a tedious, time-consuming, frustrating process. Every application needs to be installed and configured individually, and managing all of the patches and updates and configuration files makes the process a perpetual time drain.

Researchers at Sun Labs asked a simple question: why not bring the simplicity of the firmware model to datacenter systems?

Sun Labs launched Project Live* [live-star] to explore that question. Led by Principal Investigator Olaf Manczak, the Project Live* team has devised — and is now prepared to share with the open source community — a new approach that has the potential to radically simplify software distribution and configuration in the datacenter.

Live* technology distributes major software components as immutable disk images that contain pre-installed, ready-to-run software. Users can simply pick and click to create fully functional software appliances. Here's a closer look at how it works — and what it means.

Bridging the Gap between Ease of Use and Customizability

It's easy to understand why different software installation models evolved for consumer electronics gizmos and datacenter systems.

Consumer electronics products typically are designed to provide very specific functionality and require very little customization. They're usually produced in high volumes by a single vendor that has tight control over the software that is deployed.

Datacenter systems tend to be multi-purpose systems that run software from multiple vendors. They are produced in relatively low volumes and require a high degree of customization, introducing more complexity and more potential software compatibility issues.

The challenge for the Sun Labs team was to find a way to blend the ease-of-use and reproducibility of the consumer electronics model with the modularity and customizability of the datacenter model. And according to Mr. Manczak, finding a solution was urgent.

"The complexity of software installation and configuration is already a big problem for datacenter administrators, and it's going to get bigger if we can't find a better approach," said Mr. Manczak. "The number of physical datacenter servers deployed is currently growing at a rate of about 15 percent per year, and with the increasing popularity of server virtualization the total number of installations is actually growing much faster than that. We don't need small incremental improvements to current processes; we need to re-think the way software is deployed."

Overcoming the Limitations of Traditional Tools

A number of tools and techniques are currently available to streamline various aspects of software installation, configuration, and management. For example:

- Software packaging technologies such as AdminStudio from Acresto (formerly Macrovision) simplify installation
- Management systems such as Sun™ xVM Ops Center or HP OpenView make it easier to manage virtualized installations
- “Live CD” distros boot up an operating system via CD without requiring installation or affecting previously installed files

These solutions are very useful in increasing administrative productivity but do not fully address the software deployment and maintenance problem.

Block-level packaging technologies, for example, help with initial deployment but still require updates to be applied individually on each system. Management systems help automate routine maintenance tasks but only for one application at a time. And LiveCD distros give users the flexibility to mix and match various components, but cannot yet cope with structured data such as system personalization files.

“The traditional model still leaves far too much burden on administrators and end users,” he said. “What we’re trying to accomplish with Project Live* is to push the burden of software installation back to the factory—while still enabling users to easily customize their software appliance.”

“Pick-and-Click” Software Appliances

Project Live* introduces a technology that gets beyond the limitations of all previous technologies. With Live*, major software components such as operating systems, databases, and application servers are distributed as immutable disk images. Each image contains pre-installed software that is ready to run. Users simply select the combination of software components they want. Like Lego blocks, the components can be assembled into complex systems, or software appliances.

Want to run MySQL™ on OpenSolaris™ and do backup and recovery with Zmanda Recovery Manager? All of those components could be made available in their most current versions as disk images. Users would simply pick and click among those and other components to create the specific software appliance they want.

At boot time, Live* technology creates a virtual file system that dynamically combines the software from multiple images. This process is entirely transparent to existing software, and the content of the virtual file system is identical to the result of a traditional installation.

Through this approach, systems can be upgraded simply by replacing entire images with updated versions of that image. And through management systems such as Sun xVM Ops Center, this process can be automated so that multiple servers could be updated through a single command.

Secret Sauce: Virtualization with Personalization

A unique capability delivered by Live* technology—the capability that overcomes the limitations of Live CD—is the ability to virtualize the software (as a block-level snapshot) and handle both structured and unstructured data. This gives users the ability to install the software they need and maintain their personalization settings even when upgrading or patching the software.

To do this, Live* takes advantage of a technique called “file system unification” that is part of the Live CD concept. Essentially, applications see a combined file system that looks like a writeable CD-ROM, but there are actually two file systems—one that provides access to existing files on the CD, one that handles all the modifications, such as writes, deletes, and renames.

With this technique, it becomes possible to change/upgrade software dynamically; you simply replace the image with the new one and the software appliance still runs exactly as you want it.

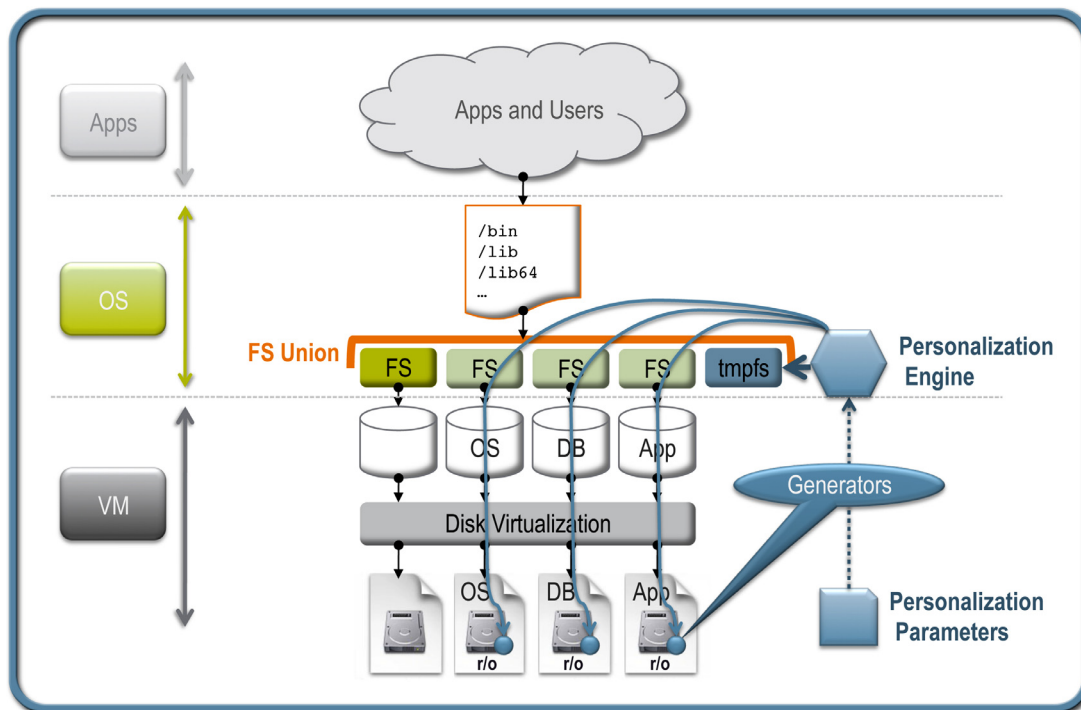
To deal with the structured data, Live* technology goes a step beyond the Live CD approach and incorporates “personalization engines”—programmable infrastructure for generating configuration files.

Unlike management tools such as Sun xVM Ops Center or VMware, the personalization engines have no knowledge

up front of how to deal with all the configuration files. Instead, the disk images have a set of “generators” that can be used to transcribe the personalization parameters—all the answers you normally give to software installers—into actual configuration files. It creates an in-memory empty file system, and uses the personalization engine to apply all the generators and create actual configuration files that go in the empty memory file. And when this is all done, you end up with something that looks like a single writable disk that has all the software installed in the right place with all the configuration files already there.

“It’s like genes in a DNA structure,” says Mr. Manczak. “No one vendor has access to the full structure, but each vendor can extract needed parameters. At early boot time, the personalization engine goes through all the images needed to create a software appliance, identifies the generators and informs them. The output is the generation of configuration files, which go into the in-memory file system. So you end up with an immutable image with both software and personalization information. Now you can make changes to the structured or unstructured data and they’re preserved.”

One of the huge benefits of the Live* approach is that it becomes possible to do operating system updates quickly. You no longer have to install a patch or security update on every machine; you just get the latest image, re-start the virtual machine, and everything runs correctly.



Live* virtual appliance

Learn More—and Participate

Combined with Xen or VMware hardware virtualization, Live* technology provides a one-click way create robust, secure, and reproducible software appliances.

Sun Labs urges the development community to take a closer look at Live* technology—and help expand its capabilities. The Project Live* team is making the core technology accessible to the open source community via GPL 2.0 licensing. For more information or to download the Live* software, go to http://research.sun.com/projects/live* . To see a demo, visit <http://www.youtube.com/watch?v=rwFYKsvEJ3g>.

About Sun Labs

Established in 1990, Sun Microsystems Laboratories is the applied research and advanced development arm of Sun Microsystems, Inc., with locations in California and Massachusetts. Sun Labs is one of the ways Sun invests in the future, and is responsible for many of the technology advancements that have made Sun a technology powerhouse—including asynchronous and high-speed circuits, optical interconnects, 3rd-generation Web technologies, sensors, network scaling and Java technologies.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN Web sun.com

