;login:
APRIL 2013 VOL. 38, NO. 2
ELECTRONIC SUPPLEMENT

Conference Reports

LISA '12: Real World Configuration Management Workshop

San Diego, CA
December 9, 2012

Real World Configuration Management

Summarized by Nick Anderson (nick@cmdln.org) and Aleksey Tsalolikhin (aleksey@verticalsysadmin.com)

The LISA 2012 Real World Configuration Management Workshop was chaired by Narayan Desai (Argonne National Laboratories), Cory Lueninghoener (Los Alamos National Laboratory), and Kent Skaar (VMware). Thirty-seven people attended.

Narayan, Kent, and Cory opened the day with introductions. The 37 attendees introduced themselves, each sharing a pain point with the group. Introductions consumed the entire morning but there was consensus that it was beneficial and worthwhile. Culture and Secrets dominated this discussion.

Culturally, acceptance of configuration management and automation tools is still not universal. Common roadblocks include perceived time constraints (takes too long and or too difficult to use the tool), and low business value by small or isolated groups. Fear and distrust of automation was also noted as an acceptance issue. Although most people were already using configuration management tools, it is not uncommon for tools to be bypassed intentionally. These manual changes are not always ported back into the configuration policy. Several people indicated they were not running their CM tools continuously, and for those who do, running in noop/dry-run/warn_only modes is not unusual. For configuration policy activation to be an infrequent and even manual action is relatively common.

Multiple attendees cited managing security domains as an issue. Questions were raised including how to manage "many hands" of different skill and trust levels, how to divide policy and grant access only to authorized individuals, how to manage secrets in policy, how to separate data from policy, and how to share common policy between disparate environments (typically government). To manage people, some are using ACLs (Access Control Lists) provided by a version control system, others are using approver-based gating that encourages peer review of policy. Many people are tying custom systems together or have custom tooling built around their configuration management systems. Separating data from policy varies widely. Sneakernet is still required to deal with disparate environments. Many attendees said that a unified view, or single source of information would be preferable to the many sources currently used.

Orchestration and performance concerns about CM tools and their ability to scale and manage complexity rounded out other common pain points. Orchestrating policy deployment is burdensome. Attendees voiced a desire for staged and slow controlled dispersion of policy, and there was discussion about the importance of promoting policy based on its own stability and having it roll out automatically. Integrating with other tools, namely monitoring systems, is difficult or at the least, not straightforward. It is not uncommon to perform management of monitoring and inventory systems separately from a CM tool. Mark Burgess pointed out that having CM and Monitoring separate is legacy thinking as modern "continuous maintenance" CM systems are continuously monitoring the systems on which they run. OS patch management continues to be painful. Some sites don't patch at all, or just roll out new OS images periodically. Orchestration of larger systems brought discussion of performance overhead from configuration management tools as well as talk of how different tools themselves manage scalability (push vs pull, centralized vs decentralized).

After returning from lunch and wrapping up introductions, attendees were asked what they would fix if they could snap their fingers and have it done. The list included: complete buy-in from customers and IT counterparts, more separation of policy and data (with easier policy metadata extraction), better tooling around version control to simplify workflows, easy discovery of possibly conflicting policies, and scope of impact of a policy change. Automatic risk-aware policy deployment over a period of days with workflows that include teaching coworkers what is happening was a common desire. Mark Burgess wished for a shift from deployment steam roller mentality to comprehensive design thinking as well as more reuse of existing parts instead of reinventing the wheel.

We kept notes on which configuration management tools attendees are using. People are primarily using one of the major frameworks (BCFG 2, CFEngine 2, CFEngine 3, Chef, Puppet), but homegrown systems are not uncommon, and new ones are still being built (e.g., Ansible, Cdist, and SaltStack).

During the final segment of the workshop, informal statistics were collected. A show of hands using the "never have I ever" (even a little bit counts) model was used to survey the attendees. Most startlingly, there were only two attendees with official QA processes for their configuration management. The weirdest things under configuration management included laser cutters (inside a 3D printer), Android phones, Raspberry Pi, routers, switches, robots, and a QNAP storage appliance.

The workshop concluded with discussion of the future of configuration management. The main concerns were enabling system orchestration (state transitions with dependencies—for example, an admin could tell systems X to transition from state A to state B, but to wait until systems Y reach a certain state first); separating development and production (Paul Krizak explained how he builds a system of dev VMs simulating production, with Jenkins used to test the change); and network configuration management (Tom Limoncelli spoke about OpenFlow, which is a way to centrally control routers).