



Announcer You're listening to the Sun Microsystem's Podcast Network. Welcome to another edition of *Innovating@Sun* with your host Hal Stern. Today's topic: *Open high-availability cluster*. And now here's Hal Stern.

Hal:

Hello and welcome to another edition of Innovating at Sun. I'm your host Hal Stern, Vice President of Global Systems Engineering, and with me today is Keith White who is the Director of Availability Engineering inside the Solaris Software Engineering organization. So, Keith, welcome to the show.

Keith:

Thank you very much. Great to be here.

Hal:

So you have a fun job in that you've designed the software that keeps things running online as much as it possibly can. So why don't you tell us a little bit what it is you do.

Keith:

Well, fortunately, I have a team of extremely intelligent people who helped the design of that software and it's not just me. But we put out availability products, in particular, the Solaris cluster and Solaris cluster geographic edition and the agents. I've been at this company for over 18 years. It's a great place to be. And I've been involved in all aspects of software development. We've just recently released the Sun Cluster and geographic edition 3.2 release, which is the largest feature release in over five years probably. So a great feature set we've just come out with very recently. So it's design and development of those products that I'm involved with here.

Hal:

And just to be specific, when you talk about the high-availability cluster, this is the software that lets machines work in an active rack of standby configuration, and basically it's failover software but not, I would say, the way we thought of it 10 or 15 years ago which was some simple script, sort of ping machines. Or rather, a robust agent framework that you can go plug a variety of application monitoring into.

Keith:

That's exactly right. I mean, this isn't just a restart mechanism. This is actually understanding the hardware, the interconnects, the network, doing heartbeats and making sure that things are actually alive and extremely fast in-kernel failure detection for either failover or application handing off to another one which is already active. So like you said, active standby including disaster recovery which a worldwide solution to problems that may come up in a particular data center or environment like that. So it's not your mom's HA of old days. This is very sophisticated technology.

Hal:

At the same point, I think that the fact that it is the foundation for application availability means that it touches an awful lot of the ecosystem of developers and partners outside of Sun.

Keith:

Sure. The primary piece of this is keeping your services alive and available where you want them. So it interfaces to all the ISVs or even home-grown applications where you need data or the service available 24 x 7. Everybody in the community, I think, and probably most people in the world touch this type of technology every day.

Hal:

Certainly with the amount of emphasis and interest right now in web-next services and horizontal scale, lots of those small servers, many people seem to be focusing on replication as opposed to the vertical scale in the failover, but do you see that this goes away? Or is this just something that we'll probably take for granted over time?

Keith:

It probably has been somewhat taken for granted in that we've actually become so successful at this technology, but the systems that you're talking about in horizontal scale, the kind of availability we're talking about there is significantly different in that it's not failures they're concerned about, but more about capacity on demand, which is a very, very different type of availability. So the failures happen and that's where you need to protect the data and the services from outages, either hardware, software or any networking or like type of problem.

Hal:

Given that, we have clustering software that is really the heart of what, I'd say, most online services. If it's not at the edge of them, it's somewhere behind there in terms of running one of those transactional or stateful services. The ability to have developers target that platform clearly is important and probably gets us to the heart of what we really should talk about which is the fact that you're going open source with this product set.

Keith:

That's absolutely right. So on June 27th, we're actually announcing the open sourcing of what we consider the first commercially successful high-availability application in the world. So we're open sourcing Solaris Cluster. We've created a community on opensolaris.org called HA Clusters. We're seeding that community with the open source from the Solaris Cluster product under the name of open high-availability cluster, and we're providing a binary distribution of Solaris Cluster that runs on open Solaris.

Hal:

So the first part of this alpha shoot into the hands of the developer community is going to be the agent framework.

Keith:

Right. We're actually open sourcing this in three phases, so we're starting with the agents themselves, starting that right now. We're going then in December to open source the geographic edition, our disaster recovery product. And in December of '08, we're going to have the entire core infrastructure completely open sourced, and that point, the whole product set and our product development will continue as open source completely.

Hal:

Hearing of this rather than waiting 'til all the encumbrances are removed and everything's ready to go, for you to get out the agents first, in talking with some of the folks on the open source marketing side, they're quick to point out that we have people who are developers and deployers in open source and people who are co-developers. And I would say system administrators, for as long as I can remember when I was one, I think, have always been kind

of in the co-developer camps. They're always people who are tooling, who are building tools or building on the good work of others. So making this available seems that we're addressing a system administrator development community right off the bat, and probably the people who would have the greatest interest in seeing how the agents work.

Keith:

And that's - when we sat back and thought about, what is the purpose of doing this open source? And who are we trying to benefit? We thought immediately that it doesn't make sense to just target a developer for a developer's sake. We're really looking at the end users and consumers of this technology to make sure that they get the benefit first. And that's why we open source the agents themselves so that you get the insight into, you know, a very specific bit of knowledge from the people who developed a platform about how we created agents. And the ISVs also get to look at that so that they can create highly available agents for their applications, hopefully expanding their markets into places they hadn't been yet. So you get the tips and tricks that we've done internally. We're looking at the sys-admin developers, slash, ISV community to expand that portfolio, to give everybody what they need in the end before we move on to other parts of open sourcing.

Hal:

As we've seen more of our users deploy open source software for a time-to-market advantage - they want to just go rapidly deploy what's available and build services on top of it, the need to make those services available, the need to be always on has increased. I think there's been probably a lack of conversation between the application developers and the application deployers, again, the system administration, systems management crowd. And having a common, open-source set of tools here that both can look at, I hope, stimulates some better conversation about how you make an application highly available, what it needs to do, what it needs to provide, how you have to instrument it so it fits into a framework like this.

Keith:

Yeah, that's exactly right. We're seeding the community with not only the software, but also the testing infrastructure so you can actually test the things that you've done are functional in this environment, but we're putting actually people into this environment as well. So you get access to the experts in the field who can tell you what is right, what is wrong and where they've been and what they've learned so that you can also grow as a developer to make sure that this type of community builds successfully upon itself.

Hal:

You said something at the opening that what you have with the high-availability product is really one of the largest, most commercially viable and successfully software products in this space. Is that the big deal in terms of open sourcing it?

Keith:

Certainly one of the big deals is that we're going with the first commercially successful product in open sourcing. There's other things that are unique about this in that we're using a CDDL license just like Open Solaris so that it's a very open license. It doesn't require a bunch from anybody who tries to contribute or put more back into their environment or our environment. We're also open sourcing the test code so you have that infrastructure available to test highly available scenarios that maybe you're not aware of, or you get the documentation as open source so that you can contribute to any parts of that like a regular open-source community. But we're also reinforcing Sun's commitment to open source, but not just open sourcing something that we don't make money on, because Sun does actually make a significant amount of money on this product. So it's taking that risk there. We're continuing to drive change into the marketplace, so we get a lot of competition where, you know, you make

claims or they make claims, but do you really know that these things are true? So we're driving an openness into the environment where HA is really important. So this is not the stuff you want to mess with. IT and people like that get a window into the technology that they wouldn't have gotten otherwise to validate whether claims we make are true or untrue by looking at the source code itself. And this is a very specialized field, so it's not just random source code that, you know, you can pick up in a week and understand everything there is to know about it. But it's a highly specialized field with information, IP as well as the guys and gals that do this day-to-day.

Hal:

And what you said about open sourcing the test tools and the test environment, I think, is absolutely critical. The hardest part of making something highly available is not writing the code in the first place or design the framework; it's making sure that it does what you want it to do when the sweat is pouring off of your forehead because something's broken. And the only way you're going to get there is if you test it really thoroughly. And testing is one of those black arts. People that do it well, really think of all the corner cases and the odd stuff and providing as much of a starting point there as we can. That's to make the adoption easier for the people who are going to be building on top of the framework that you provide.

Keith:

Well, it's certainly allows them to be a lot more confident in what they're delivering as agents for their solutions. It's not something you want to discover at 3 a.m. on a Sunday morning that something didn't quite work the way you expected. So this infrastructure does cover a significant amount of corner cases that people wouldn't normally think of unless this is the world you live in. And, of course, we've been living this for many years at this point and have a lot of understanding of what it takes to make something highly available when the chips are down.

Hal:

Okay. In terms of a call to action, what you want people to do after the first drop of the open source products is available, what's your next steps you want to offer up to the developer community?

Keith:

Go to opensolaris.org/os/community/ha-clusters. Get the how-to guides and the getting-started guides. Understand what it is that we're offering there. If you're an ISV or sys-admin developer type and want to contribute your agent for the applications that you're interested in, please contribute that back. They'll be available on the community site. An interesting tidbit here that we haven't touched on yet is that any HA agent that's built on top of the open source platform, will actually run on a commercially available product, namely, Solaris Cluster 3.2. So you get the benefits that your application is now available on a supported distribution. Become part of that community. Contribute. Understand where things are. Make your applications available. Go to blogs.sun.com/sc to talk more about it or understand more directly.

Hal:

Great. Well, Keith, I'd like to thank you for giving us an update on where we are with open sourcing the high-availability framework and really giving us a view of where we're going to see this open-source project progress over the next year. Again, that's for making life interesting for developers and also keeping everything the developers do up and running on the web. So you've been listening to another edition of Innovating at Sun, and I'm your host Hal Stern.

Announcer You've been listening to *Innovating@Sun*. Join us next time for the latest in innovation from Sun Microsystems, only on the Sun Microsystems Podcast Network.

<http://blogs.sun.com/innovation>