

Sun Expert Exchange

“Fast Track to Solaris 10 Adoption: Predictive Self-Healing”

Wednesday, August 18, 2004

Q: What is the mechanism that PSH is using to detect an application crash? Will it be able to automatically upload crash dump to Sun and initiate a service call to tech services?

Angel Camacho (A): This involves a daemon (fmd) that collects information on events inside the system that updates an event list that is "published" for use for the other agents inside the system, and so on. Regarding the Service call, this is an ongoing process with Services and monitoring agents, our plan is to leverage some customized agents to do so. Please follow up at <http://www.sun.com/software/solaris/10/>

Q: Are there any guides on how developers can take advantage of PSH

Michael Shapiro (A): The major way we want developers to take advantage of PSH is by writing an XML manifest for the Service Management Facility (SMF), the second part of PSH, which will be delivered in the next Solaris Express release. As part of that release, there will be tons of SMF documentation made available.

Q: So upgrading to Solaris 10 is a sure way to help the Red Sox win. Count me in!

Michael Shapiro (A): I agree!

Q: Is work being done to send the alerts from PSH to the Sun Field Support to handle CPU/memory replacements?

Smita Thakur (A): Yes we are working on making sure that there is a good connection with the remote facilities.

Q: Will the agent dev tool for integration be another d-type of language or scripting capable?

Michael Shapiro (A): We are working on a PSH module, which will allow administrators to run a custom script when an automated diagnosis occurs.

Q: How will or can PSH interact with Sun or Veritas Cluster services?

Cindi McGuire (A): SunCluster and Veritas can register to receive detailed fault information from cluster nodes to make decisions about resource failover and management and use SMF to restart node-based services. Additionally, the PSH architecture can be leveraged to create a consistent administrative model and experience for a single node or a cluster of nodes.

Q: I like the concept surrounding PSH and the dev of APIs enhancing better logic in fault isolation. What's next?

Michael Shapiro (A): In the short term, our immediate goals are to port everything we've done so far to x86/AMD. Going forward, we want to both build a more complete profile of "system-level" objects (e.g. harden more driver) and then move up the stack to higher level abstractions (e.g. file systems, network, Java).

Q: Are trap messages the only SNMP 'article' produced? - Is there a history/log that can be polled at the MIB level?

Michael Shapiro (A): We're in the process of developing the PSH->SNMP connections, so this is a great topic for giving us more feedback to we can meet your requirements. If you have time, go to sun.com/bigadmin/content/selfheal and describe more about what you would like to see.

Q: I see PSH handled memory and CPU potential problems. Are there any plans for predicting disk failures, perhaps through the use of SMART technology built into many drives?

Cindi McGuire (A): Yes! We are planning to harvest SMART data and create an error telemetry that allows us to predict disk failures.

Q: Where can more information be found on SMF? A quick search of the main Sun site didn't find anything interesting.

Smita Thakur (A): Please check out: <http://blogs.sun.com/sch>

Q: So SMF is not in Beta 6, but will be in Beta 7?

A: Yes, it is scheduled to be in Beta 7.

Q: What is SMF in full name?

Smita Thakur (A): Service Management Facility

Q: is sun positioning PSH for future integrations that could compete with traditional NMS tool suites

Michael Shapiro (A): We're actively exploring connections to higher-level network management software. Stay tuned to PSH announcements for more information.

Q: Will there be a relationship between PSH and Solstice DiskSuite or will we continue to have to rely upon mdlogd?

Cindi McGuire (A): There are plans to instrument SVM and our future file systems for PSH. Migration of tools such as mdlogd will occur at that time.

Q: Does the white paper touch on configurable behaviors and instantiating child processes of them?

Brian Ellefritz (A): It does touch on configurability but doesn't go as far as "how to" levels of detail.

Q: Will PSH data be available through kstat?

Michael Shapiro (A): Kstats are somewhat orthogonal in that they are bean counters measuring numbers of certain kinds of events in the kernel. We update these in addition to sending telemetry for PSH when errors occur. And then PSH itself has its own bean counters for statistical purposes: fmstat(1M) lets you view these.

Q: Given PSH supports systems as individuals but will it handle RAID levels with customizations?

Cindi McGuire (A): Integration of PSH outside of a single system is planned. We are looking to bring PSH technology out to the network, RAID and fabric-based storage.

Q: Is access to the SUNW-MSG-ID limited to the internet or will it be hosted locally with the FCM?

Michael Shapiro (A): We're considering providing a way for customers to cache the information locally and customize it at their site. Stay tuned to future PSH announcements.

Q: What sort of tools will be available to configure PSH?

Michael Shapiro (A): Most of PSH is self-configuring: the only things that we want you to tune are things like custom actions to take when a diagnosis occurs (e.g. e-mail your pager or something like that). Then the Service Management Facility (SMF), which will manage application services, has many configurable behaviors for administrators -- stay tuned to Solaris Express to learn more.

Q: How will PSH handle fan failures or CPU overheats?

Smita Thakur (A): CPU Overheats: Yes Fan Failures: Yes, however dependent on the platform

Q: Does self-describing data imply that there are XML schemas?

Cindi McGuire (A): XML can be used to marshal our self-describing error and fault protocol event data.

Q: How can third party applications use PSH features. Is there any way of checking third party software logs for errors?

Michael Shapiro (A): Yes. The major way application software vendors can plug into PSH is using the Service Management Facility (SMF), coming in the next Solaris Express release. By defining an SMF manifest, a service will be automatically restarted upon failure (be it software bug, administrator error, or hardware failure) and it will be given an individual service log (in addition to any logging features that the application developer provides).

Q: Why isn't Sun back porting this to Sol9 and 8? Not everyone will be moving to Sol 10 on [Sun's] timeline, but will still benefit from this technology in their existing Data Centers.

Smita Thakur (A): The intention is to have Sun's customers take advantage of not only this key technology in Solaris 10, but also other ground breaking technologies in that release. These technologies demanded a radical change in the kernel and hence it will take a big effort to backport them Solaris 8 & 9.

Q: How does PSH respond to Oracle in a clustered environment? - ex: Oracle runs out of swap space.

Cindi McGuire (A): Oracle, like other application or middleware software, may register as an SMF service and/or receive detailed fault information. If a fault should occur on any of the resources it uses, Oracle may respond by releasing its hold on the affected resource and failing over to a duplicate or restart its services.

Q: Are PSH capabilities applied system wide or is there a finer granularity (ex. Solaris 10 zones / grid containers)?

Michael Shapiro (A): Actually the granularity of retiring faulty resources is much finer than even zones: we can in Solaris 10 offline faulty CPUs, individual physical pages of memory, I/O devices, and kill processes and restart the affected service. And this works in local zones as well as global zone.

Q: When will be the first training or update for Solaris 10?

Brian Ellefritz (A): Training will begin rolling out around our ship date at the end of 2004

Q: This is what I got when I selected download software: That is why I am asking exactly how to get the software. Any help would be greatly appreciated.

A: Sorry for the confusion. In short, there are two flavors of the Solaris Express program. If you want to receive support, you get the "Commercial Use" version and pay \$99. If you just wish to begin testing Solaris 10, we'd recommend the free version. At: <http://www.sun.com/software/solaris/solaris-express/get.html> Things might be a bit more clear. You can get CD images for either SPARC or x86, Commercial or Non-Commercial use.

Q: I see in the documentation that fmdump will tell you the part number that needs to be replaced. Is this true for third party components, or strictly Sun supplied ones?

Michael Shapiro (A): We will tell you the location path for the FRU and that will work regardless of whether, for example, the PCI card there is from Sun or another vendor. Some other information, like the part number and serial number, require that the part have that information encoded in it in a standard form (e.g. Sun FRUID) that we can read. You'll need to check with the vendor to see if they support that.

Q: Will api calls be available to tie into NMS tools such as openview

Angel Camacho (A): Our plan is to leverage customized agents to integrate those functions, please visit <http://www.sun.com/software/solaris/10/> in the near future for up to date information.

Q: Will PSH be able to interact across zones? For example, offlining one zone and reallocating a nic card to replace one that failed in a different (more critical) zone?

Michael Shapiro (A): This is exactly the type of functionality we will be able to build by connecting PSH with Solaris's IPMP feature and the ability to export virtualized network interfaces into zones. We don't have this in Solaris 10 yet, but as we convert our networking subsystems to PSH we will be able to do that and that is exactly the type of thing we want to be able to deliver to you.

Q: Will JES installers be "fully" integrated with Solaris 10 installers?

Smita Thakur (A): The focus of this event is Predictive Self Healing. Later we will have a Solaris 10 event where we will take up such questions. Stay tuned for the Solaris 10 event.

Q: How does PSH interact with Containers to help prevent downtime? Can a hardware failure be transparent to a container?

Michael Shapiro (A): PSH interacts with containers in that we try to isolate errors to a user process if possible and restart its containing service. If we can't do that but we can isolate the problem to a zone (container) then it can be restarted. Hardware failures are "transparent" to a container in that containers typically depend on virtualized resources, such as a pool of CPUs or a filesystem. Depending on the failure mode of the underlying resource and how that manifests through the virtualized resource exported to the container, that problem may be "visible" or not. Finally, the diagnosis results and suggested repair actions communicated to syslog, for example, are always transparent to containers -- those are only logged to the global zone for the system administrator, and are not seen by users in the local zones.

Q: How does PFH relate to SNMP?

Angel Camacho (A): SNMP is a way to produce traps that Monitoring software use, and Predictive Self Healing is a framework inside the Operating System that trigger actions and corrective procedures, you can find more information at <http://www.sun.com/msg>

Q: Is the sun acquired cluster technology going to be leveraged for PSH to applications? What is SMF?

Cindi McGuire (A): Not yet, but there are plans to move PSH technology up our software stack. This includes middleware software such as Java, N1, Cluster, SunCluster, etc.. SMF is the Service Management Facility. There is a technical whitepaper that describes PSH and SMF available at <http://www.sun.com/bigadmin>.

Q: My server had a panic on CPU2 write back check error. Should I say that Solaris 10 PSH will save the server from this type of error?

Michael Shapiro (A): I'd need the complete error message with the context of where we detected the error to tell you whether it would be recoverable or not, but yes, PSH would have automatically diagnosed this problem for you. Errors such as the one you describe now produce automated telemetry events to be diagnosed by PSH, and we've made a continuous effort across our Solaris 8 and 9 patches and in Solaris 10 to harden Solaris against all such errors to the degree permitted by the hardware.

Q: will you be publishing more articles on tip and tricks for PSH

Brian Ellefritz (A): We're working on a more consolidated location for emerging resources for PSH; for now we'll be updating the "bigadmin" site at : <http://www.sun.com/bigadmin/content/selfheal/>

Q: Is there a way to view PSH if one is visiting Palo Alto?

Smita Thakur (A): Please visit Sun's Executive Briefing Center in Menlo Park, and you can get a demo on Predictive Self Healing.

Q: Can PSH send traps to a monitoring solution when there is a fault?

Angel Camacho (A): Our plan is to leverage custom agents to do it, please check often our site for new information. <http://www.sun.com/msg>

Q: Will the SUNW-MSG-ID be server centric with updates

Cindi McGuire (A): As appropriate, the message ID will be platform specific. For example, a fault message that is

specific to a SunFire 6900 will contain a message ID that is unique to that platform. The message and its ID will direct the admin to platform-specific response and repair actions.

Q: Will this be included in the Solaris open source distribution?

Michael Shapiro (A): Yes, Sun's intent is to make Solaris open source, and that would include the source for PSH features.

mark (Q): I am unable still find Solaris 10 at the web site. You get to download and the only thing out there is Solaris 8 and Solaris 9. What does one have to do exactly to get the software ? Thanks

Brian Ellefritz (A): At the bottom of this screen you can click the link "Get Solaris 10 Now!", or use this URL to go to the Solaris Express program site: <http://www.sun.com/software/solaris/solaris-express/>

Q: are PSH features intended to be used by app developers?

Angel Camacho (A): Absolutely, if you are using it to take corrective actions on your development, or taking advantage of the full set of API's there is something you can use if you are a developer.

Q: How is PSH integrating with Monitoring tools/agents (like Sun Cluster) , HP Openview, and Tivoli. Is there an SNMP agent to send alerts to a Control Station?

Michael Shapiro (A): We're working with all our ISVs to take advantage of the new APIs and technology offered by PSH. We will be providing a module in the Solaris 10 timeframe to permit SNMP traps to be sent triggered when an automated diagnosis occurs.

Q: What are the filesystem features if any, fir backups in Solaris 10? E.g. getting list of changed file since a given time

Smita Thakur (A): Please attend the Expert Exchange Session on 9/15/04 for the session on Dynamic File System

Q: is there a good out-of-the-box overview for S10 admin, presenting both new and old tools?

Brian Ellefritz (A): There are a number of Solaris 10 admin resources being posted to the "bigadmin" site at :<http://www.sun.com/bigadmin/products/sol10.html>

Q: Will this provide the ability to dynamically reconfigure domains based on load. I have a batch process that runs each evening and would like the domain to borrow resources when the other domains are idle and return them when needed.

Cindi McGuire (A): No. All PSH responses (DR, included) are based upon the diagnosis of a system fault.

Q: Does PSH work or react any differently on real (Intel) processors, compared to AMD processors?

Smita Thakur (A): No it is just the same.

Q: Is it correct that every event that used to result in a message to syslog will now (also) result in a message to fmd?

Michael Shapiro (A): No. The transition from old-style software that simply spews error messages to syslog to self-healing telemetry is a gradual one. So we've focused on some of the key areas for RAS in the first release (e.g. CPU, Memory, I/O) and we'll be working on others in priority order. We also want our partners and ISVs to plug in as well.

Q: can you briefly discuss the idea surrounding telemetry probes and how you reached that concept

Cindi McGuire (A): In traditional Unix systems, error information is haphazardly output to syslog for humans to diagnose the underlying fault. The error messages are confusing and difficult to write software to interpret and respond to. With PSH, we have moved away from unstable error messages to a new protocol for describing error information using stable, self-describing data. The error data forms the basis of the new error stream (telemetry) that can be dispatched to software for programmatic diagnosis.

Q: Will PSH detect that I am underpaid and get me a raise? Will it detect that I am tired and buy me a coffee? If yes, then I am sold.

Smita Thakur (A): Absolutely as long as you can configure it appropriately.

Q: Will I be able to add new errors - and responses - to PSH?

Michael Shapiro (A): In the first release of PSH, the way we will permit you to plug in to PSH is using the Service Management Facility (SMF) for user applications. Later, we will begin exposing APIs for other types of plug-ins to device driver developers and for other uses. We will also deliver modules that permit administrators to configure custom responses to diagnosis results such as e-mail messages, SNMP traps, and so forth.

Q: Is there a possibility for PSH to detect failures in services on other machines? Like lets say database failure on other server.

Angel Camacho (A): As far as the system is running Solaris 10, yes, just remember this is a per system feature.

Q: Please compare aggressive page retirement (Solaris 8/9) and PSH (Solaris 10) with regards to memory error management.

Michael Shapiro (A): Solaris 10 includes all of the underlying technology we use in the VM system such as page retirement found in previous releases. So with PSH, you get aggressive page retirement but we have more sophisticated diagnosis algorithms determining when to apply it, and we have all the other stuff such as an improved administrative model.

Q: Is Admin Guide available for PSH?

Smita Thakur (A): The manpages are already posted at: docs.sun.com The admin guide will be available shortly. For more info please check the bigadmin website.

Q: how to do get the Solaris 10 software?

Brian Ellefritz (A): You can download Solaris 10 builds by joining the Solaris Express program at: <http://www.sun.com/software/solaris/solaris-express/>

Q: are self-healing activities configurable

Angel Camacho (A): yes, they are, for further information see the Technical introduction to Predictive Self Healing

Q: The tech intro talks about Self Healing Ecosystem. The software components talked about there seem to be specific to diagnosis engines for different hardware modules (I/O devices, etc.) How does Sun envision this Technology being used with software products (web, App, directory, etc) to monitor the resources these services use on a system. Will the Pred. Self-Healing extend to the memory allocated to a software process and recognize faults within that memory space? Will it be able to trigger a service restart or is the vision to use this only on kernel level faults?

Michael Shapiro (A): For software services, the next release of Solaris Express will include the second part of PSH, which is the Service Management facility. This performs monitoring of software applications and uses PSH to know when to automatically restart them. You can look at Stephen Hahn's blog (blogs.sun.com/sch) for a preview. For your memory question, we will recognize faults within the memory space of any user process. If the fault causes correctable errors (CE's) we'll diagnose the issue and potentially remap the physical page under the process. If the process hits an uncorrectable error (UE) in memory, it will be killed and then SMF will restart the containing service. So PSH addresses faults that affect the kernel and those that affect user processes.

melefant12 (Q): How long have you been planning this release?

Angel Camacho (A): Some of the technology inside Solaris 10, has been under development for up to 3 years, while the design and architecture, for a little bit longer, it's been a long, but gratifying road.

Q: How will developers interact with predictive self-healing system? Especially Real-time Java (RTSJ) developers that may have to give up allocated hardware resources?

Cindi McGuire (A): Application developers will be allowed to participate as an SMF service that may be restarted in the face of a fault on one or more resources it is using. Resource abstractions via Java will continue to work but may not be fully connected to SMF resource service contracts.

Q: Will PSH have a defined message set logging to /var/adm/messages and will there be a string that can identify all PSH messages?

Michael Shapiro (A): Yes, it does. If you download the white paper from <http://sun.com/msg/> you will see an example screen shot of the diagnosis message. The message always starts with "SUNW-MSG-ID" in the upper-left-hand corner. You can also configure the syslog-msgs PSH module to send the PSH diagnosis results to one of syslogd(1M)'s LOCAL0-7 facilities, and then set up syslog.conf to segregate that facility into a separate file (i.e. other than /var/adm/messages). Finally, we will be delivering in the Solaris 10 timeframe a module that will permit administrators to forward such messages to custom scripts (e.g. to e-mail them).

Q: Where can I get the PSH Technical Guide? Was on vacation and just registered this morn.

Brian Ellefritz (A): We will send the Technical Guide again in the follow up email from this event. (Hope you had a great vacation!)

Q: Is Solaris 10 will be also ported on Fujitsu hardware?

Smita Thakur (A): Yes.

Q: If it auto-recovers, does it also auto-patch and upgrade?

Cindi McGuire (A): Not currently, but we are researching diagnosis of software defects and automated responses and self-healing of broken software packages.

Q: Will Solaris 10 run on a spark 10?

Angel Camacho (A): It is not a tested configuration and won't be supported.

Q: What errors that today would panic a system will be handled by PSH? What won't?

Michael Shapiro (A): In both past releases and in Solaris 10, we've been actively working to make the system recover and continue from as many types of errors as possible. One place we've made great progress is in I/O: Solaris 10 will not panic from any PCI bus transaction where the hardware maintains system coherence; this was not true in previous releases. We'll also be bringing major improvements to the resilience of Solaris x86 on AMD processors during Solaris 10. There are always cases where the kernel must panic to preserve the integrity of your user data because the h/w error is so severe that it cannot capture enough state or maintain coherence so that the OS can recover. All CPUs and I/O h/w have cases like this. Our goal is to make Solaris be able to survive all the others.

matt (Q): what is different between Solaris 9 and 10?

Angel Camacho (A): There are many new features inside Solaris 10 that you don't find in Solaris 9, like Trace, Predictive Self Healing and N1 Grid Containers, you can read more of it at <http://www.sun.com/solaris/10> but remember, there is binary compatibility between Solaris 9 and 10

Q: How does PSH decide on the corrective action to take? Is there a fault-action list published? Can the admin control the actions taken, especially something that affects performance of the system?

Cindi McGuire (A): There is not (currently) a way to get a full 'action' list. The admin may configure PSH agent activity via configuration options. The current options are course-grain: on or off. Future enhancements will give the admin more fine-grained control over agent actions.

Q: Can I configure how PSH reacts to errors?

Smita Thakur (A): Yes you can do that.

Q: what will be the FCS date for Solaris 10?

Brian Ellefritz (A): Solaris 10 will ship at the end of calendar 2004

Q: Will this be integrated into cluster technology? Example: Two of four CPU's fail and the system is running in a degraded state but it is still up. The cluster detects this and failover based on a policy.

Smita Thakur (A): This is being reviewed.

Q: With so many new features in S10, will there be more focused training courses (i.e. for PSH, Zones, Trace)?

Angel Camacho (A): Yes, there is going to be Solaris 10 focused training material, where DTrace, N1 Grid Containers and all the new features will be covered, please check our course catalog at <http://training.sun.com/index.html> later on for our newest information.

Q: Can PSH help the Red Sox predict their problems and correct them, to win the World Series?

Michael Shapiro (A): Yes it can. Since our fmd(1M) architecture permits any diagnosis algorithm to be plugged in, we should be able to incorporate both the brain of Bill James as well as relevant sabermetrics to make the right thing happen. On the other hand, did we really need a new software architecture to know to replace Pedro when he was gassed in game 7?

Q: are the messages of Software Express Program showed before occur the problem, starting from symptoms

Smita Thakur (A): Yes.

Q: Will PSH have a connection to Netconnect (the successor to Sun Management Center).

A: This is part of the roadmap for PSH but will not be available in the first release of Solaris 10

Q: What will the PSH do out of the box and what sort of configuration is needed to fully use it?

Cindi McGuire (A): PSH provides error handling, fine-grained fault diagnosis, predictive analysis of failing components and self-healing responses for faults detected in a system. Applications and services will continue to run while the system transparently self-heals and resources are failed over or restarted. PSH is available on all S10 systems with special instrumentation of the CPU/mem and PCI subsystems available for US-III based systems.

Q: What are the data sources that feed Error Reports? Are any of these new relative to Solaris 9?

Michael Shapiro (A): All of PSH is new relative to Solaris 9. Prior to Solaris 10 and PSH, all error information from the kernel was transmitted by means of cmn_err(9F), which just sends a text string for humans to syslog. In Solaris 10, PSH uses a structured event transport to send telemetry events for automated diagnosis from the kernel to fmd(1M). And we have enhanced the ability of many subsystems, such as our bus nexus drivers, to be able to capture error reports for automated diagnosis.

mark (Q): How do I get the hold of Software Express - Is it available through download on both Sparc and Intel?

Angel Camacho (A): yes, it is available for both platforms through <http://www.sun.com/software/solaris/solaris-express/>

Q: Can be PSH installed on top of Solaris 9?

Angel Camacho (A): That is a new feature of Solaris 10, and will not be part of Solaris 9

tim (Q): write to future delivery of predictive-self healing are plans to expand the functionality outside of systems?

Smita Thakur (A): Predictive Self-Healing is a new fault management technology being implemented across Sun's product line. It is a Sun wide effort.

Q: What happens when a fault occurs on a CPU where the kernel is running or where permanent memory resides?

Michael Shapiro (A): If a CPU fault occurs while a thread is executing in userland, then the user process will be terminated and the Service Manager will restart the containing service. If a CPU fault occurs while a thread is executing in the kernel, then it depends whether or not the thread is in a protected code region or not: one example of such a region is copying in data (or out) as part of system call processing; there are others. In these cases, we can similarly contain the problem and continue. If we are not in a protected region, the kernel will panic and reset, and the problem will be diagnosed on the way back up (or by a service processor). Memory faults have similar different cases: the Solaris VM system has to look at the state of the page, whether a read or write is being attempted, whether the page is clean or dirty, and so on to determine the degree of isolation and recovery. We're actively working on improving every area of Solaris to handle errors by isolating and retiring the bad resource to the degree possible given the hardware platform, whether it be SPARC or x86/AMD.

Q: Will this be back ported to Solaris 9?

Smita Thakur (A): No plans of backporting to Solaris 9.

Q: What does Dynamic File Systems do for me

Smita Thakur (A): Please attend the Expert Exchange Session on 9/15/04 for the session on Dynamic File System

Q: Will the dynamic file system have the capability of increasing _and shrinking_ file systems, similar to what Veritas VM and vxfs can do today?

Angel Camacho (A): Please come to our next Sun Expert Exchange, on September 15th, that will be our next topic.

Q: What about UltraSparc II chips. I would like to leverage this technology in my existing midrange servers where I experience many faults. What about SBUS?

Cindi McGuire (A): We have implemented PSH capabilities for the US-II PCI subsystem. Future instrumentation of the US-II CPU/memory subsystem is planned but not available in S10. There are no plans for sbus instrumentation.

Q: What will it be able to heal? Anything like (or better than) a system restore function?

Michael Shapiro (A): Assuming by "system restore" you're talking about disaster recovery features, then yes, the new Service Management Facility (SMF) coming as part of PSH will include features to help. SMF permits administrators to take snapshots of the service configuration, give them human-readable names, and restore them on demand. SMF will appear in the next release of Solaris Express.

mark (Q): There was a reference to "Test Driving Solaris 10", I tried to download the software but the only software that was out there was Solaris 7/04. Can we download it for both Intel and Sparc ?

Smita Thakur (A): Yes. The new release was released yesterday through Software Express.

Q: can you provide an example of the predictive self healing technology in action?

A: Please refer to the PSH Technical Guide that was sent out with the invitation for this event. There is a great example in there with screen shots.

Q: Can customized agents be created by the customer, in particular with respect to monitoring zones, or is fmd restricted to the global zone?

Cindi McGuire (A): In the current implementation, fmd runs in the global zone. Future connections with zones will allow PSH diagnosis and response (agent actions) in local zone configurations.

Q: does this include kernel parameter tuning? ex; number of handles for users and such (being auto-tuned)?

Michael Shapiro (A): One of our major goals in Solaris has always been to make the system self-tuning. Every release we've taken away more and more of these: for example, in Solaris 9 we made the number of ptys scale automatically, and in Solaris 10 we've taken away the need to tune IPC and shm tunables, making them dynamically-scaling resource controls. In PSH, we've designed our new features from the ground up not to require custom tunables.

kagreen (Q): Can you explain N1 Grid Containers?

Angel Camacho (A): Please refer to the transcripts from previous Expert Exchanges (last month), or <http://www.sun.com/software/solaris/10/>

Q: are there going to be enhancements to utilizing systems calls to better capture the full features provided with self healing? this also ties in to use of d-trace functions as well.

Cindi McGuire (A): Future enhancements to PSH will allow applications and services to receive fault information and respond based upon the resources it uses. There are no direct ties with DTrace.

slava (Q): How about General Overview of Predictive Self-Healing?

Smita Thakur (A): Sun has developed a new architecture for building and deploying systems and services capable of Predictive Self-Healing. Self-healing technology enables Sun systems and services to maximize availability in the face of software and hardware faults and facilitates a simpler and more effective end-to-end experience for system administrators, reducing cost of ownership.

Q: Can PSH be used to monitor and correct problems in user apps? If so, can I use the framework to run custom scripts? Is there is an API for the framework?

Michael Shapiro (A): PSH includes our new Service Management Facility (SMF). You can see some previews of this in Stephen Hahn's blog (<http://blogs.sun.com/sch>). SMF monitors all running services on the system, and can automatically restart them. APIs are provided for writing custom monitoring scripts, including the ability to wait for a service to change state. SMF will appear in the next Solaris Express download.

Cindy (Q): Will there be a live upgrade path from Solaris 8 directly to 10?

Smita Thakur (A): Yes.

Q: How is this different from Solaris 9?

Smita Thakur (A): Predictive Self Healing is not part of Solaris 9. It is a new technology in Solaris 10.

Q: What are some of the new features in Solaris 10?

Angel Camacho (A): Some of the new features you can find in Solaris 10 are DTrace, Predictive Self-Healing, N1 Grid Containers and Dynamic File Systems. For more information please check <http://www.sun.com/solaris> and click on the Solaris 10 link.

Q: What, if any overhead does fmd introduce, or will this be similar in nature to the use of DTrace?

Michael Shapiro (A): DTrace is a facility for dynamic instrumentation, so it has no overhead when not in use, and an overhead proportional to the question you ask when you use it. FMD is a continuously running daemon, but it only does something when an error is detected on the system for which self-healing telemetry is present. So its cost is proportional to whether or not the system is experiencing a fault.

Q: how is sun positioning the education and training associated with predictive self-healing?

Smita Thakur (A): Predictive Self-Healing is covered by the training courses on Solaris 10. Stay tuned for more.
