# Sun Expert Exchange

Technical Knowledge Base for Sun Inner Circle Members

## Fast Track to Solaris 10 Adoption: Predictive Self-Healing

### A Sun Expert Exchange Discussion

Predictive Self-Healing (PSH) is an innovative capability that automatically diagnoses, isolates, and recovers from many hardware and application faults. PSH allows critical applications and services to continue uninterrupted in the event of software and hardware component failures, and even software mis-configuration problems.

This summary includes highlights of the hour-long Q&A,* organized into the following sections:



In addition to questions and answers, you'll also find references and links to additional resources provided by Sun.

*Note: The information contained in this transcript, taken directly from a live Sun Expert Exchange event, has been edited for clarity and adherence to trademark guidelines.



# Fast Track to Solaris 10 Adoption: Predictive Self-Healing

## General Information



---

**Q: How about General Overview of Predictive Self-Healing?**

**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, thus reducing cost of ownership.

**Q: How long have you been planning this release?**

**A:** Some of the technology inside the Solaris 10 OS, has been under development for up to three years, a little bit longer for the design and architecture. It's been a long, yet gratifying road.

**Q: How can I get the software?**

**A:** 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 the Solaris 10 OS, we recommend the free version. You can get it here. Things might be a bit more clear. You can get CD images for either SPARC or x86 platforms, Commercial or Non-Commercial use.

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

A: Some of the new features you can find in Solaris 10 OS 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: Is work being done to send alerts from PSH to Sun Field Support to handle CPU/memory replacements?**

**A:** Yes, we are working on making sure that there is a good connection with the remote facilities.

**Q: What is different between the Solaris 9 and 10 operating systems?**

**A:** There are many new features inside the Solaris 10 OS that you don't find in the Solaris 9 OS, such as DTrace, Predictive Self-Healing, and N1 Grid Containers. You can read more about it at http://www.sun.com/solaris/10. Remember, there is binary compatibility between Solaris 9 and 10 operating systems.

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

**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 drivers), then move up the stack to higher level abstractions (e.g. file systems, network, Java).

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

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

**Q: What is SMF in full name?**

**A:** Service Management Facility.

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

**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?**

**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: How does PSH relate to SNMP?**

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

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

**Q: Can you briefly discuss the idea surrounding telemetry probes and how you reached that concept?**

**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: 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 PSH 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?**

**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 (CEs), we'll diagnose the issue and potentially re-map the physical page under the process. If the process hits an uncorrectable error (UE) in memory, it will be killed, then SMF will restart the containing service. So PSH addresses faults that affect the kernel and those that affect user processes.

**Q: How will developers interact with the PSH system, especially Real-time Sun Java (RTSJ) developers who may have to give up allocated hardware resources?**

**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: What will be the FCS date for the Solaris 10 OS?**

**A:** The Solaris 10 OS will ship at the end of 2004.

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

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

**Q: How do I get the hold of Software Express? Is it available through download on both SPARC and Intel platforms?**

**A:** Yes, it is available for both platforms through http://www.sun.com/software/solaris/solaris-express/

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

**A:** That is a new feature of the Solaris 10 OS.

**Q: Write to future delivery of PSH — are plans to expand the functionality outside of systems?**

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

**Q: How is this different from the Solaris 9 OS?**

**A:** Predictive Self-Healing is not part of the Solaris 9 OS. It is a new technology in the Solaris 10 OS.



# Fast Track to Solaris 10 Adoption: Predictive Self-Healing

## Documentation & Training

1. Can you provide an example of the PSH technology in action?
2. When will be the first training or update for the Solaris 10 OS?
3. How is Sun positioning the education and training associated with PSH?
4. Are there any guides on how developers can take advantage of PSH?
5. Where can I find more information on SMF? A quick search of the main Sun site didn't reveal anything interesting.
6. Does the white paper touch on configurable behaviors and instantiating childs procs of them?
7. Will you be publishing more articles on tip and tricks for PSH?
8. Is there a way to view PSH if one is visiting Palo Alto?
9. Is there a good out-of-the-box overview for Solaris 10 OS admin, presenting both new and old tools?
10. Is an admin guide available for PSH?
11. Where can I get the PSH Technical Guide?
12. With so many new features in the Solaris 10 OS, will there be more focused training courses (e.g., for PSH, Zones, DTrace)?

---

**Q: Can you provide an example of the PSH technology in action?**

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

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

**A:** Training will begin rolling out around our ship date at the end of 2004.

**Q: How is Sun positioning the education and training associated with PSH?**

**A:** Predictive Self-Healing is covered by the training courses on the Solaris 10 OS. Stay tuned for more details.

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

**A:** The main 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 plenty of SMF documentation made available.

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

**A:** Please check out http://blogs.sun.com/sch

**Q: Does the white paper touch on configurable behaviors and instantiating childs procs of them?**

**A:** It does touch on configurability, but doesn't go as far as "how to" levels of detail.

**Q: Will you be publishing more articles on tip and tricks for PSH?**

**A:** We're working on a more consolidated location for emerging resources for PSH; for now we'll be updating the "BigAdmin" site.

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

**A:** Please visit Sun's Executive Briefing Center in Menlo Park, where you can get a demo on Predictive Self-Healing.

**Q: Is there a good out-of-the-box overview for Solaris 10 OS admin, presenting both new and old tools?**

**A:** There are a number of Solaris 10 OS admin resources being posted to the "BigAdmin" site.

**Q: Is an admin guide available for PSH?**

**A:** The main pages are already posted at: docs.sun.com. The admin guide will be available shortly. For more info please check the BigAdmin Web site.

**Q: Where can I get the PSH Technical Guide?**

**A:** We will send the Technical Guide again in the follow up email from this event.

**Q: With so many new features in the Solaris 10 OS, will there be more focused training courses (e.g., for PSH, Zones, DTrace)?**

**A:** Yes, there is going to be Solaris 10 OS-focused training material, where DTrace, N1 Grid Containers, and all the other new features will be covered. Please check our course catalog later on for the latest information.



# Fast Track to Solaris 10 Adoption: Predictive Self-Healing

## Performance Issues

1. What mechanism does PSH use to detect an application crash? Will it be able to automatically upload crash dump to Sun and initiate a service call to tech services?
2. 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?
3. Given that PSH supports systems as individuals, will it handle raid levels with customizations?
4. How will PSH handle fan failures or CPU overheats?
5. Does self-describing data imply that there are XML schemas?
6. Are PSH capabilities applied system wide, or is there a finer granularity (ex. Solaris 10 OS Zones/ grid containers)
7. 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?
8. How does PSH interact with Containers to help prevent downtime? Can a hardware failure be transparent to a container?
9. My server had a panic on CPU2 write back check error. Should I say that Solaris 10 OS PSH will save the server from this type of error?
10. Can PSH send traps to a monitoring solution when there is a fault?
11. Will the SUNW-MSG-ID be server-centric with updates?
12. Will this provide the ability to dynamically reconfigure domains based on load? I have a batch process that runs each evening; I would like the domain to borrow resources when the other domains are idle and return them when needed.
13. Does PSH work or react any differently on real (Intel) processors, compared to AMD processors?
14. Will every event that used to result in a message to syslog will now (also) result in a message to fmd?
15. Will I be able to add new errors and responses to PSH?
16. Is there a possibility for PSH to detect failures in services on other machines, e.g., a database failure on other server?
17. Will PSH have a defined message set logging to /var/adm/messages, and will there be a string that can identify all PSH messages?
18. If it auto-recovers, does it also auto-patch and upgrade?
19. What errors that would panic a system today will be handled by PSH? What won't?
20. How does PSH decide on 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?
21. Are the messages of Software Express Program showed before the problem, occurs starting from symptoms?
22. What are the data sources that feed error Reports? Are any of these new relative to the Solaris 9 OS?
23. What happens when a fault occurs on a CPU where the kernel is running or where permanent memory
24. Does this include kernel parameter tuning, e.g., the number of handles for users and such (being auto-tuned)?
25. 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?
26. What, if any overhead does FMD introduce, or will this be similar in nature to the use of DTrace?

**Q: What mechanism does PSH use to detect an application crash? Will it be able to automatically upload crash dump to Sun and initiate a service call to tech services?**

**A:** This involves a daemon (fmd) that collects information on events inside the system that updates an event list that is "published" for use by 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: 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?**

**A:** Yes. We are planning to harvest SMART data and create an error telemetry that allows us to predict disk failures.

**Q: Given that PSH supports systems as individuals, will it handle raid levels with customizations?**

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

**A:** CPU Overheats: yes; Fan Failures: yes, depending on platform.

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

**A:** XML can be used to marshal our self-describing error and fault protocol event data.

**Q: Are PSH capabilities applied system wide, or is there a finer granularity (ex. Solaris 10 OS Zones/grid containers)**

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

**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?**

**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 the Solaris 10 OS 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: How does PSH interact with Containers to help prevent downtime? Can a hardware failure be transparent to a container?**

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

**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 OS patches and in the Solaris 10 OS, to harden the Solaris OS against all such errors to the degree permitted by the hardware.

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

**A:** Our plan is to leverage custom agents to do it; please check our site often for new information: http://www.sun.com/msg

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

**A:** As appropriate, the message ID will be platform specific. For example, a fault message that is specific to a Sun Fire 6900 system 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 provide the ability to dynamically reconfigure domains based on load? I have a batch process that runs each evening; I would like the domain to borrow resources when the other domains are idle and return them when needed.**

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

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

**A:** No it is just the same.

**Q: Will every event that used to result in a message to syslog will now (also) result in a message to fmd?**

**A:** No. The transition from old-style software that simply spews error messages to syslog to self-healing telemetry is a gradual one. 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.

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

**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 on.

**Q: Is there a possibility for PSH to detect failures in services on other machines, e.g., a database failure on other server?**

**A:** As far as the system is running the Solaris 10 OS, yes, just remember this is a per-system feature.

**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?**

**A:** Yes, it does. If you download the white paper from http://www.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 OS timeframe a module that will permit administrators to forward such messages to custom scripts (e.g., to e-mail them).

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

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

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

**A:** In both past releases and in the Solaris 10 OS, we've been actively working to make the system recover from as many types of errors as possible. One place we've made great progress is in I/O: the Solaris 10 OS 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 OS on AMD processors during the Solaris 10 OS. 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 the Solaris 10 OS able to survive all the others.

**Q: How does PSH decide on 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?**

**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: Are the messages of Software Express Program showed before the problem, occurs starting from symptoms?**

**A:** Yes.

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

**A:** All of PSH is new relative to the Solaris 9 OS. Prior to the Solaris 10 OS 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 the Solaris 10 OS, PSH uses a structured event transport to send telemetry events for automated diagnosis from the kernel to fmd(1M). Also, we have enhanced the ability of many subsystems, such as our bus nexus drivers, to be able to capture error reports for automated diagnosis.

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

**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 on whether the thread is in a protected code region. One example of such a region is copying in data (or out) as part of system call processing, and there are other examples. 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 the Solaris OS 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 technology.

**Q: Does this include kernel parameter tuning, e.g., the number of handles for users and such (being auto-tuned)?**

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

**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?**

**A:** PSH includes our new Service Management Facility (SMF). You can see some previews of this in Stephen Hahn's blog. 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.

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

**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 the system is experiencing a fault.



# Fast Track to Solaris 10 Adoption: Predictive Self-Healing

## Installation & Configuration



---

**Q: What tools will be available to configure PSH?**

**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: Are self-healing activities configurable?**

**A:** Yes, they are. For further information see the Technical Introduction to Predictive Self-Healing.

**Q: Can I configure how PSH reacts to errors?**

**A:** Yes.

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

**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 Solaris10 OS systems with special instrumentation of the CPU/mem and PCI subsystems available for US-III based systems.

**Q: Will there be a live upgrade path from the Solaris 8 OS directly to the Solaris 10 OS?**

**A:** Yes.



# Fast Track to Solaris 10 Adoption: Predictive Self-Healing

## Compatibility Issues

1. Will the agent dev tool for integration be another d-type of language, or scripting capable?
2. How will or can PSH interact with Sun or Veritas Cluster services?
3. Are trap messages the only SNMP “article” produced? Is there a history/log that can be polled at the MIB level?
4. Will PSH data be available through kstat?
5. How can third-party applications use PSH features? Is there any way of checking third-party software logs for errors?
6. Why isn’t Sun back porting this to the Solaris 9 and 8 operating systems? Not everyone will be moving to the Solaris 10 OS on [Sun’s] timeline, but could still benefit from this technology in their existing Data Centers.
7. How does PSH respond to Oracle in a clustered environment? Example: Oracle runs out of swap space.
8. 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?
9. Will API calls be available to tie into NMS tools such as OpenView?
10. How is PSH integrating with monitoring tools/agents such as Sun Cluster, HP OpenView, and Tivoli? Is there an SNMP agent to send alerts to a control station?
11. Please compare aggressive page retirement (Solaris 8 and 9 operating systems) and PSH (Solaris 10 OS) with regard to memory error management.
12. Will the Solaris 10 OS also be ported on Fujitsu hardware?
13. Will the Solaris 10 OS run on a spark 10?
14. 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?

---

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

**A:** We are working on a PSH module that 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?**

**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: Are trap messages the only SNMP “article” produced? Is there a history/log that can be polled at the MIB level?**

**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: Will PSH data be available through kstat?**

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

**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 the Solaris 9 and 8 operating systems? Not everyone will be moving to the Solaris 10 OS on [Sun's] timeline, but could still benefit from this technology in their existing Data Centers.**

**A:** The intention is to have Sun's customers take advantage not only of this key technology in the Solaris 10 OS, but also of other groundbreaking technologies in that release. These technologies demanded a radical change in the kernel; hence, it would take a big effort to backport them the Solaris 8 and 9 operating systems.

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

**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: 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?**

**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?**

**A:** Our plan is to leverage customized agents to integrate those functions, please visit in the near future for up-to-date information.

**Q: How is PSH integrating with monitoring tools/agents such as Sun Cluster, HP OpenView, and Tivoli? Is there an SNMP agent to send alerts to a control station?**

**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 OS timeframe to permit SNMP traps to be sent triggered when an automated diagnosis occurs.

**Q: Please compare aggressive page retirement (Solaris 8 and 9 operating systems) and PSH (Solaris 10 OS) with regard to memory error management.**

**A:** The Solaris 10 OS 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: Will the Solaris 10 OS also be ported on Fujitsu hardware?**

**A:** Yes.

**Q: Will the Solaris 10 OS run on a spark 10?**

**A:** It is not a tested configuration and won't be supported.

**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?**

**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 the Solaris 10 OS. There are no plans for SBus instrumentation.



# Fast Track to Solaris 10 Adoption: Predictive Self-Healing

## Functionality & Usability Issues

1. Is access to the SUNW-MSG-ID limited to the Internet, or will it be hosted locally with the FCM?
2. Is the Sun acquired cluster technology going to be leveraged for PSH to applications? What is SMF?
3. Are PSH features intended to be used by app developers?
4. What will it be able to heal? Anything like (or better than) a system restore function?
5. Can customized agents be created by the customer, in particular with respect to monitoring zones, or is FMD restricted to the global zone?
5. 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 DTrace functions.

---

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

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

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

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

**A:** Absolutely. Whether you're using it to take corrective actions on your development, or taking advantage of the full set of APIs, there's something you can use if you're a developer.

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

**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.

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

**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: 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 DTrace functions.**

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