Whitepaper

SGI™ Embedded Support Partner

SGI Embedded Support Partner

Managing today's enterprise systems is not merely more difficult than managing the corporate computing environments of the past, but is also fundamentally different. Managing key software applications, systems, and the corporate network now requires an open, highly scalable, cross-platform solution. To be competitive and to differentiate your service, you must be able to ensure continuous system availability in a highly distributed changing environment. Monitoring purely for computing resources does not protect levels of service and availability. SGI Embedded Support Partner (ESP) addresses this challenge. This paper provides an overview of the features and benefits of ESP with a primary focus on its architecture.

ESP is a set of facilities embedded in the IRIX® operating system that provides an integrated support environment. ESP enables SGI customers to achieve high levels of availability through reliable, proactive, automated support. For distributed systems, ESP's scalable architecture provides an efficient, consistent, centralized management and support capability. Personnel can be notified of developing system conditions automatically so that issues can be resolved before they develop into catastrophic failures.

ESP can be configured in two modes:

- Single system mode
- System group manager mode

In both modes, ESP supports the following facilities:

- ESP monitors
- ESP notifiers
- ESP administration
- ESP reports

ESP Monitors

An event is any abnormal system condition—for example, performance degradation in the system, a memory error, a panic, or a hang. ESP's monitors allow the detection of these conditions.

The default monitors in ESP include:

- Kernel monitoring
- Performance monitoring
- Availability monitoring
- Configuration monitoring
- Application monitoring

Kernel Monitoring

Kernel-level events are monitored by a special event-monitoring daemon that watches syslogd directly and filters events as they occur, based on ESP's event profile setup. SGI's unique sequence number scheme allows kernel-level events to be monitored with efficiency and scales as the number of event types increases. For operating systems that do not follow the unique sequence scheme, a template-based monitor takes precedence.

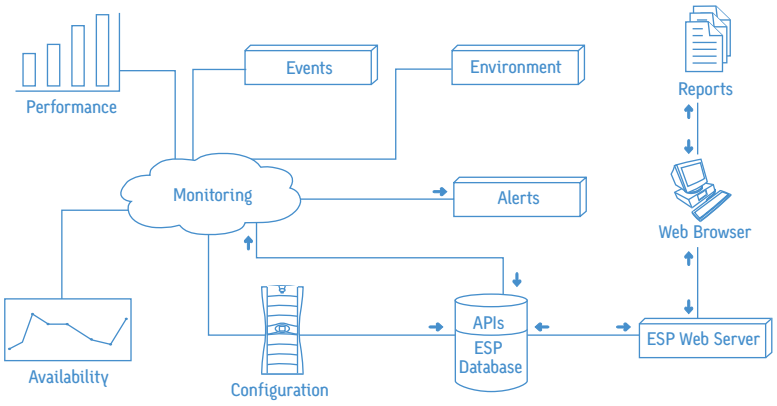

Figure 1: Embedded Support Partner architecture

Update Events. Class "Kernel"

strlab04.csd.sgi.com 10 records per page

No	Event Description	Status	Occurrence	Assigned Actions
1	Allocated more memory than cleared	Enabled	1	None
2	Bad device	Enabled	1	None
3	Bad free size for bitmap(1)	Enabled	1	None
4	Bad free size for bitmap(2)	Enabled	1	None
5	Bad prom swap	Enabled	1	None
6	Biophysio Failed userdma	Enabled	1	None
7	Bitmap overflow	Enabled	1	None

Figure 2: Event Monitoring

Performance Monitoring

Performance monitoring is accomplished using the unique SGI Performance Co-Pilot™ application. This client-server architecture facility has the ability to monitor hundreds of kernel metrics. After extensive research and experimentation with systems in database environments, Web servers, computational environments, file servers, etc., more than 30 performance monitoring rules have been defined. The performance monitor inference engine logs events when thresholds defined in the rules are triggered. Performance Co-Pilot utilizes SGI's unique sequence numbering scheme to log the events into the event management scheme.

Availability Monitoring

Availability monitoring is accomplished by the SGI availability application, which runs after every reboot. The monitor looks for several types of availability events, including administrative reboots, system hangs, system panics, and interrupts. The availability monitor uses unique SGI crash analyzers to automatically generate crash analysis files that include stack traces. The capture of various types of interrupts gives ESP the ability to generate accurate availability metrics.

Configuration Monitoring

System hardware and software monitoring is accomplished using the unique SGI configuration monitor. The monitor walks through kernel data structure to generate a unique component-to-component relationship on every reboot. Using ESP's local database, the configuration monitor has the capability to distinguish configuration changes and record change events and change records. This capability enables reconstruction of the system hardware and software configuration for any given instance of time as well as asset management. Configuration information includes serial numbers, part numbers, revision numbers, board locations, installed software, and patches.

Application Monitoring

ESP can monitor user applications by virtue of the simple event-monitoring application programmer's interfaces (APIs) distributed with ESP. Using the ESP infrastructure, a user can define application events and integrate the API into the applications. A single call to the API within exception codes of the application can trigger a sequence of corrective actions. Script-oriented applications can also be monitored using ESP's logging facility esplogger.

ESP's monitors support a unique event profile mechanism that allows dynamic loading and monitoring of specific events. A variety of profiles can be set up, loaded, and unloaded based on the type of environment the system is set up to operate. An example would be setting up performance monitoring profiles for CPU-intensive environments. The profiles are simple human-readable text file formats that can be hand-edited and installed. Profiles also make it easy to distribute event setup to several systems for identical monitoring.

ESP Notifiers

ESP provides a variety of notifiers. These are proactive notifications, executed in real time, that can be set up by the system administrator. The proactive notifications can be defined on a per-event basis and can include any or all of the notification schemes described below. The notifiers can be configured to provide detailed information about the event. This includes type of event, error message associated with the event, time of occurrence of the event, host on which the event occurred, etc. The various mechanisms for notification include:

- ASCII notification to the administrator's console
- Graphical pop-up notifiers to any system console (local/remote)
- E-mail notification (plain text or encrypted)
- Pager notification

ESP Administration

As with any application, ESP requires administration for effective use of the system. ESP users define the administration policies within which ESP operates at all times. All components of ESP that are administered can be modified dynamically. ESP administration can be accomplished with a user-friendly browser-based interface or a command line interface (CLI). The policies defined for administering ESP include:

- ESP users and user privileges
- ESP Web server administration
- Setup of system monitors, rules, and responses

ESP Users and User Privileges

ESP maintains a record of a set of users. ESP users are not UNIX® users and need not have UNIX accounts to be able to use ESP. The administrator defines ESP users when ESP is administered for a system. Privileges that govern ESP users include the ability to perform the following activities:

- Set up an ESP environment
- Configure and enable ESP event structures
- View events, actions, and diagnostic reports
- Generate system availability reports
- Generate hardware and software inventory reports, including historical views
- Create and view ESP logs

The above privileges are implemented via page privileges that are managed by the ESP Web server.

ESP Web Server Administration

As part of ESP setup, the administrator must set up users and user privileges. ESP also maintains access lists of systems defined by the IP addresses that allow them to connect to the Web server. The access list can be set to filter at subnet levels. For added security, ReverseDNS lookup is also implemented to validate connections to the ESP Web server. Access paths and directory resolution by the Web server are all defined within a root-owned configuration file that provides additional security levels.

Setup of System Monitors, Rules and Responses

Setting up system monitors entails many functions, including the recording of data, throttle, and thresholds. Recorded data is extremely useful for troubleshooting system problems. This data provides a view of all events as seen by the system plus the ability to identify problematic areas. All data in ESP is managed on a local central relational database. To reduce the impact of data volume and any performance degradation, throttle rules can be to identify each type of event that governs when new records are created. The proactive feature of ESP comes from the action processor built into the event-management scheme. Rules can be applied to events to govern what actions need to be taken. Such rules also include the time and count thresholds at which actions will be taken. The setup of system monitors, rules, and responses is a function of the fine tuning that the system administrator/support professional will perform to optimize ESP for the needs of the operation environment.

ESP Reports

Armed with a comprehensive data set captured by the various monitors in ESP, administrators can generate a variety of reports. All reports can be generated on the command line interface and on ESP's browser-based graphical user interface. Access to the various reports is controlled by the privileges set by the administrator for each user. Reports can be generated for a range of dates where applicable. Reports include:

- System overview report
- Event report summary
- Availability report summary
- Hardware report summary (inventory and change reports)
- Software report summary (inventory and change reports)
- Diagnostic report summary
- Action report summary

Availability Report

strlab04.csd.sgi.com **05/07/2001 to 06/06/2001**

Interrupts	Count	Downtime	MTBI	Availability
Unscheduled	none	0 min	N/A	100.00%
Scheduled	1	5 min	720 hrs	99.99%
administrative: reboot	1	5 min	720 hrs	
Scheduled and Unscheduled	1	5 min	720 hrs	99.99%
Average uptime	359 hrs 57 min			
Least uptime	485 hrs 25 min (current epoch)			
Most uptime	841 hrs 21 min			
Average downtime	5 min			
Least downtime	5 min			

Figure 3: Availability Report

Hardware Inventory Report

strlab04.csd.sgi.com **06/06/2001 16:43:05**

50 records per page

No	Part Name	Location	Part Number	Serial Number	Revision	Installation Date
1	1	N/A	N/A	K0006011	N/A	05/17/2001
2	8P12_MPLN	N/A	030-0762-006	DAS008	E	05/17/2001
3	IP27	n1	030-0733-003	DDR845	J	05/17/2001
4	IP27	n2	030-0733-003	DAY395	L	05/17/2001
5	IP27	n3	030-0733-003	DDR661	J	05/17/2001
6	IP27	n4	030-1266-001	DPM396	H	05/17/2001
7	ROUTER-IR1	r1	030-0841-002	DPC631	G	05/17/2001
8	ROUTER-IR1	r2	030-0841-002	DPC691	G	05/17/2001
9	MENET	io4	030-0873-003	GPS598	G	05/17/2001
10	DIVO	io11	030-1046-002	DEF451	H	05/17/2001
11	MSCSI	io8	030-1243-001	ESJ390	C	05/17/2001
12	BASEIO	io1	030-0734-002	FSV192	J	05/17/2001
13	PCI_XIO	io2	030-1062-002	GPP343	D	05/17/2001

Figure 4: Inventory Report

System Group Manager

Providing continuous, predictable, and reliable support service in a dynamic and highly distributed environment can be challenging and expensive. ESP's scalable architecture provides an efficient, consistent, and centralized way to manage and control distributed systems for support and service. System Group Manager (SGM) provides enhanced functionality and is a licensed feature. SGM and the centralized browser-based management console provide centralized monitoring and management capabilities, including:

- Centralized support administration
- Centralized event processing
- Centralized automated response and notification
- Centralized site reporting

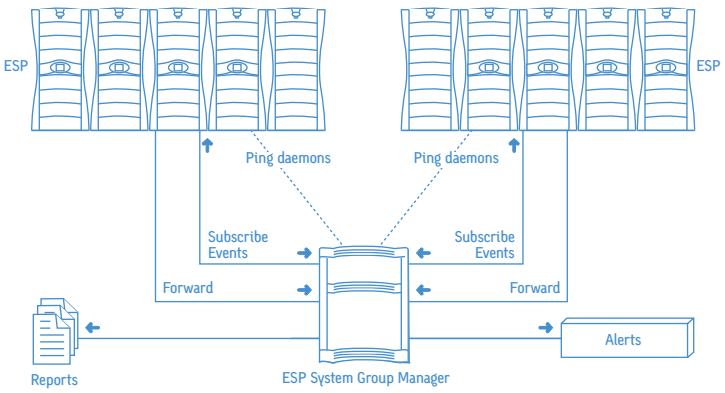

Figure 5: Embedded Support Partner System Group Manager

SGM adopts a simple policy that enables it to selectively subscribe to events from group members. Events are logged to SGM using Remote Procedure Calls (RPC). Data is digitally signed using HMAC (MD5) to control message integrity. ReverseDNS lookup is incorporated for group member validation. SGM can be set up to receive proactive notification on a per-event basis. An SGM agent also allows monitoring for system daemons on the group members. As a result, SGM can recognize any system in the group that goes down and generate notifications. SGM can generate group-level reports or reports for any single member within the group.

Security within ESP

In order to mitigate any potential security breaches that ESP may cause, SGI asked RSA Security, Inc. to perform an evaluation of ESP. Pursuant to the evaluation, SGI implemented a wide range of recommendations from RSA Security, including:

- Validating user permissions of process for proactive actions and disabling actions by root
- Implementing ReverseDNS lookup for both the Web server and ESP SGM
- HMAC/MD5 digital signature of all data transfers to the ESP SGM
- Disabling of login attempts with time-out periods
- Implementing a CLI for all ESP configuration/reports, disabling the use of the ESP Web server
- Restricting all ESP database transactions locally

While no system can offer absolute security, we are confident that SGI provides a highly secure environment with ESP.

Performance Characteristics of ESP

Applications that are built as facilities are always expected to consume the least amount of system resources, including CPU, memory, and disk utilization. ESP's two daemons, eventmond and espdbd, are event-driven and thereby consume CPU resources only when an incident is seen by ESP. However, when ESP receives an event, the time taken to process the event is less than 2 milliseconds of CPU time. This time includes the complete processing and storing of the event into ESP's database. The memory utilization by eventmond and espdbd is ~200KB and ~500KB, respectively. Most of

the disk utilization comes from the storing of system configuration data and may therefore vary based on the configuration of the system. Disk utilization is less than 30MB for 64-processor systems having anywhere from 75 to 125 boards. An archiving facility of ESP (esparchive) can also be run to compress the database. Compression is between 40% and 60% of the original size of the database.

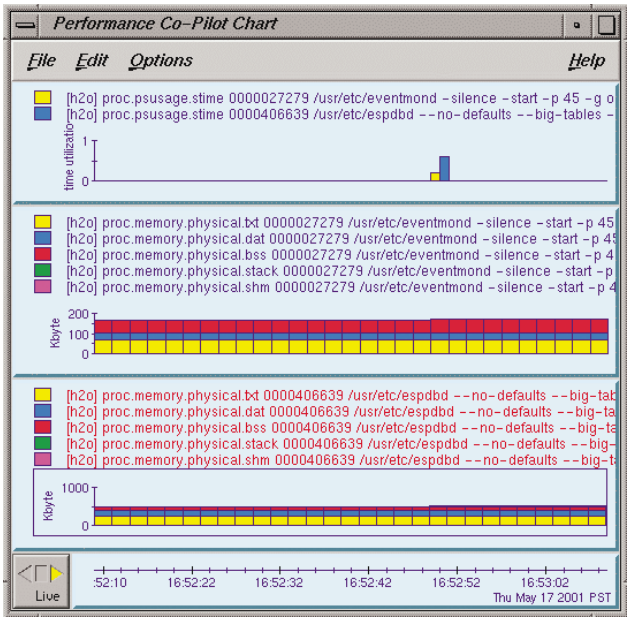

Figure 6: Performance chart on a 1-CPU, 128MB system for ESP daemons.

End-to-End Support—The Works

Once configured, ESP can send incident notifications via plain text or encrypted e-mail to the SGI Customer Support Center. ESP's proactive feature set enables SGI to implement best-in-class services that can increase a customer's system availability to a much greater degree. The notification consists of basic system information and incident information against which a range of business rules are applied. These business rules determine severity and priority for the incident in order to generate a trouble ticket automatically. A parallel search is initiated in the SGI Knowledgebase, which contains thousands of solutions. With uniform, granular search criteria and previously reported incidents from similar system configurations, solutions are extracted with a high degree of accuracy. The trouble tickets are routed through designated queues to specialists who respond quickly to system problems. At the same time, the trouble ticket information and solution IDs are sent to the customer electronically.

The Bottom Line: Automated Problem-Solving and Reduced Downtime

Managing a large enterprise is not easy. Sifting through syslogs to determine what went wrong is cumbersome. In these times, when IT resources are in short supply and systems are becoming much more complex, off-loading mundane monitoring activities to a tool is very appealing. This is exactly what ESP makes possible. Once set up, the systems report anomalies to designated personnel. Most importantly, the capture and notification of events take place in real time, potentially circumventing system downtime. With the integration of ESP to an SGI Customer Support Center and Knowledgebase, ESP can call on SGI specialists to solve system problems before they impact the costomer's business.

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America 1(52) 5267-1387
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 771.0290

3092 (6/01)

J12816