

ParaStation GridMonitor 5.0

Software Product Detailed Description

Product: ParaStation GridMonitor 5.0

Date: May 2011

Document number: PSGM-1.2.2en

This software product description documents the functionality provided by the ParaStation GridMonitor 5.0 as well as the system prerequisites required for installation and operation, licensing scheme and other useful information.

Overview

The *ParaStation GridMonitor* as part of the *ParaStationV5* cluster suite is a versatile system monitor for Linux-based compute cluster. A multiplicity of information from different devices and services from a cluster may be read, evaluated and stored. The *GridMonitor* provides the administrator with various aspects of the available information, from an overall status of all configured clusters to in-depth details of nodes and devices. Data can be grouped with respect to different aspects and visualized using a web browser.

Furthermore, parameters may constantly be monitored and the administrator may be informed, if required.

Technical features

Data gathering process ('collector')

All available data is retrieved and managed by a so called collector. All data is retrieved using dedicated agents for each device or service. These agents support various protocols, like e.g. SNMP or IPMI. To minimize overall system and network load, only data requested by a client application is read from the agents by the collector. If no one is interested, no data is transferred and therefore no compute cycles and network bandwidth is used. The collector is especially designed to handle problems like dead nodes, broken network connections and limited network bandwidth, commonly found in cluster environments.

The collector gathers various data from different information sources available within a cluster, describing:

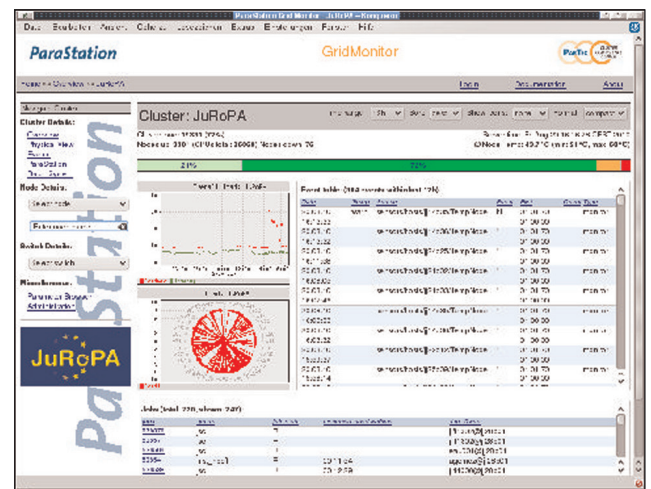
- compute nodes,
- file servers and front-end nodes,
- network devices,
- baseboard management controllers,
- storage devices or
- runtime systems.

Available parameters not only include operating system values, like system load, network counters or temperatures, but also parameters supplied by runtime systems like batch queuing systems or ParaStation. If available, they may even include information provided by network switches, rack or room environment monitoring devices, uninterruptible power supplies or similar devices.

Each parameter is cached within the collector for the configured decay time. Using intelligent caching algorithms, multiple reading of data is avoided.

Data can be stored to and retrieved from a database. Thereby, a data history is available, e.g. for plotting diagrams. Parameters and sample frequencies can be configured independently. 'Virtual' parameters can be computed, monitored and stored to the database based on actual read data, e.g. the total system load as sum of all node load values.

Each known numerical value can be compared against an upper and lower limit. In case this value under-runs or over-runs those limits, an event will be generated. To constantly monitor these parameters, reading cycles can be defined.



Events describing abnormal situations within a cluster can be generated by monitoring parameter limits, node availability, etc. Events will be stored within the database and reported by email.

Beside the actual data, the *collector* also provides information about the type of available data. Based on this parameter type system, it's easy for a graphical user interface to construct dynamic selection boxes without actually reading the data and thereby wasting network bandwidth and compute cycles. This parameter type system also enables a user interface to include new parameters without modifying the scripts or page layout. E.g., newly added nodes will be recognized and shown automatically.

Agents

All data collected by the *collector* is retrieved using agents. Dedicated agents for different services and devices are available. Agents may be installed on monitored nodes (see package *pscollect-agents*) or on the collector node to gather information from remote devices.

Agents are available for

- host parameters (psbase),
- accounting parameters (psacc),
- IPMI parameters (psipmi),
- PBS parameters (pspbs),
- SNMP parameters (built in to collector),
- ParaStation process management parameters (pscollect_psi)

Agents run as non-privileged user *pscd*. If necessary, particular commands are executed as user root using the *sudo* command.

Fast import

Sometimes, importing data via *pscollector* agents infrastructure may not be feasible, especially when data comes in at high frequency. For these cases, a fast-import tool exists that allows loading large chunks of key/value pairs (together with time stamps) to be loaded into the collector's round-robin database at once.

Database connectors

At some sites, interesting data like hardware performance counters, or energy consumption measurements may already be available in an external database. The collector allows on-the-fly reading and visualizing data from such databases, without special agents.

Currently, reading from SQLite, MySQL and Postgres databases is supported.

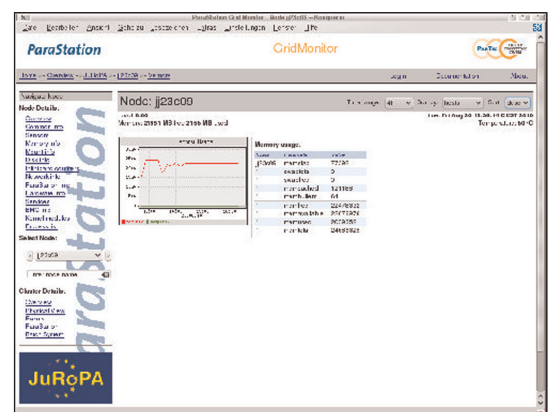
Graphical client (GUI)

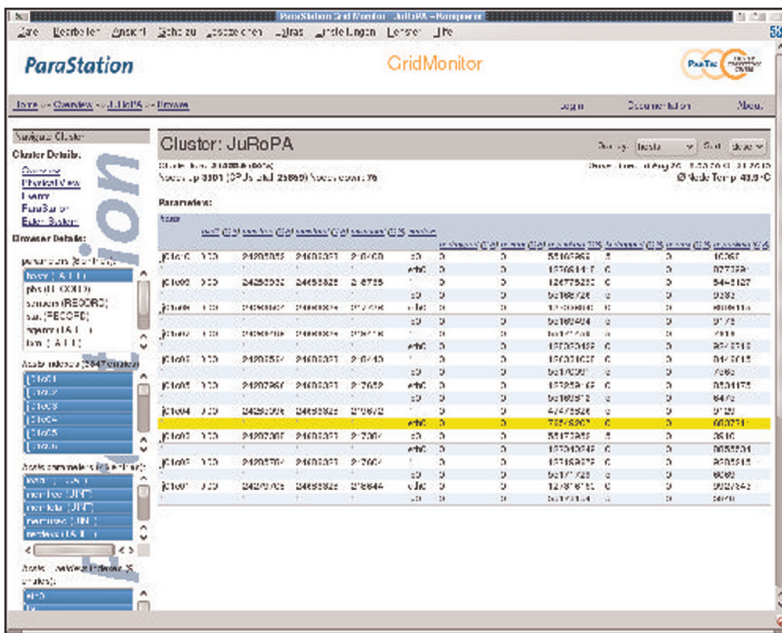
The graphical user interface (GUI), based on a web server and PHP scripts, gives comfortable access to the data, provided by one or more *collectors*. The information is grouped within various views, like multi-cluster or cluster dedicated overview and details, node overview and details, or cluster-wide parameter lists. Each web browser may be used to display this pages.

Views may be modified by pull-down menus defining the time range for lists and graphs, which icons to show, sort order, or columns to sort by.

The graphical user interface also provides information about currently pending events and event history. Parameters can be graphically shown as history diagrams for periodically sampled data, or bar and radar charts for current data.

Beside the predefined data views, a "parameter browser" is implemented within the graphical user interface. The parameter browser is a generic tool to display each particular parameter known to the collector.





Parameters are organized in a hierarchical way by using table and record entries. Each table provides indices and parameters, whereas each record only holds parameters. These parameters may be scalar parameters, like integers, floats or strings, or again may be tables or records. Examples for indices are host names or switch names.

Parameters may be provided by each data source known to the collector, therefore the parameter browser is e.g. also a SNMP (or MIB) browser. The data source is transparent to the parameter browser; therefore the user does not have to take care from which source the actual data is read. Arbitrary parameters may be shown together, like network counters of a node and switch.

The selected data is shown using a matrix layout, similar to a calculation sheet. It is

organized in rows and columns, where the scalar parameters and indices show up in columns. Using pull-down menus, the matrix may be sorted increasing or decreasing by columns. Column data may be visualized using diagrams.

Software packaging

The *GridMonitor* is available as four RPM packages, ready for installation. These packages are:

- *pscollect*: includes the central collector, database and local agents,
- *pscollect-agents*: commonly used agents installed on monitored nodes,
- *psgridmon*: comprises all PHP scripts and configuration files for the web server providing the graphical user interface,
- *psgridmon-doc*: provides all the documentation (HTML, PDF and man pages).

The *pscollect* package is installed on the management node, whereas the *pscollect-agents* package is required on all cluster nodes. The *psgridmon* packages must be installed on the web server node. Installing the *psgridmon-doc* packages is not required, but highly recommended. In particular, installing this package on the web server will provide online documentation within the *GridMonitor*.

Documentation

The documentation is available as a separate, installable package called *psgridmon-doc*. It includes:

- *GridMonitor Administrator's Guide* (English, PDF and HTML format),
- Manual pages describing the *GridMonitor* programs (English, groff format).

Supported Environments

Cluster nodes, management node

The management node requires a recent Linux version. Supported are x86 and x86_64-architectures. For cluster nodes, all versions and architectures are supported. For a cluster node, the GridMonitor requires a Linux kernel version 2.6, but no particular distribution or architecture.

The management node does not require an installed database, as this is part of the *pscollect* package.

Network switches

Additional monitored devices, in particular network switches, have to provide information using SNMP. Standard MIBs (SNMPv2-MIB, IF-MIB) must be supported. Manufacturer specific MIBs may be supported upon request.

Baseboard Management Controllers (BMCs)

BMCs providing an IPMI V2.0 interface may be queried to retrieve system information for servers like SDR, SEL or FRU, independently of the state of the server.

Other monitored devices

More devices like RAID arrays, Infiniband switches or similar providing SNMP-based status information may be supported upon request.

Runtime environments

The *GridMonitor* supports the *ParaStation* process environment. Job information will be shown within the appropriate views. In addition, OpenPBS, Torque and PBSPRO batch queuing systems are supported. Information about jobs and queues managed by those systems are shown within the *GridMonitor*.

Web server, Browser

The *GridMonitor* requires an Apache web server, version 2.0 or later, providing mod-php support. Currently, only Linux-based environments are tested.

A decent browser may be used to display the information. Javascript is required for navigation, but no frames are used.

License

The *GridMonitor* is copyright © by ParTec Cluster Competence Center GmbH. Usage is only permitted with a valid license agreement.

Media

All *GridMonitor* packages are available in binary format for x86 and x86_64 architectures. The packages are available in the download area of ParTec Cluster Competence Center's web server www.par-tec.com.

Support

After signing a support contract, support for all packages is granted for the agreed period of time. The maximum response time is one working day. Support is performed by telephone, email, and/or remote login. Onsite support at the installation site is not included.

The support comprises all *GridMonitor* components as well as the open source software utilized (if applicable).

Scope of delivery

The *ParaStation GridMonitor* comprises of the following components:

- Software packages (see above),
- Documentation (*ParaStation GridMonitor Administrator's Guide*) in electronic format,
- Support as agreed in the support contract.

Copyright

ParTec, ParaStation and *ParaStationVS* are registered trademarks of ParTec Cluster Competence Center GmbH. All other product and brand names are trademarks or registered trademarks of their respective owners.

The information in this version of the software product detailed description is valid as from the time of publishing. Errors & omissions excluded.

Further information

For further information about the *ParaStation GridMonitor* visit <http://www.par-tec.com> or send an email to sales@par-tec.com.