

Hyperic Test Ride Guide

3.1 Edition



Hyperic, Inc., 609 Mission Street, 4th Floor, San Francisco, CA 94105

Copyright © 2007 Hyperic, Inc.

All rights reserved.

Printed in the United States of America.

No part of this publication may be stored in a retrieval system, transmitted, or reproduced in any way, including but not limited to photocopy, photographic, magnetic, or other record, without the prior agreement and written permission of Hyperic, Inc.

Other identifiers may be trademarks of their respective owners.

U.S. GOVERNMENT RESTRICTED RIGHTS. Programs, Ancillary Programs and Documentation, delivered subject to the Department of Defense Federal Acquisition Regulation Supplement, are "commercial computer software" as set forth in DFARS 228.0202, Commercial Computer Software and Commercial Computer Software Documentation, and as such, any use, duplication and disclosure of the Programs, Ancillary Programs and Documentation shall be subject to the restrictions contained in the applicable Siebel license agreement. All other use, duplication and disclosure of the Programs, Ancillary Programs and Documentation by the U.S. Government shall be subject to the applicable Siebel license agreement and the restrictions contained in subsection (c) of FAR 52.227-19, Commercial Computer Software - Restricted Rights (June 1987), or FAR 52.227-14, Rights in Data – General, including Alternate III (June 1987), as applicable. Contractor/licensor is Hyperic, Inc., 609 Mission Street, 4th Floor, San Francisco, CA 94105.



Test Ride Guide Table of Contents

5	Navigating Your Dashboard
6	About This Guide
7	Hyperic HQ 3.1 Product Overview
7	Hyperic HQ 3.1
7	Hyperic HQ 3.1 Enterprise
9	Glossary
11	Start Your Engines
11	Setting Up Hyperic HQ
11	Login to HQ
11	Import a new platform
11	View platform metrics
13	Enable Event/Log Tracking
13	Event/Log Tracking Setup
14	View the Event/Log Entries
15	Configuration Tracking
15	Configure an Alert
16	The Scenic Route
16	Operations Dashboard
16	Auto-Discovery
17	Creating a Dashboard Availability Summary
18	Favorite Resources
18	Recent Alerts
18	Metric Viewer
18	Organizing Resources
18	Auto-Group
19	Compatible Groups
19	Mixed Groups
20	Applications
21	Monitoring

21	Detailed Metrics
22	Monitoring Network Services
24	Monitoring Agentless SNMP Devices
25	Monitoring a Terminal Services Plugin
26	Charts
28	Alerts
28	Escalations
29	Alert Center
30	RSS
30	Control Actions
32	Fully Loaded Model
32	Support for Roles
32	Role Based Access Control
33	Role Based Alerting
33	Enterprise Alerting
33	Multi-conditional Alerts
33	Global Alert Templates
34	Recovery Alerts
34	Group Based Alerting
35	Follow-the-sun Alerting
35	Scheduled Control
35	Report Center
36	Smart Trend Auto-Baseline Feature
36	External Authentication
36	High Availability
36	Failover Clustering
36	Linear Scalability
37	Fahrvergnügen

Navigating Your Dashboard

The screenshot shows the HYPERIC dashboard interface. The top navigation bar includes links for Dashboard, Browse Resources, Alert Center, Report Center, Administration, Recent Resources, and Help. The main content area is divided into several sections:

- 1. Navigation for entire HQ application:** The top navigation bar.
- 2. Search by resource type:** The 'Search Resources' section with a search bar and filters.
- 3. Saved Charts for quick view:** The 'Saved Charts' section showing a chart for 'hq.hyperic.net: Free Memory'.
- 4. Availability counts for resource types:** The 'Availability Summary' table showing counts for various resource types.
- 5. Recently added resources:** The 'Recently Added' section showing a list of resources.
- 6. Add content portlet:** The 'Add content to this column:' section with a 'Select Portlet' dropdown.
- 7. Recent alerts scroll:** The 'Recent Alerts' section showing a list of alerts.
- 8. Auto-discovery for resources to import:** The 'Auto-Discovery' table showing discovered resources.
- 9. Favorite resources:** The 'Favorite Resources' table showing a list of favorite resources.
- 10. Recent alerts:** The 'Recent Alerts' table showing a list of alerts.
- 11. Recent control actions:** The 'Recent Control Actions' table showing a list of control actions.
- 12. Recent control frequency:** The 'Quick Control Frequency' table showing the frequency of control actions.
- 13. Problem resources:** The 'Metric Viewer Demo CPU Util' table showing CPU utilization metrics for various resources.
- 14. Add content portlet:** The 'Add content to this column:' section with a 'Select Portlet' dropdown.

1. Navigation for entire HQ application	8. Auto-discovery for resources to import
2. Search by resource type	9. Favorite resources
3. Saved Charts for quick view	10. Recent alerts
4. Availability counts for resource types	11. Recent control actions
5. Recently added resources	12. Recent control frequency
6. Add content portlet	13. Problem resources
7. Recent alerts scroll	14. Add content portlet

- denotes RSS feed to import into your reader of choice



About This Guide

Welcome to the Hyperic HQ 3.1 Test Ride Guide!

At Hyperic, we realize that the beauty of open source is not just that it's free, but that it's accessible and useful. You, the user, want to be in the proverbial drivers seat. You can download it, get it running, and enjoy the beautiful ride it gives you all by yourself. If you were out buying a car, you wouldn't need someone to teach you how to drive. But you may want someone to give you a quick orientation on the car to tell you all about the incredible new features it has. That's why we wrote this guide.

The idea is we'll give you an overview (in the cleverly named Hyperic HQ 3.1 Product Overview section). Then we'll get you through some basic steps for installation. This is the Start Your Engines section, which really, is available on our documentation site, but we copied it here just to make your life easy. What can we say? We like you already. Next up is the Scenic Route section, where we discuss some of the major features and provide some sample activities. This isn't an all-inclusive map though! Nor do you have to follow it like a cookbook, every step of the way. Choose your own adventure! You're in the drivers seat now!

Finally, at the end, we'll give you a preview of the Fully Loaded Model. This is a high level overview of Hyperic HQ Enterprise. Built with the same style and same engine, this hot little number sports some extra bells and whistles that are a must-have for organizations of scale. For that full tour though, you'll need to download the enterprise trial.

Finally, all automobiles need service now and then. The beauty of open source is though, that you don't always have to go to the dealer to get it. The last section, Fahrvergnügen, will introduce you new driving enthusiasts to the power of the Hyperic Community.

So, ladies and gentlemen, get ready and start your engines!

Good luck, and happy driving.

-Hyperic Team



Hyperic HQ 3.1 Product Overview

Hyperic HQ is the only systems and applications management software designed and built expressly for the fast-moving online services market, where the demands of SOA, virtualization and composite applications are changing the pace and complexities of managing the data center. Everything is more compressed, more utilized, more integrated, more critical. The traditional tools IT operators used to monitor systems feebly cope with this new market landscape.

Hyperic enables IT managers to expand past classical monitoring activities to drive availability and improve the overall health of their IT infrastructure. Purpose-built for web infrastructure, and architected to consider all layers of infrastructure including hardware, middleware, virtualization and applications, Hyperic HQ delivers system monitoring, trending, and analysis with one click of a button. Hyperic HQ is the first and only enterprise-class open source systems management software that lets administrators manage the IT technologies in the market today as well as those appearing tomorrow.

Hyperic HQ 3.1

Designed to supply maximum visibility for web infrastructure with minimal overhead, Hyperic HQ software delivers comprehensive systems management. Available through the open source GPL 2.0 license, users of Hyperic HQ 3.1 can:

- Auto-Discover asset inventory with one click
- Monitor metrics of 65+ technologies across 9 OSES & 4 architectures
- Track performance, configuration and security changes
- Manage availability with alerting and corrective control actions to address problems before they occur
- Extend, customize visibility to best manage your unique environment

Hyperic HQ 3.1 Enterprise

Hyperic HQ Enterprise provides extensive systems management coverage that saves time and money when managing large infrastructure environments and operations teams. Hyperic HQ Enterprise includes extended automation and control designed to help enterprises achieve operational efficiency, including:

- Role-Based Access Control for managing user permissions and visibility
- Follow the sun alert assignment notifies users based on time of day and user's scheduled availability
- Global Alert Templates create alert definitions once and apply across large resource groups
- Enterprise Alerting includes multiconditional and recovery alerts, as well as role-

based and group-based alerting

- Scheduled Control enables administrators to automate maintenance activities and conditional responses to alerts
- Smart Trend baselining automates metric analysis to flag deviations from normal performance for fast problem ID
- Enterprise Reporting provides management easy analysis of performance and service level activity
- External Authentication leverages existing LDAP or Kerberos directories
- High Availability allows multiple HQ Servers to assume workload for one another in the event of a downed server
- Distributed monitoring allows for linear scalability of the HQ Server

HQ Enterprise is available as a free trial for download directly from Hyperic under a Commercial License. The enterprise trial does not expire, but is limited to three managed devices. For businesses seriously considering Hyperic HQ Enterprise, a full Enterprise Evaluation is available where for 60 days, you can enjoy all the benefits of the enterprise features on an expanded set of resources, two support calls and a one hour meeting with our Technical Account Manager staff. To find out more about this, contact Hyperic Sales at sales@hyperic.com or call 1-888-4Hyperic.

This guide is designed to demonstrate the capabilities of the Open Source product. However, as applicable, we will call out where additional items and capabilities exist in HQ Enterprise. We'll keep that contained in the Fully Loaded Model though, so it'll be easy to tell what's in the open source software.



Glossary

Before you begin, it's important to establish some common definitions of terms so you can better understand the information now visible to you through Hyperic HQ. Systems management offers us many different definitions for the same terms depending on context. The Hyperic definitions are important to consider as we group information for monitoring and analysis within the application.

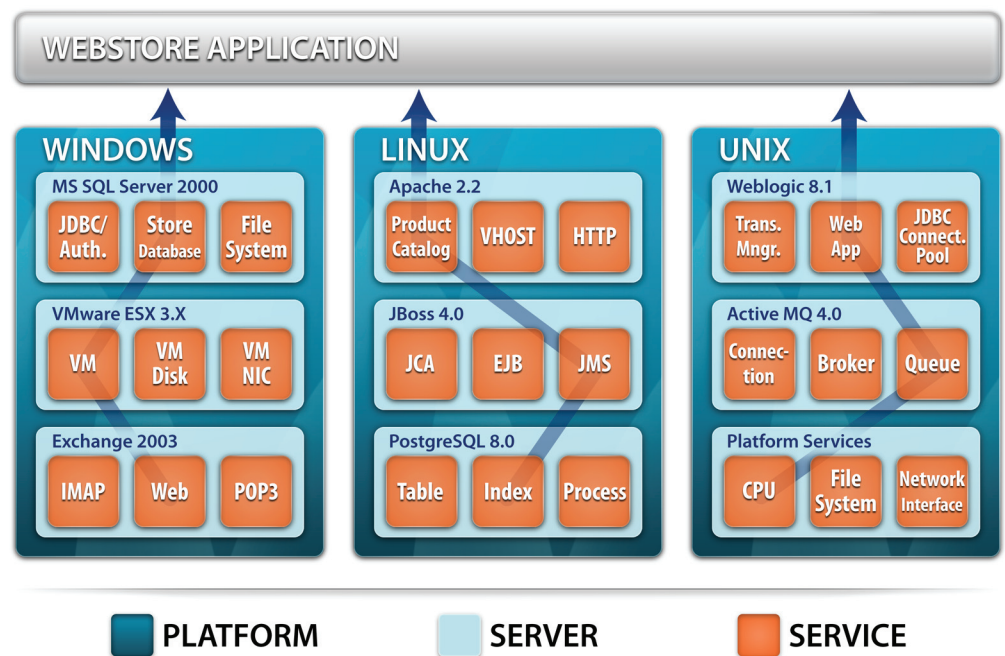


Figure - components of a webstore application illustrate the definitions of Platform, Server and Service within HQ.

- **Application** – A collection of platforms, servers and services organized to fulfill a single business purpose. Applications typically have many servers and services and can also run on more than one platform. A typical J2EE application modeled in HQ can contain Apache Virtual Hosts, Tomcat Webapps, JBoss Connection pools, and Oracle Instances.
- **Platform** – A machine/operating system combination or any network or storage device. Platforms are the lowest level of management, and include components such as CPU's, Network Interfaces, and File Systems.
- **Server** – Any software that is installed on a platform under management. Databases, middleware, virtualization, application and web servers are all examples of servers. Servers run on platforms. Platforms host multiple Servers. Examples of servers include any installations of JBoss, Tomcat, or MySQL on a given platform.
- **Service** – A component of a server dedicated to a specific purpose. Typically services are represented by the units of work of a given server. Different types of servers each define a list of one or more types of services they provide. Examples of services include Webapps deployed in Tomcat, or Virtual Hosts configured in Apache.

Services can also be attached directly to a platform in the case of CPU, Network Interfaces and Filesystems.

- **Resource Type** – A categorical description of a type of resource that can be managed. Upon installation, Hyperic HQ describes hundreds of resource types which include specifications for Name, Version, Available Metrics, and Available Control Actions. Resource types span Platforms, Servers, and Services.
- **Managed Resource** – A specific instance of a resource type within the managed environment.



Start Your Engines

The purpose of this section is to get your installation of Hyperic HQ up and running. Starting your engines, if you will. The next section, The Scenic Route, will revisit some of these setup sections in more detail, and point out some other areas of interest.

Setting Up Hyperic HQ

General set-up of the Hyperic HQ product has not changed. The getting started guide has been copied in this document for your convenience and is commonly available on the Hyperic Documentation site.

Login to HQ

Using a web browser, navigate to port 7080 of the host you installed HQ on. An example might be `http://localhost:7080` or `http://hq.myhost.net:7080`. You will be presented with the HQ Login screen. Enter the default HQ Administrator username and password of `hqadmin` and `hqadmin`. Upon successful login, you will be at the HQ **Dashboard**.

Import a new platform

The **Dashboard** consists of several portlets that provide different types of information. The one portlet that will stick out is the Auto-Discovery portlet. It will have several entries for the newly registered HQ Agent on the local machine. The radio button next to this host should be selected. If there are other entries and the radio button is not selected, please select it. Now click Add to Inventory to import the platform and its servers. The process adds the platform and servers to the HQ inventory as well as enables the default metrics for collection on the newly imported resources. You will know this process has been started when you see the hostname for the platform in the **Recently Added Resources** dashboard portlet.

View platform metrics

The default metric collection intervals for platform metrics range between 1 minute and 10 minutes depending on the metric. So since we just scheduled these metrics for collection, we need to give it a few minutes to collect the first data points. After the first metrics have been collected, we can view them from the platform metrics page.

To view the metric data, it is necessary to navigate to the platform's HQ Indicators page. There are several ways to do this. The easiest way to do this is to use the **Recently Added Resources** portlet on the **Dashboard**. All new or updated platforms will be listed on this portlet. Alternatively you can click the **Browse Resources** link at the top of most any page. Next click on the link of the platform you wish to view the metrics for. This is the **Platform Indicators** page. This page displays a wealth of information on the health of the current resource (in this case, a platform) as well as information on resources related to this platform.

The indicators page is used to define correlated “indicator views” that include metrics from various resources plotted along a common timeline. By default, the indicator view of this platform will include a handful of key metrics from the machine itself. The bottom portion of the indicators panel (a black bar with dots) defines a clickable timeline which represents individual time slices. Each indicator chart displays a blue rectangle whose top and bottom indicate the high and low value observed for that metric and a dark blue line indicating the average.

Indicator views can be customized to help diagnose and understand the interaction between different resources. As an example, let’s say you want to build a view that correlates network interface performance with other metrics from the platform.

In order to do this:

1. Select the checkbox next to the “Network Server Interfaces” found under the “Platform Services Health”.
2. Now click the **View Metrics** button to bring up the list of available metrics for the platform, as well as the network interfaces. Scroll down the list of metrics to find “Packets Transmitted Per Minute”. Click on the arrow icon to add the metric to the bottom of the current indicator view. You can choose to move the individual indicator charts up and down to your chosen order
3. Finally, save your custom view as the default by clicking on the arrow icon in the top right hand corner of the indicator view. You can create custom views with different names. This will help you correlate multiple combinations of metrics and help you diagnose the interactions between different aspects of your environment.

Now that you’ve defined a custom view, you can look at the indicators page to see how metrics such as the 5 Minute Load Average correlate to the amount of network traffic being put through the network interface(s) of your platform.

In order to get more specific metric information or to enable more metrics for collection, click on the Metric Data mini-tab. This list of metrics represents all the metrics for which HQ has data in the current display range. In our case, we just started metric collection a few minutes ago, so there will be very little metric data here. There are still a few things you can do to explore the possibilities in HQ. Each metric has an information icon on the right side of the page. Clicking this will open a window with detailed information on the specific metric. You can chart the data for any metric by clicking on the metric name. Alternately, you chart multiple metrics at the same time by checking the checkbox next to the desired metrics and clicking the Chart Selected Metrics button.

By default, only the “indicator” metrics are configured to be collected. To enable more metrics for collection, click the arrow button next to Show All Metrics at the top of the Metric Data page. This will list all metrics available to collect for the resource type. To

enable metrics or change their collection interval, simply check the box next to the metrics you wish to change, enter the collection interval at the bottom of the page and click the arrow button.

Enable Event/Log Tracking

1. Often an IT problem can be traced to the logs of the operating system, application server, or middleware. These logs are scattered throughout the infrastructure depending on the product. HQ provides the Event/Log Tracking feature to centralize this data to provide the following benefits:
2. Report real-time and historical details of any log event generated by any managed resource
3. Track the configuration of any host or application to facilitate impact analysis and change control
4. Enforce security policies by detecting and logging any physical or remote access into any host in your environment
5. Define alerts to track to specific log messages anywhere in your environment
6. Correlate any log event, configuration change, or security event to the availability of your environment

Notes on Log and Configuration Files:

- Multiple files may be specified in the Configuration and Log Files fields. Separate multiple files with commas.
- The files must be readable by the HQ Agent, so make sure file permissions are set accordingly.
- Not all resources are capable of supporting configuration tracking. Typically this is limited to the platform or server.

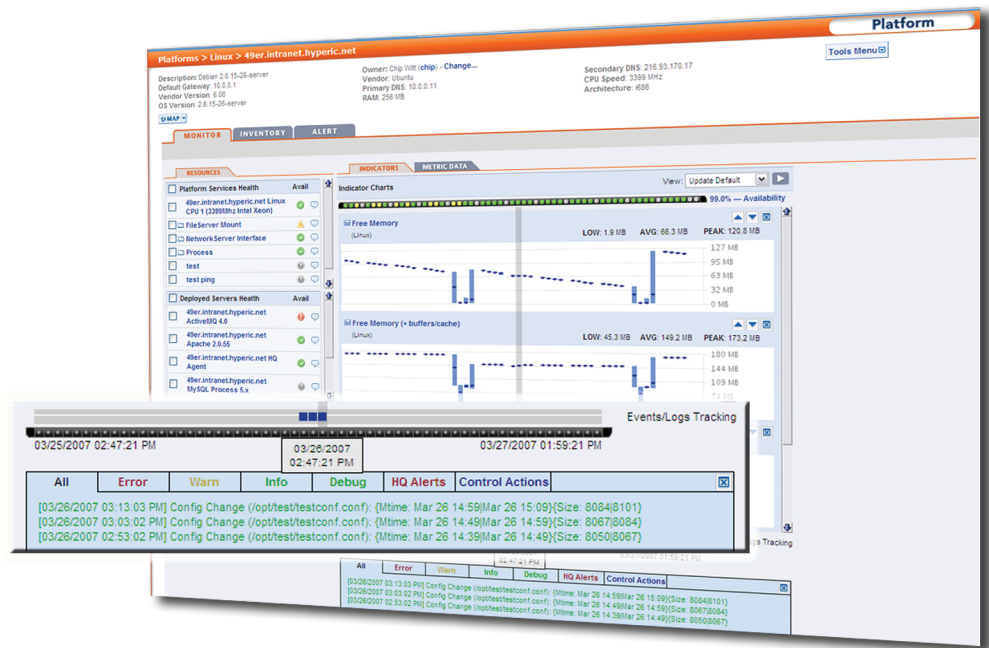
Event/Log Tracking Setup

1. From the **Dashboard** or **Browse Resources** select the Platform that contains the log you want to track.
2. "Tools Menu" found at the top right of the screen and select "Configure Platform"
3. Check the "Enable Log Tracking" checkbox.
4. Specify the full path to the log you want to track in the "Log Files" field and select the log level (Error, Warn, Info, Debug) you want to track.
5. Click OK to save your changes.

View the Event/Log Entries

5. Go to the Monitor tab for a platform
6. Look at the bottom time line in the Indicator charts section of the page
7. If any Event or Log entries are available, they appear in a gray area labeled Event/Log Tracking above the bottom time line. Blue squares represent time slices that contain Event/Log data. Note that if there haven't been log items matching your log level, there may not be any events available to view.
8. Clicking on the blue square expands the Event Tracking information box, which details the events that occurred (including the event/log or configuration file change for which you enabled for tracking) within the time slice displayed. See the screen capture below.

Figure - clicking on the blue squares shows event and log detail, providing critical activity detail necessary for root cause identification.



Optionally, you may want to setup an Alert based on the contents of this log:

9. Select the same Platform that you configured in the steps above
10. Click the "Alert" tab
11. Click "New" to create a new alert
12. Select "Events/Logs Level:" and the level you want, choose Any if you are not sure which level to choose
13. In the text box next to the log level put in the string you want to alert on
14. Click OK to save your alert

Configuration Tracking

1. Browse to the resource for which you want to establish configuration tracking in the HQ UI.
2. Click on the **Tools Menu** found at the top right of the screen and select **Configure Platform**.
3. Check the box to enable configuration tracking (config_track.enable or server.config_track.enable), and specify the configuration files you wish to track (config_track.files or server.config_track.files) as a comma-delimited list.
4. Click "OK" to set the changes. The files must exist before you set this configuration, and the user under which the agent on the machine is running must have permissions to at least read the files you are specifying.

Collection of the change information occurs every 15 minutes. When a change to a specified configuration file occurs, it shows up as a blue square on the bottom of the Indicator tab of the Monitoring page for the resource where the tracking was enabled.

Configure an Alert

To configure an alert for your platform, click on the Alert tab. Click the New button to create a new alert definition. We will configure an alert to let you know if this platform ever goes down.

Enter Platform Down! as the name of the alert, select !!! - High priority, optionally enter a description and leave Active set to Yes. Next, select Availability in the metric dropdown. Select < (Less than) and enter 100% in the (absolute value) line. Scroll down, leave Each time conditions are exceeded or met selected and click OK to create the Alert Definition.

Now that the alert definition is created, we need to add notification to it. Click on the Notify HQ Users tab and click Add To List. Check your user (select hqadmin if that is the only user), click the right arrow to move the name over to the notification list and click OK to set notification.

Now the email address associated with the selected user will get an alert notification if this platform ever goes down.



The Scenic Route

Congratulations! You now have a running installation of Hyperic HQ 3.0. Now its time to get out in the open and drive like you want to drive. The following sections outline some of the local scenic spots, and can be reviewed in any order. They are designed to give you a more in-depth look at Hyperic HQ's most important features, but feel free to explore even more on your own!

Operations Dashboard

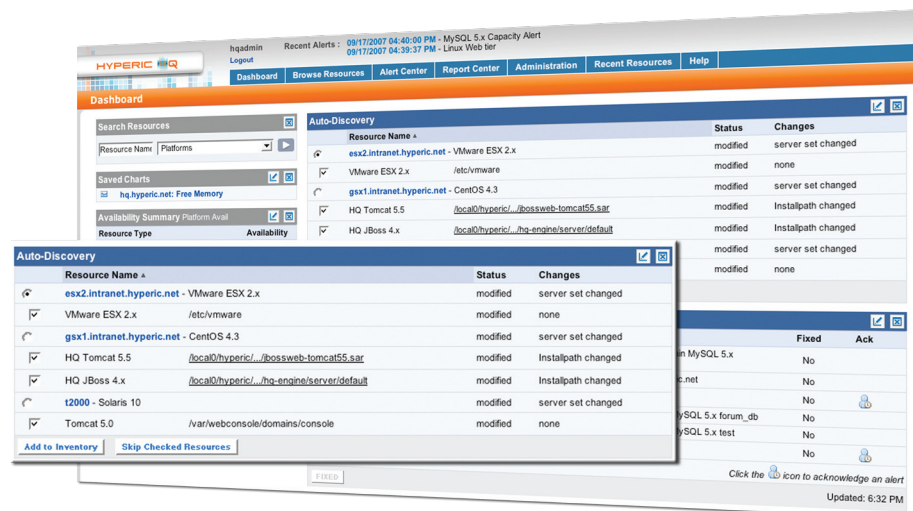
The first view you immediately see when logging in is the HQ Operations **Dashboard**. To get back to this **Dashboard** at any time, just click on the Dashboard link in the application masthead. The information on this page lets you know which of your resources are down, the alerts that need your attention, and show you at a glance detailed metric data.

The **Availability Summary** in the left hand column is usually the first place administrators look at to assess their environment. For more on configuring this summary portlet, continue to the following sub-section "Creating a Dashboard Availability Summary". This portlet groups resources by type and shows us what is healthy. Any red icons or text in this portlet is cause for immediate concern.

In the main frame on the right hand side are more portlets including **Auto-Discovery**, **Favorite Resources**, **Recent Alerts**, **Metric Viewer**, and **Control Actions**. If you haven't yet configured them, these portlets will be empty. The order of these portlets can be rearranged into a custom layout by drag and drop. They can be deleted by clicking the X icon in the top right of the portlet, or add them by selecting the type and clicking the + icon at the bottom of each column. The pencil icon on the portlets allows you to edit the settings of the specific portlets to populate or modify your portlets.

Auto-Discovery

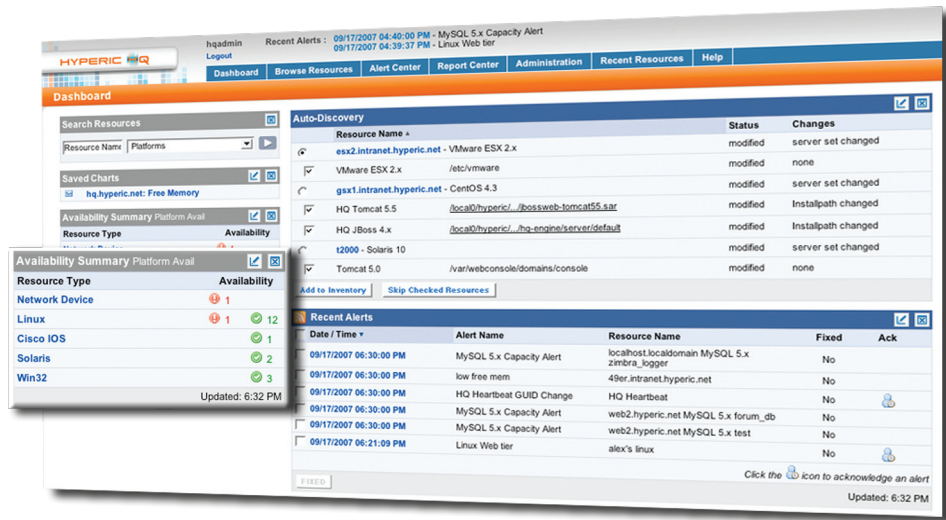
Figure - Auto-discovery portlet shows available resources to be managed. Adding these to inventory immediately starts metric collection on these items.



As described in the Getting Started section, once you deploy an agent, this portlet will auto-populate any new managed resource on that platform. In order to start managing the available resources, you need to import them. This portlet will remain empty until new resources are discovered. If your environment is not expected to see frequent changes, you may want to drag this portlet to the bottom of the list on the **Dashboard**.

Creating a Dashboard Availability Summary

Figure - Availability Summary provides quick reference for the general state of your IT infrastructure's availability.



To create a new summary that will contain all your resources, summarized by type, and will be displayed indicating the number up (green) and the number down (red) on the **Dashboard** do the following:

1. On the bottom of the left side of the **Dashboard**, select **Availability Summary** from the drop-down menu, and click on the "+" sign to the right of the drop-down.
2. Click the edit icon (pen graphic) on the new **Availability Summary** portlet.
3. Click on **Add to List...** in the **Selected Resources** section.
4. Filter the view of your entire resource inventory using the drop-down lists containing resource categories and types.
5. Select the check-box next to resources you wish to add to the summary, and click the blue arrow (>) to move the selected resources from the **Resources** table to the **Add Resources** table. Do this for each page of resources until you have added all the resources you want to add.
6. Once all your desired resources are added, click "OK".
7. Enter a description (name) for your **Availability Summary** in the description field, then click "OK".

Favorite Resources

This portlet lets you pick and choose from individual platforms, servers, and services. If you click the edit icon for the Favorite Resources portlet you can add or remove items. Clicking the Tools Menu, which can be found on any specific Resource's page, and clicking on 'Add to Dashboard Favorites' will add Favorite Resources to your Dashboard Favorite's list.

Recent Alerts

This portlet lists the Recent Alerts that have fired on your system. If you click the 'Edit' icon on this portlet you can configure it so that only the High priority alerts are listed. Or, if you prefer, set it to include Medium or Low level alerts instead. The most recent alerts are also always displayed on the application masthead regardless of where you are in the application.

Once you have acted on an alert, perhaps you rebooted a machine that was not responding, you can use the Recent Alerts portlet to mark that Alert as "Fixed." This is a sign to every one else using Hyperic HQ that the problem has been resolved.

Metric Viewer

You can have multiple Metric Viewer portlets on the Dashboard each reporting on different aspects of your network. Metric Viewer portlets can be arranged to look at specific resource types and specific metrics, such as CPU usage.

Organizing Resources

Organizing Resources

Auto-Group

An auto-group is created by HQ when a resource contains a collection of like services that can be rolled-up into a single line-item link. A perfect example is on a network switch Platform. A switch can contain many interfaces, and to list them all one-by-one at the Platform page would make it hard to manage other data easily. By rolling all the interfaces up into an auto-group, the user receives immediate visual status of everything in the group. The user can drill-down by following the link to get an aggregate view of collected metrics across the entire group of resources, and can drill-down further to gather additional detail on the individual interfaces when necessary.

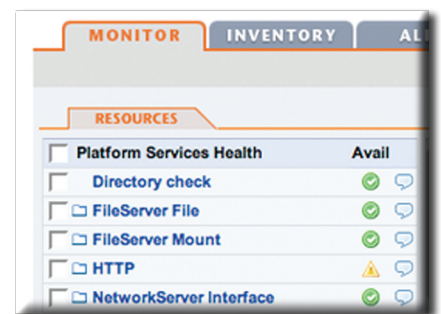
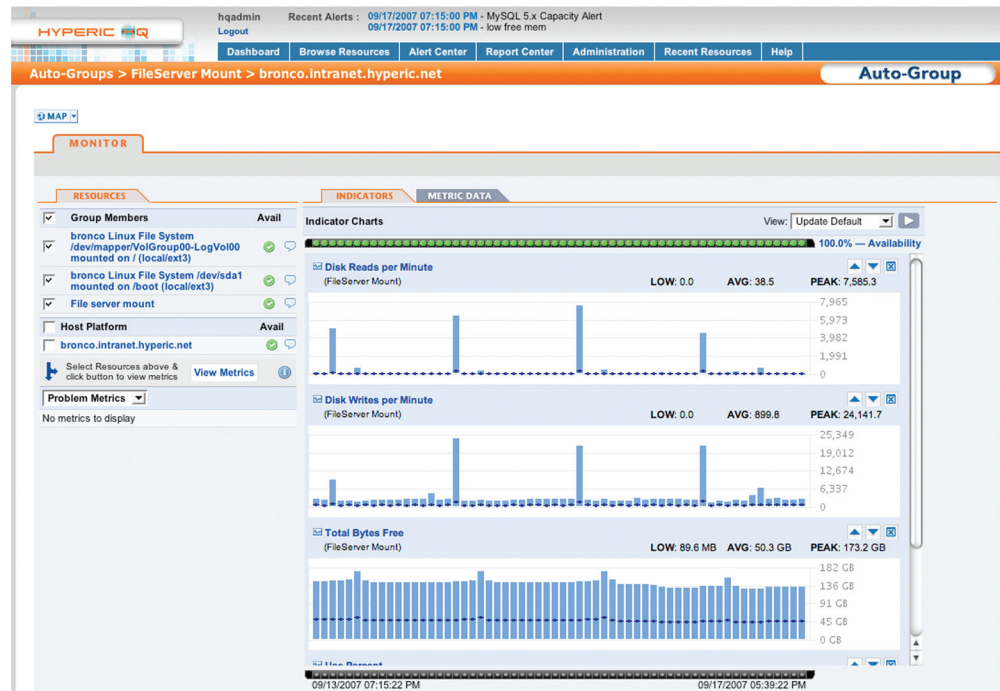


Figure - Auto-group provides a quick way to group like resources for monitoring. Folders to the left of the resource indicates an auto-group.

Figure - Monitoring auto-groups aggregate metrics and demonstrate overall performance of a bank of resources.



Compatible Groups

You can create groups of resources through the **Browse Resources** pages for Platforms, Servers, and Services. Find the resources you want to group together and check the checkbox next to those resources, then click “Group” at the bottom of the table. If you have more than one page full of resources, you can use the drop-down menu at the bottom-right of the table to show “All” resources.

HQ will intelligently group items you select: if the resources have common metrics, they are classified as a “Compatible Group / Cluster”.

When you drill-down into a compatible group, you will see indicator charts in much the same way you do when viewing individual resources, but for the group metrics for all resources defined in the group are aggregated together. If you drill-down further into a group indicator chart, you can see detailed metrics for each resource all on one chart for that metric. You can also “Save Chart to the Dashboard” for easy reference later by clicking on the button available at the upper-right of the detailed chart view.

Mixed Groups

If the resources do not have common metrics or control actions, they are categorized as a “Mixed Group”. Only availability is shown for mixed groups.

Applications

Hyperic uses the term “Application” to describe a collection of services which, together, fulfill a single business purpose. This concept is used to represent the idea that an application, from the business point of view, is composed of many different pieces, and those pieces are usually distributed across different platforms and provided by different servers. HQ’s user interface offers users the ability to manage infrastructure from an application point of view as opposed to a hardware point of view. A typical J2EE application modeled in HQ can contain Apache Virtual Hosts, Tomcat Webapps, JBoss Connection pools, and Oracle Instances.

Through this very powerful concept, IT Folks are able to manage beyond the basic system-level, and gain an enterprise perspective on the resources for which they are responsible.

In order to setup your application, you must first create your empty application container. Since this is a custom collection of platforms, servers and services, Hyperic can not auto-discover your application for you. To create the application:

1. Select **Browse Resources** from the HQ Dashboard masthead.
2. From the **Browse Resources** page, select New Application from the **Tools Menu** in the upper-right-hand corner. A name for your application definition is required. Optionally, you can add information about Location, Description, an Engineering Contact, a Business Owner, and an IT Operations Contact.
3. When you click OK, you are immediately taken to the Inventory tab of the new application.

Next you must add services to your new application:

4. On the Inventory tab for your new Application, you should see the fourth header down labelled “Services”. At the bottom of the now empty Services table, click on the button labelled “Add to List”.
5. You are taken to a page that will contain several pages of services about which HQ has information. Check the boxes next to the services you wish to list as dependencies for your Application.
6. Click the blue right-pointing-arrow to move them from the Services list to the Add Services list. You can select several services at a time from the page, but can see only one page of 15 services at a time during this process.
7. Once you’ve got all your services in the Add Services list, click OK. The newly added services will be listed in the Inventory page’s Services list.

Back on the Monitor tab, you will likely see only the availability indicator for your defined application in its Charts. If you want to add Indicator Charts for monitoring your new application then we’ll need to do a few more things.

8. If you have added some services, you should also now see some Deployed Services and Host Servers listed under Resources on the left side of the page that match your inventory selections for the application. Check the boxes next to these resources.
9. Click Set Resources to get a list of available metrics for them to appear below.
10. From this list, you can click on the blue arrow next to the available metric to add its indicator chart for the application.

DON'T LOSE YOUR NEW VIEW!

At the upper-right-hand side of the Charts is a View drop-down menu. You can create a new view, give it a name in the value box, and click the white "play" button, or you can just click the white "play" button to update the Default view. Either way, you will not lose the work you've done to specify indicators for the application you've defined. It is common to create several views on the Indicators tab so that different comparisons can be made between metrics under various circumstances. This use of views works for any Indicator tab for any resource defined in HQ, so you may find use for it outside of Applications too.

Monitoring

Typically system administrators start with the **Dashboard** but quickly find information they would like to drill down on.

Detailed Metrics

Hyperic HQ is an excellent tool for answering the simple question of "what resources are up or down in my network?" However, on top of that, it can provide you with an amazing level of detail about the state of your hardware and software.

The following are three examples of the technical detail that Hyperic HQ provides. It's a powerful tool for analyzing the state of resources.

MACHINE_1 -- Swap Used metric

One of your application servers has been performing poorly. The sys admin for this box said to keep an eye on how much Swap space was being used. To configure this, go to the **Dashboard** and in the **Favorite Resources** portlet select MACHINE_1. Click the Metric Data tab that is just to the right of the Indicators tab. This lists all of the metrics we are tracking for MACHINE_1. Scroll down to the Utilization section and look for Swap Used.

SERVER_2 -- JBoss Threads metric

As a preventative measure, you have been told to keep an eye on the JBoss Thread count for your enterprises application server. Too many threads in use is a sign that trouble is on the horizon.

In the **Favorite Resources** portlet select SERVER_2. You will now be on the detail page for our JBoss server. You don't have to go any deeper to answer the question, as the first chart in the Monitor window is the Active Thread Count for JBoss.

SERVICE_3 -- Tomcat requests per minute

Your VP wants to know if the tomcat server is getting any traffic. "How busy is the Tomcat server?" is the question.

Click the **Browse Resources** link and then select the "Services" link that is between "Servers" and "Compatible Groups". Select the pop-up menu that says "All Services Types" and choose the first item that includes "Tomcat Connectors". You will now see the list of Tomcat servers. The final column tells you exactly how busy your web servers are.

Monitoring Network Services

HQ is an agent-based monitoring solution, but it is able to monitor many services over the network without having an agent physically on the hardware running that service. The monitoring HQ does on these network services is a quick check to make sure the specific service is working. Supported services include:

- HTTP
- SSH
- IMAP
- POP3
- NTP
- DNS
- FTP
- SNMP
- LDAP
- SMTP
- DHCP
- RPC
- TCP Socket

The "TCP Socket" service is simply a generic check that a specific TCP port on a platform is listening for traffic. That is useful in the case where you need monitoring on a custom service or a service that is not specifically supported.

To configure HQ to monitor a network service:

1. Determine where the check should originate from. You will be selecting an HQ Agent to do the check, so the network connection to the service will originate with the agent you select. You select this agent by navigating to the platform monitored by that agent within HQ.

2. Navigate to the platform you select to do the check, drop down the **Tools Menu** (blue box near the upper right) and select "New Platform Service".
3. On the resulting page, enter a name for your new service (i.e. "www.mysite.com HTTP check") and in the drop-down, select the resource type (i.e. "HTTP"). You will notice there are many service types in addition to the network services specified above. These are additional services that can be managed at the platform level in HQ, we are just concerned with the network services for now.
4. Click OK to create the service.
5. On the resulting page, you will see a message highlighted in yellow that says "This resource has not been configured. Please set its Configuration Properties". Click on the "Configuration Properties" link in the message to configure the service.

Each type of service will have configuration specific to it. The configuration parameters should be readily apparent to anyone familiar with the implementation of the particular service. In our example, we will configure an HTTP service accessed via the URL "http://www.mysite.com/".

6. We know that it is not SSL (since the URL is "http" and not "https"), so we will leave that checkbox unchecked.
7. Change the hostname to the hostname in the URL ("www.mysite.com"). Leave the port the default "80" since that is the default HTTP port and the URL did not specify a different one.
8. It is safe to leave the timeout at 10 seconds. For our example, we can leave the default path of "/" since that is all the URL specified. If we wanted to check a different path, like "http://www.mysite.com/img/myimage.jpg", we could specify "/img/myimage.jpg" here in the path setting.
9. We will assume the URL is not password protected. If it was, we would simply specify the realm, username and password required.
10. The method can be either GET or HEAD. HEAD will result in less network traffic. However, you may want to use GET if you want to take advantage of the pattern matching feature. Specify a "hostheader" which allows you to set your own "Host" HTTP header in the request. This is useful for name based virtual hosts, which are common among Apache users.

Note: If the HTTP request we are making will result in a redirect, the redirect checkbox should be checked. This is important, because an HTTP server returns a different code for a redirect and HQ will assume that the HTTP service check is not available if it is a redirect, unless this redirect configuration is set.

11. Next we look at pattern configuration. This allows you to configure a pattern or substring that HQ will attempt to match against the content in the HTTP response.

This is particularly useful as it allows you to verify that not only is your HTTP service running correctly, but it is serving the content you expect it to be.

12. Last is any proxy configuration that may need to be in place for the HTTP request to be successful. We will assume there is none for our example, and click OK to apply our configuration.

At this point, the HQ Agent will attempt to make a request to the HTTP service we just configured in order to validate our configuration. Assuming we set up everything correctly, we should get a green highlighted message telling us that our configuration has been saved. HQ is now monitoring the HTTP service.

Monitoring Agentless SNMP Devices

HQ can monitor just about any network device via SNMP. There is specific monitoring for some, such as "Cisco IOS" and "Cisco PIXOS", and there is generic monitoring for "Network Device" and "Network Host". The types of devices that can be monitored range from firewalls, switches and routers to hosts that can be queried SNMP. All of these things are considered to be a "platform" in HQ and they are all setup and configured in the same general way.

To add a network device to HQ for monitoring:

1. Navigate to either the **Dashboard** or the **Browse Resources** page in HQ. There will be a link for "New Platform" in either the Resource Summary portlet or the **Tools Menu**. Click that link.
2. On the resulting page, you will name your new platform resource (i.e. "myswitch1.mycompany.com") and select the platform type. For this example, we will select "Network Device".
3. Specify the fully qualified domain name (FQDN) for the device (can be the same as the name - i.e. "myswitch1.mycompany.com").
4. Select an HQ Agent that will be monitoring this device. This is important because the agent you select will be where you will monitor this resource from. For instance, you may want to select an agent that is on the same network segment as the device you are monitoring for performance. Or you may want to monitor the device from a specific location in your network. This must be taken into consideration when selecting which agent will be monitoring your network device.
5. Enter the network information for the device (must at least enter an IP address) and click OK to create the resource and add it to your HQ inventory
6. On the resulting page, scroll down and click "Configuration Properties" to open the configuration properties section, then click "Edit" to configure the resource.
7. If your device does not use ifDescr to uniquely identify the interfaces, you can configure HQ to use any of the columns in the drop down to get unique names

for your interface services. The interface index is how HQ will name the interface resources on a network device like a router or switch. Most devices work with the default ("ifDescr") just fine, so we can leave it at that for now.

8. The SNMP IP will be the IP address the HQ Agent will use to contact the device. Give it the IP address where the device responds to SNMP requests, the port number and the SNMP community string.
9. The SNMP version is specific to your device configuration. Many times, the default ("v2c") will work just fine. If you do end up using "v3", that is likely because you have authentication configured. If that is the case, you will need to also configure the username, password and auth type.
10. When you have the configuration set, click OK to save it.

HQ will attempt to validate your configuration, and when you get the green highlighted success message, you will be monitoring your network device.

Monitoring a Terminal Services Plugin

Detailed background and example to show how Windows service plug-ins are constructed can be found on the Hyperic Documentation site.

The following serves as an example of deploying a plug-in, configuring and monitoring a Windows Service, and using Windows performance data.

Deploy the Windows Terminal Services Plug-in:

1. On the windows agent, move `terminalservices-plugin.xml` from `C:\hyperic\agent-3.0\pdk\examples` to `C:\hyperic\agent-3.0\pdk\plugins`, and then restart the agent. The agent log (in `C:\hyperic\agent-3.0\log\agent.log`) will show that it registered.
2. On the HQ server, you can place the plug-in one of two places:
`C:\hyperic\server-3.0.0\hq-engine\server\default\deploy\hq.ear\hq-plugins\`
OR
`C:\hyperic\hq-plugins`
3. Placing the file in the first location will require a server restart. Restart as necessary.

The second location does not exist by default. The benefit of creating the `hq-plugins` directory is that, after the server is restarted once with this directory in-place, you can do "hot" (no server reboot required) server-side deployments for subsequent plugins.

When the plugin registers properly on the server, you should see something like the following in `C:\hyperic\server-3.0.0\logs\server.log`:

```
2007-02-09 15:51:35,193 INFO [org.hyperic.hq.product.server.mbean.ProductPlugin
Deployer] HQ plugin terminalservices registered
```

2007-02-09 15:51:35,452 INFO [org.hyperic.hq.product.server.session.ProductManagerEJBImpl] terminalservices unknown -- registering
2007-02-09 15:51:41,996 INFO [org.hyperic.hq.product.server.session.ProductManagerEJBImpl] terminalservices deployment took: 1.865 seconds
2007-02-09 15:51:42,725 INFO [org.hyperic.hq.product.server.mbean.ProductPluginDeployer] HQ plugin terminalservices deployed

Use the Terminal Services Plug-in:

4. Browse to the Platform page for the Windows machine on which you are wanting to configure the Terminal Services monitoring.
5. On the right, toward the top of this page, click on "New Server".
6. Give it a name, say, "HOSTNAME Terminal Services".
7. In the "Server Type" drop-down, select the new Terminal Services type, which should be there if the plug-in is doing its thing. "Install Path" field is required, put "C:\WINDOWS\System32\svchost" in there. Click OK.
8. You will be told you have to go into "Configuration Properties" to finalize the configuration. Set "Auto-Discovery for Sessions" to On, and click OK.

A few minutes later, the Availability and Active Sessions will show up as indicators under Monitor. You can add "Inactive Sessions" and "Total Sessions" as well through the Metric Data tab (click on "Show All Metrics", and set collection intervals to begin collecting).

When you browse the Platform, you will now see your Terminal Services Server listed as a Deployed Server.

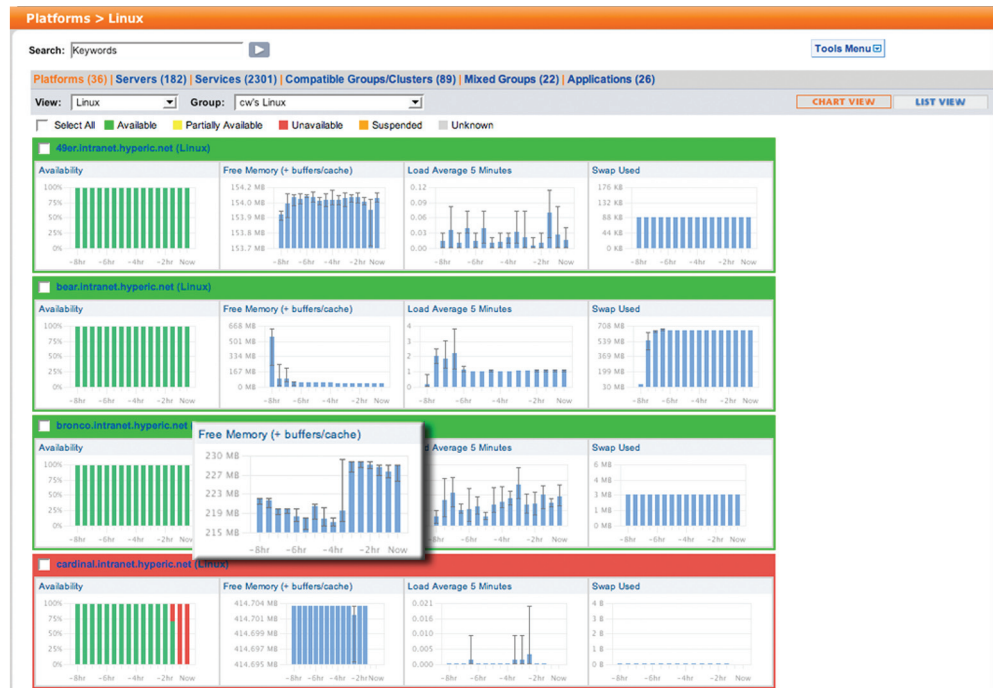
Charts

An alert when box goes down is useful but is a reactive condition. To be proactive you need evidence that a problem is brewing. Hyperic HQ has a great number of charts that can aid in this effort.

Click the Browse Resources link and then click the Chart View link below the Tools Menu. This will load all of the "Indicator" charts for the resources that were listed on the screen previously. A green background is quick indicator that the platform is healthy whereas red is a sign of trouble, and yellow means that it is partially available.



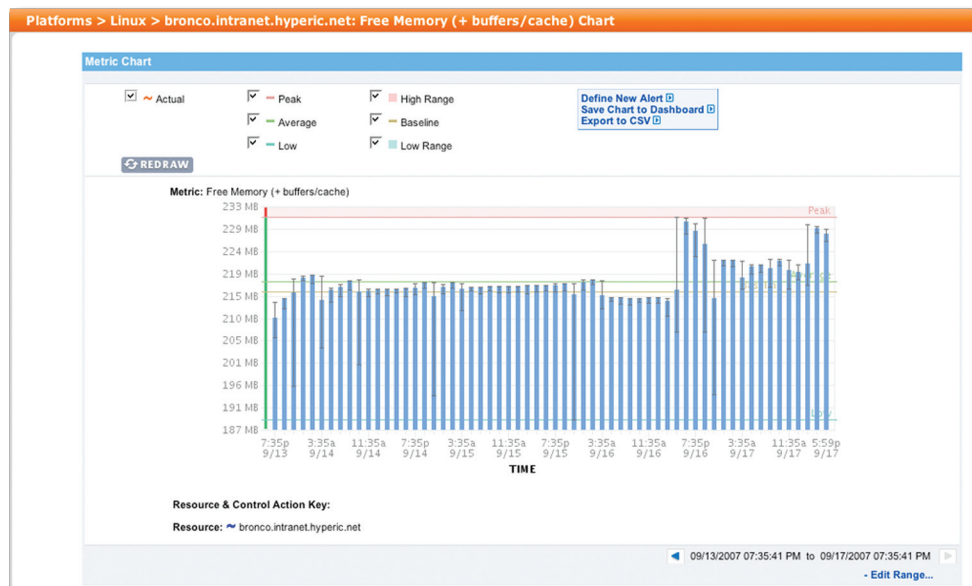
Figure - charts provide a quick view on the overall trending of multiple resources.



As you browse the many charts on the screen you check each platform and visually identify machines that are slowly running out of Free Memory or whose Load Average is much higher than the rest.

Figure - Detailed charts provide greater visibility into performance trends of specific resources.

Note: Baseline, high and low range trends are HQ Enterprise features.



When you see a particular chart that causes concern you can drill down to just that chart with a single click. Click a chart now and a larger version of that chart will load. The single chart page gives you a much more detailed chart for that single metric. You can change the time frame with the Edit Range tool or even export the data to CSV using the link above the chart.

Alerts

Alerting is a key component of any monitoring tool. In the Getting Started section, we have already set up one or more alerts for a platform going down. You may have noticed that there were a number of other conditions in Hyperic HQ that you can set alerts for on any platform, server, or service you are monitoring. Additionally, there are multiple ways to drill down on a resource to conveniently administer or view alerts, and to receive alerts. Additionally, Hyperic HQ alerts can be sequenced to form escalations. This section is designed to demonstrate a few of these additional combinations.

Let's say that you want to create an alert off of one of your most recently viewed resources. In the Operations Dashboard, you should now have a populated list of 'Recently Viewed Resources'. Alternatively, if you wish to use another platform not recently viewed, click the **Browse Resources** link at the top of the page. By default you will be presented with the list of platforms on your network. Select one of the platforms in that list and then click the Alert tab. From there, click the New button and we can start working on our alert.

To create your new alert, perform the following:

1. Give it a name
2. Choose "Availability" in the Metric menu
3. Change "Greater than" to "Less than" and put a one in the adjoining value field
4. Click OK at the bottom of the screen

You have now created an Alert for your platform that will fire immediately and repeatedly if the box goes down. The following screen lets you set an Escalation (covered in the next section) or lets you specify email addresses or users that should receive the alert.

Escalations

Hyperic HQ 3.1 allows you to send alerts on a timed basis so that different individuals are depending on the length of time a resource has been unavailable.

Note: This example assumes that you have set up multiple users already in the Hyperic HQ system. These users are referred to as either "Frontline Support" or "Backline Support". Where you see this, substitute individual users in for these roles. A formalized concept of roles which controls visibility, and alert assignment exists in Hyperic HQ Enterprise. For more information on that, see the section "The Fully Loaded Model".

To access the Escalation Scheme configuration click the **Administration** link at the top of the screen and then click Escalation Schemes Configuration on the following screen.

To create an Escalation that notifies front line support first and after 30 minutes notifies the IT manager, follow these steps:

1. Name your Escalation
2. Select the radio button "Allow user to pause escalation" and leave the menu set to "5 minutes".
3. Click Next Step
4. Click Create Action
5. In the "Select who to notify" menu choose "Notify HQ Users"
6. A dialog box will appear, choose users who are designated as "Frontline Support"
7. Change the last menu from "continue" to "wait for 20 minutes"
8. Click Save

The first step of your Escalation is now recorded. Get started on the second step using the Create Action button again. Repeat the above steps but this time notify users who are designated as "Backline Support" and leave the bottom menu set to "continue".

Your Escalation is now ready to go. Navigate to the **Browse Resources** page, choose the platform you used in the Alert section above, and click the Alerts tab. Select the Alert you created earlier and in the Escalation tab choose your Escalation Scheme.

This alert is now smarter than it was before. Rather than bother the entire IT Operations team when a your box goes down it will just notify your Front Line techs. The Back Line techs are only notified if the first group fails to rectify the problem in a 20 minute window.

Alert Center

One of the new features for HQ 3.1 is the Alert Center, where admins can see all of the current alerts in one easy-to-find place. You can always access the Alert Center via the top navigation, to the right of Browse Resources.

Once in the Alert Center, you can quickly scan current alerts and click on those that you wish to investigate. Alerts not yet acknowledged will be highlighted in red. Controls on the left allow you to filter the list based on alert type, priority, and time frame.

Figure - Alert Center
provides a central place
to manage all alerts in a
flat heirarchy.

Note: Sorting alerts
by groups is an HQ
Enterprise features.

The screenshot shows the Hyperic Alert Center interface. At the top, there's a navigation bar with links: Dashboard, Browse Resources, Alert Center (selected), Report Center, Administration, Recent Resources, and Help. Below this, the 'Alert Center' section is active, showing an 'Alerts' tab. On the left, there's an 'Alert Filter' panel with 'Alert type' set to 'Resource', 'Minimum priority' set to 'High', and 'In the last' set to 'All Time'. The main area displays a table of 'Resource Alerts'.

Date	Alert Definition	Resource	Fixed	Acked By	Priority
9/17/07 10:35 AM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/16/07 10:35 PM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/16/07 10:30 AM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/15/07 10:25 PM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/15/07 10:20 AM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/14/07 10:15 PM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/14/07 2:49 PM	Site Unavailable	SF Login Check	No		High
9/14/07 10:10 AM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/13/07 10:05 PM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/13/07 10:04 AM	High Disk Usage	web2.hyperic.net Linux File System /dev/sda2 mounted on / (local/ext3)	Yes		High
9/13/07 10:00 AM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/12/07 9:50 PM	High Disk Usage	vmlin-008 Linux File System /dev/mapper/VolGroup00-LogVol00 mounted on / (local/ext3)	No		High
9/12/07 9:24 PM	High Disk Usage	web2.hyperic.net Linux File System /dev/sda2 mounted on / (local/ext3)	No		High

RSS

Watching the **Dashboard** or your email inbox isn't the only way to know if an important Alert has fired. Hyperic HQ 3.0 includes an RSS feed so that while you're getting your daily tasks taken care of you can watch for alerts in any RSS reader.

To access the RSS feeds you can click any of the familiar orange XML icons in the **Favorite Resources**, **Recent Alerts**, or **Control Actions** portlets that appear on the **Dashboard**. In an RSS friendly browser such as Firefox this will cause a pop-up window to appear allowing you to save the RSS feed. Alternatively you can right click on the XML icon, choose Copy Link Location, and paste the link into the RSS reader of your choice.

Once you start receiving RSS feeds, you will note that you can Acknowledge or mark as Fixed the alerts you receive directly from the RSS notifications.

Control Actions

Hyperic HQ is primarily a monitoring tool but it also has the ability to send commands to servers that it monitors. For instance, if you notice that your web server is "Available" but that it is rejecting connections it may be time to restart it. Or, perhaps, you have noticed that an application based on JBoss has become sluggish -- you can use Hyperic HQ to run the Garbage Collection routine for JBoss.

Apache Restart Example

1. Click the Browse Resources link at the top of the page
2. Select the "Servers" link that is just to the right of "Platforms"

3. From the View menu choose "Apache 2.2"
4. Pick one of the servers listed that is "Available"
5. Click the Control tab
6. From the Control Action menu choose "Restart"

While Apache is processing that command proceed to the next example.

JBoss Garbage Collection

1. Click the Browse Resources link at the top of the page
2. Select the "Servers" link that is just to the right of "Platforms"
3. From the View menu choose "JBoss 4.0"
4. Pick one of the servers listed that is "Available"
5. Click the Control tab
6. From the Control Action menu choose "RunGarbageCollector"

To see the results of these two examples, along with a complete history of every control action for a resource, navigate to that specific resource's Control tab. Additionally, you will notice when you return the Dashboard and scroll down to the Control Actions portlet in the right hand column, it will now list the most recently executed commands and the number of times it was executed.



Fully Loaded Model

In this section, we provide an overview of Hyperic HQ 3.1 Enterprise. It's got the same look and feel and the same powerful engine under the hood, but it has some fancy features meant to support bigger teams managing larger scale deployments. These features are available in a paid-for annual subscription, but we assure you - if you are managing business-critical services at any significant scale - these features will save you time and money by streamlining management tasks, and automating behavior and analysis. HQ Enterprise will provide you a level of operational efficiency that is sure to get noticed! Who knows, buying that Fully Loaded Model may even secure you that dedicated parking space you've been after! Check it out, see if it's of any interest - and if it is, you can take it out on a test drive by downloading the trial directly from our website, or call your friendly neighborhood sales rep for a chauffeured tour.

Support for Roles

Roles in HQ provide a means to grant granular access to information on resources, and to delegate authority and responsibility for managing those resources.

Role Based Access Control

Role Based Access Control (RBAC) is used to assign permissions to users to view or manage resources based on their responsibilities in your organization. You can assign groups of resources to a role, and can control whether members can view, create, modify, delete, receive alerts, or control resources to which they are assigned.

This can be used in a variety of ways:

- Application developers can be given read-only access to Oracle Database metrics on database servers providing service to their development environments without the Database Administrator having to worry. This can reduce requests to the DBA while enabling better software development testing.
- Managers that want to see the "big picture" on their IT environment can be provided with access to view the status of resources that define business processes in their organization, and can tune their **Dashboard** settings to display key metrics and availability summary information.
- Provide users with key metrics on resources they use and care about without overwhelming them with unnecessary, and reduce support call burden. Is the mail server down, or is it just me? Are we having problems with connecting to the Internet? Let them answer these basic questions through self-service access to status via HQ while the IT staff works to resolve the problem.

Role Based Alerting

Role Based Alerting is used to set notification for an alert or in an escalation scheme to a collection of users defined in a role instead of having to add users one-by-one. This is a critical feature in large environments, where there tends to be greater segmentation of resource responsibility and accountability. It makes adding new users to alerts destined to specific group in the organization easier to manage, as you have to add the user only once...to the role that represents the group of users, instead of having to add the to every alert for which their role will have responsibility. Role-based alerting can be leveraged further when setting up Policy Control.

Enterprise Alerting

Enterprise Alerting extends alert definitions to account for scenarios typical of large scale, large team deployments. These extended alerts include multi-conditional, recovery and group based alerts.

Multi-conditional Alerts

Multi-conditional alerts allow an administrator to add additional depth to conditions which cause alerts to fire. A standard alert looks at a single condition, evaluates it based on real-time metrics, and then decides whether or not to fire. Multi-conditional alerts, as the name implies, simply keeps track of multiple conditions, and fires based on evaluation of them all. Conditions can be strung together using AND or OR Boolean logic.

An example of this in action would be setting an alert to fire if "free memory drops below 10 MB AND the 5 minute CPU load is greater than or equal to 2.5". Both the free memory and CPU load could be set as conditions for independent alerts. Coupled together, they may indicate a specific trouble condition an administrator needs to take action on immediately, and become a powerful tool for proactive systems management.

Global Alert Templates

Through the use of policy-based alerting, HQ makes it possible for administrators to set default alerting policies for all resources of a particular type. This means that, instead of having to set discrete alerts for each individual resource, administrators can set a global alert template once for all Linux Platforms, or all Apache Servers, or all Oracle tablespaces, etc. across the environment. This makes control of alerts in larger scale enterprises more manageable, and the team more flexible to change alert policies globally when the business demands.

Examples of this in action would include setting an alert for 90% used disk space across all file system mount resource types, setting an alert to fire if CPU load exceeds 2.0 across all CPU resource types, or setting an alert to fire if the number of connections to the Apache resource type exceeds 150. Essentially, any alert that can be set on an individual resource can be set for all resources of that type using policy-based alerting.

Like policy-based monitoring templates in HQ, policy-based alerting is managed through the "Monitoring Defaults" section of HQ "Administration".

Recovery Alerts

HQ Enterprise gives administrators the ability to define a special type of alert for when a system returns to a normal level of operations after exceeding a threshold. This concept actually involves a pair of alert definitions - one that defines an error condition and a special linked alert that defines the recovery condition. When these two alert definitions are set up, the administrator will see one alert notification when the error condition occurs and then another notification when the recovery condition occurs.

An example of a good use for an alert/recovery alert combination would be an administrator who wants to know when a particular Linux platform is under a heavy CPU load. The initial alert to indicate a problem could be "5 Minute Load Avg > 2.0". The recovery alert could be "5 Minute Load Avg < 0.5".

In the event that the Linux 5 minute system load average goes over 2.0, the administrator would get an email alert notification about this event. That alert would automatically become disabled in HQ so no more alert notifications will be sent about that alert, no matter how long the system load average remains over 2.0. However, when the system load average drops to < 0.5, the recovery alert will notify the administrator and HQ will automatically re-enabled the first alert.

This feature automates the process of reconciling alert conditions as fixed. With this support, operator efforts are better focused on investigating alert conditions that are still outstanding.

Group Based Alerting

Collections of resources can be combined into groups within HQ. Now, with Group Based Alerting, you can alert based on the status of a group instead of just the individual resources contained in the group. For Compatible Groups (groups of resources all of the same type), you can set alerts based on any metric shared across the group. For Mixed Groups (groups of resources that are of different types), you can set alerts based on Availability.

This is an extremely powerful concept for administrators responsible for maintaining service-level agreements based on users' perception of application availability. Let's say an administrator has a cluster of three application servers defined in a group. If one of those servers went down, it might be important, but would not necessarily be critical enough to wake an administrator at 3 AM in the morning to resolve the issue because redundancy is still preserved with two servers in the group still up. If, on the other hand, two of the three resources went down, an additional resource failure would represent end user downtime, and would warrant waking an administrator up at an obscene hour.

Group alerts can be set to fire based on physical counts (e.g. when 2 servers meet a condition), or on a percentage of the total number of resources in the group (e.g. 66% of the group). Alerting for groups uses escalation schemes for notification.

Follow-the-sun Alerting

Large operations teams typically service their environments in shifts. Groups of individuals rotate responsibility for responding to alerts based on time of day. With Hyperic HQ Enterprise, these groups of individuals can be organized into roles, such as "Day Shift" and "Night Shift" or "Americas Support" and "EMEA Support". These roles can further be configured for availability based on the time of day on the HQ Server. Then, when an alert fires, the HQ Server will determine which role is available and notify the appropriate role.

This functionality is critical to balancing workload across shifts or geographies, and focusing those teams on investigating alert conditions assigned to their work queue.

Scheduled Control

Control actions are a powerful tool within HQ. For the Servers and Services that support them, they can provide HQ users with a means of performing the actions necessary for server or service recovery in a variety of situations. The Enterprise version of HQ adds the ability to schedule control actions to occur at some date in the future, or as a regularly occurring process for maintaining an environment.

Administrators do not always like to admit it, but most are familiar with that one troublesome piece of their environment that requires a regular restart during off-peak times to eliminate the risk of that service going down when users might feel it most.

Scheduled control in HQ allows administrators to not only automate this process, but to log it for on-going analysis. A positive side-effect of having the ability to schedule control in HQ Enterprise is that alerts can be configured to schedule an immediate control action based on an alerting condition. This allows for automated incident response, which maintains an auditable record of actions performed, and wakes up your administrators in the middle of the night only if the recovery control action fails.

Report Center

Pre-built reports are an important communication vehicle for organizations to monitor progress, service levels and the overall health of your IT environment. The Report Center, accessible via the top navigation, provides many of these reports out of the box. From the Report Center, you will find a list of pre-defined reports, summarizing metric reporting and availability. More reports will be added as we go forward. You may change the scope of any report by specifying a resource group or the beginning and end dates. The report format defaults to PDF, but you may also choose HTML, CSV or Excel formats.

Smart Trend Auto-Baseline Feature

As HQ collects metrics on all of the resources it is managing, it keeps track of a baseline for each metric via Hyperic's Smart Trend technology. For example, let's say you are collecting the metric for "5 Minute Load Average" on a Solaris platform. After HQ has collected a sufficient set of data for this metric, it will automatically perform a baseline calculation. Let's say the baseline for this metric is "0.3". If the 5 minute load average on the machine then spikes higher, say 1.2 for a sustained period of time, the metric would then be considered "Out Of Bounds" (OOB). Resources with OOB metrics are brought to the attention of the HQ administrator via the **Problem Resources** dashboard portlet. The baseline value will continue to be automatically adjusted as the data set changes over time.

There is no user interaction required to turn on this feature, it is automatic. The size of the metric set to be used when calculating the baseline for a metric can be configured by the HQ administrator in the Server Administration section. This feature is only available in the Enterprise version of HQ.

External Authentication

Larger organizations typically have a centralized authentication system, through which they manage all user accounts for enterprise applications. This is far easier to maintain than discrete authentication systems for each application. HQ Enterprise provides support for user authentication against LDAP, Active Directory or Kerberos third-party systems. This allows HQ to be easily integrated into existing user management processes.

High Availability

When "All Systems Go" is your requirement for application availability, the one thing that can not be tolerated is downtime due to a single point of failure in your monitoring architecture.

Failover Clustering

For this reason, true fail-over clustering is built-in to the HQ Enterprise architecture. Once configured, if an HQ node goes down, other nodes in the cluster recognize the failure, and immediately assume the load of the failed node in a user-transparent way.

Linear Scalability

Another important facet of High Availability is scalability and performance of your HQ implementation. HQ Enterprise allows load balancers in front of a cluster of HQ Servers. Agents report metrics to any of the HQ Servers based on a round-robin assignment. The HQ Server will then process the metrics, issuing alerts or other responses and populating the database. This provides the linear scalability necessary to manage thousands of resources centrally.



Fahrvergnügen

Driving enthusiasts love other driving enthusiasts. The same is true of Hypericans. Hyperic believes the biggest value of open source, and Hyperic HQ, is the proud community that supports it. We recommend trying out the community support, including our forums, at <http://www.hyperic.com/community>.

We hope you enjoyed your test ride, and will be joining our community. We're eager to hear from you soon!