

A White Paper

## **Best Practices for Implementing Server Monitoring**

***HEROIX***

Heroix  
165 Bay State Drive, Braintree, MA 02184 USA  
[www.heroix.com](http://www.heroix.com), [info@heroix.com](mailto:info@heroix.com)

# Best Practices for Implementing Server Monitoring:

When it comes to server (physical or virtual) monitoring there are a finite number of performance and availability metrics which all can be effectively categorized as “data”. When considering best practices for server monitoring it is best to segregate the handling of the data into four distinct categories:

- I. Capture
- II. Store
- III. Evaluate/notify/correct
- IV. Interact

In order to have an effective and efficient server monitoring strategy special attention needs to be applied to each of the 4 categories. The goal is to have a cohesive strategy in place so that IT can readily identify and avoid present and future server performance and availability issues.

## Capture

**There are all kinds of data available for server monitoring. The challenge isn't determining what data to gather but rather the best practices surrounding how to best capture the data. A server monitoring implementation not only need to be fast, lightweight, and reliable but it also must be resilient to network disruptions.**

- ✓ **Discovery:** Ideally, the best data capture process is one that is automated and that minimizes the interaction required by IT staff to account for changes in the server infrastructure, this means employing a methodology that automatically discovers and understands key server components coming and going. For example, it could be discovery related to services on a Windows platform, file systems on Unix, or VMs in a virtual infrastructure. Automated discovery not only saves in terms of IT administrative time, but it also eliminates the possibility that a critical component within a volatile server infrastructure will miss being monitored.
- ✓ **Deployment:** Go “lightweight” whenever possible. A best practice is a data capture strategy that minimizes the footprint/overhead on the servers being monitored. Lightweight can readily be achieved by using protocols that currently exist on the servers to agentlessly capture data. There are a variety of agentless protocols available including WMI, SSH, JDBC, SNMP, JMX, etc. An agentless strategy also provides for a fast return on effort, as the time span between deployment and actual monitoring is negligible

Heroix

165 Bay State Drive, Braintree, MA 02184 USA  
www.heroix.com, info@heroix.com

- ✓ **Scaling:** As virtualization and cloud computing increase in popularity - the sheer numbers and types of servers is mushrooming. A distributed architecture where the capture of critical server metrics can be readily spread across commodity based computers (physical or virtual) allows for a monitoring approach that can't be constrained by size or scope.
  - ✓ **Resilience:** Servers and the accompanying IT infrastructures are increasingly dispersed; the challenge is ensuring data capture even during network disruptions. An underlying principle is to ensure the capture of all data even when the network is compromised. Ideally a store and forward strategy for remote locations will ensure that critical data is delivered when connectivity is lost and later re-established.
- 

## Store

**The long term archival of captured data into a database is essential to providing historical context of server related issues. Best practice is to retain all relevant data that can come in a variety of forms including:**

- ✓ **Performance** – Data related to CPU, Memory, Disk, Network and more is critical to understanding and ensuring that the server infrastructure is delivering according to expectations. Proper analysis of the performance data will identify problem areas before there is a resource deficiency or hardware problem.
- ✓ **Availability / Response time** – We can break this down to 1) availability/response time between end-user and server and 2) availability/response time between servers. Retention of this data is critical to mapping server behavior as it relates to business continuity and ultimately user productivity.
- ✓ **Service Level Compliance** – Correlates server performance data and availability/response time data to business services. Many IT departments have contractual obligations for service levels, Service Level Agreement (SLA) data documents the level of compliance. Long term SLA data can help identify patterns of non-compliance based on time of day, day of the week, etc. and allow for more effective problem resolution
- ✓ **Event Data** – Event data can come from a variety of sources including server logs and application logs. Events provide context as to what performance and availability issues occurred, when, as well as the level of persistence.

---

## Evaluate/Notify/Correct

**Essential to any server monitoring strategy is accurately assessing the captured data, proactively alerting staff to problems, and where possible initiating corrective actions.**

Best Practices include:

- ✓ **Avoid false positives:** Once the key metrics have been captured and stored, the next step is to evaluate the data and based on the results determine what problems warrant attention. Performance metrics can be tricky depending on whether the metric's calculated value is based on a moment *in* time (i.e. Disk Space used - Windows) or is based on a value *over* time (i.e. CPU busy percent – Windows). Evaluating disk space is simple, it is one value, you'll compare to a defined threshold and make a problem determination. However, rate based metrics like CPU percent are *very different* - the value is calculated based on CPU Time used / Elapsed time. This concept has applicability across platform (i.e. CPU Ready% in VMware). In order to avoid a false positive, the best practice is to evaluate the metric across multiple collections for problem determination. For example, if we calculate CPU busy percent over 5 minutes and collect 3 values over 15 minutes of 90%, 30%, and 60% then the best practice is to NOT TRIGGER on the 90%, but evaluate our CPU usage across multiple collections, i.e.  $(90\%+30\%+60\%)/3=60\%$ . Evaluation across multiple time periods provides for a far more realistic picture.

If too many false positives enter into the equation, then we get into a bit of “Boy who cried wolf”, where the problem volume is so high that it is difficult to separate serious problems from the fodder.

- ✓ **Resilient Problem Notification:** When a problem occurs and early responders' attention is warranted then mechanisms like email and SMS must be resilient. The notification mechanism(s) should be independent of the server infrastructure being monitored. For example, what if an email notification needs to be sent but the internal email server is unreachable? A notification best practice is to bypass the internal server infrastructure all together and pursue technology with a built-in email or SMS capability. Bypassing the corporate email server and sending email outside the corporate domain will ensure delivery even under the most serious of situations.
- ✓ **Automate Problem Resolution:** Avoid evaluating in a vacuum. If at all possible, maintain a context of how persistent a problem is. Where possible, execute scripts or commands as a first line of defense to fix or mitigate a problem and then escalate with additional actions and/or notifications if a problem continues.

Heroix

165 Bay State Drive, Braintree, MA 02184 USA

[www.heroix.com](http://www.heroix.com), [info@heroix.com](mailto:info@heroix.com)

---

## Interact

**A web interface to control and visualize the data is especially important. A best practice for any server monitoring strategy is the ability to display what problems are unfolding real-time as well report on what historically has occurred.**

- ✓ **Centralized Management:** The underlying monitoring strategy should operate on heterogenous environments and abstract out platform differences. All platforms need to be managed with a common look and feel so as to reduce complexity and decrease administrative overhead.
  
- ✓ **Visualization:** The presentation of critical data is an absolute must. The value of the data can only be truly realized when it can be readily consumed.
  - Graphical Display – The ability to quickly summarize and see critical problem graphically via customizable real-time dashboards is especially important. Color coding along with a hierarchical display with drill-down capability enables IT to triage problems and act quickly to the most critical of issues.
  - Event Display – When problems are documented as events, they provide critical information as to the nature and persistence of issues. A real-time event display that groups events based on any IT defined criteria allow staff to quickly diagnose how pervasive a problem is and how it affects the supporting IT infrastructure and applications.
  
- ✓ **Historical Reports:** Reporting on long term server performance, problem history, and service level compliance is essential to identifying capacity and availability issues. As the saying goes... *“Those who don’t learn from history are doomed to repeat it”*. A good practice is to be proactive about reporting and to automate the report generation process. Critical data that is disseminated regularly via email or a web portal helps keep everyone from IT staff, to management, to end users informed and also helps show the value of IT!
  - **Performance data** – Report on server performance
  - **Availability / Response time** – Report availability/response time for servers and applications.
  - **Service Level Compliance** – Report on both server performance and availability/response as a service level for a business services.
  - **Event Data** – Report on event history to provide context around performance and availability issues.

Heroix

165 Bay State Drive, Braintree, MA 02184 USA

[www.heroix.com](http://www.heroix.com), [info@heroix.com](mailto:info@heroix.com)

---

## Conclusion

**When it comes to server monitoring it is all about the data. It is about the best way to capture the data and once that data is captured how to best utilize it.**

**The monitoring process itself needn't be heavy or burdensome. Embracing a strategy that is lightweight, efficient, resilient, and automated provides for a fast return on effort.**

**Equally important to the actual collection of data are the algorithms used to evaluate the data. An approach that is efficient about identifying problem areas and isn't chatty allows for more effective monitoring and alerting of server infrastructures.**

**If you can't see the problem you can't fix the problem, visualization of the data is essential to any server monitoring strategy. Real-time displays as well as historical reports enhance the value of the data by making easy to consume and understand.**

**Incorporating best practices that take a streamlined approach and that focus on the efficient collection and utilization of server data will assuredly result in a successful long-term server monitoring strategy.**

---

## About Heroix

Heroix has a 30+-year history of proven monitoring solutions, with products running on tens of thousands of critical servers. It offers fast, easy, affordable application and networking monitoring solution for physical and virtual environments. [Download Longitude Now](#) and you'll be monitoring and planning in just 10 minutes.

Heroix believes that the information in this document is accurate as of its publication date; such information is subject to change without notice. Heroix is not responsible for any inadvertent errors.

Heroix, Heroix Longitude and their corresponding logos are registered trademarks of Heroix. All other company and product names mentioned are used only for identification purposes and may be trademarks or registered trademarks of their respective companies.

Copyright © 2017 Heroix. All rights reserved.

Heroix  
165 Bay State Drive, Braintree, MA 02184 USA  
[www.heroix.com](http://www.heroix.com), [info@heroix.com](mailto:info@heroix.com)