Design Principles for IT Monitoring Systems

INTRODUCTION

Monitoring tools are critical to quickly identify and address problems that affect an IT organization's service to its users. These tools scan for problems with network, systems, and application resources that a company relies on for its business.

However, IT managers face significant challenges deploying and managing these systems to their maximum benefit. Key issues include: 1) high costs for commercial software licenses and maintenance fees; 2) waste from unused "shelf ware" and/or multiple tools with redundant functionality; 3) dependence on expensive consultants for installation and support; 4) lengthy and expensive system deployment, and 5) high resource cost associated with system management, administration, and support.

Moreover, monitoring systems frequently are not configured to produce critically needed input to the IT business processes used for management of service levels, availability, capacity, incidents, problems, and operations.

Well designed monitoring systems, those integrated with standard IT management practices, are key to achieving world class IT service effectiveness and are a major driver of corporate productivity and competitiveness. The design principles described below can be incorporated into a system that is simple and inexpensive to deploy, manage, and operate.

Maslow revisited: the hierarchy of monitoring functions

Most operational challenges associated with monitoring are associated with deploying more monitors and features than are necessary to achieve effective systems management. Failure to methodically design the monitoring process creates systems that are expensive and unwieldy to deploy, manage and maintain.

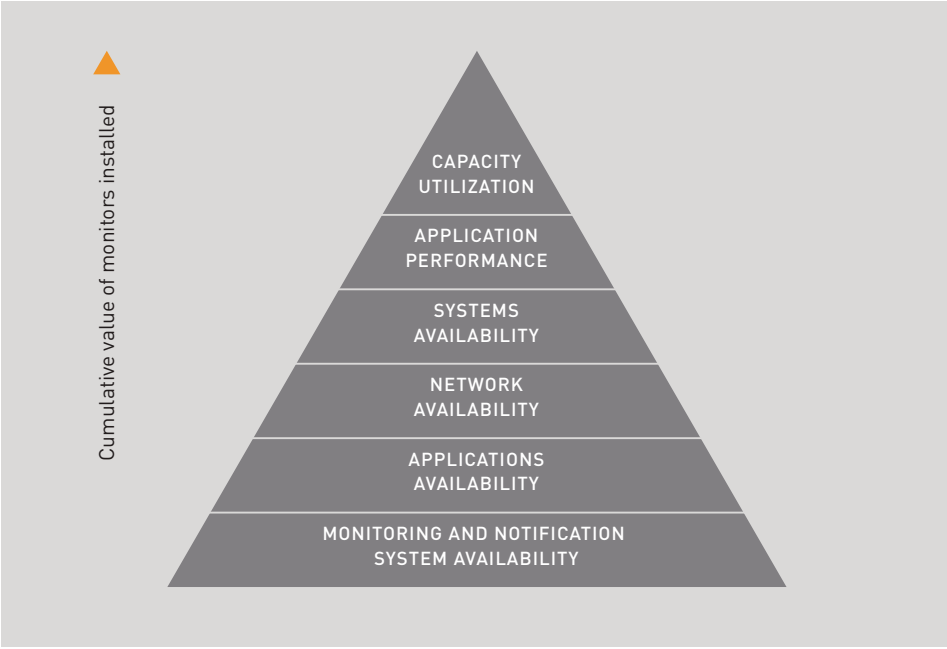

Figure 1: Hierarchy of monitoring functions

Table 1: Issues impacting monitoring system design and strategy

Key infrastructure issue	Implication for Monitoring System
Infrastructure size and complexity	For large infrastructures, simpler solutions will keep the cost of changes at predictable levels. Consider de-centralizing monitoring hosts and system management responsibility.
Cost of outages	The higher the cost of outages, the more should be invested in additional monitors (capable of detecting incipient failures) to avoid them.
Redundancy and diversity of infrastructure components	More monitors are justified to ensure that the "b channel" components are available before failover.
Expected rate of network and system change	The higher the rate, the simpler the monitoring configuration should be. Ensure that additions, changes and deletions can be made as easily.
Size and experience of staff tasked with analyzing monitoring output data	Ensure that all relevant monitoring data is included in reports, but not necessarily linked to alarms. Understand how reports will be used before developing them.
Stability of the applications and content	Stable systems require fewer monitors and, paradoxically, enable more sophisticated monitoring systems.
Support staff availability	24/7 NOC/Help Desk support permits dramatically simpler monitoring vs. pager-based, first level response support.
Heterogeneity of equipment and software	More heterogeneous environments require more complex monitoring solutions to gain the same level of business benefit.

To achieve an effective design, it is advantageous to group monitors by type and purpose into a "hierarchy of functions" as shown in diagram (Figure 1). Monitor groups providing the most benefit for the least cost should be the highest priority, followed by less valuable monitors until the marginal benefit of adding another functional group approaches zero.

Focusing on producing maximum infrastructure coverage with the fewest possible monitors reduces system complexity, deployment time and cost. As much as 80% of the business benefit from system monitoring can be obtained from as little as 20% of the total monitors in a given monitoring system.

The goal is to quickly deliver a comprehensive foundation to which additional monitoring groups can be added if and when justified. Key environmental issues and their implications for monitoring systems are detailed in Table 1 above.

CONFIGURATION AND DEPLOYMENT

Monitoring system deployment should be conservative and simple. The goal is to use as few monitors as possible to minimize complexity and cost, while delivering maximum IT management and business value to the organization. This suggests an incremental approach of deploying the most important monitors first, and configuring and deploying them to take advantage of their full reporting, analysis, and dashboard and other capabilities.

Installing incrementally allows the discoveries, experience and early wins associated with the first deployment to drive the design choices made in later deployments. Results from an initial deployment often surface unanticipated issues requiring attention in later deployment while satisfying other issues in unexpected ways. For example, availability monitors deployed in an initial project will provide a baseline of metrics that help determine which additional monitors would provide the most business benefit. Installing the highest impact functions first will also provide information that accelerates integration of the system's output data into the organization's business processes.

The resulting system should make use of all data derived from each of the monitors. Monitors left unconfigured, or those whose output is ignored, typically generate expensive and distracting false alarms. Configuration options should include group assignments, filter assignments, threshold values, retry interval and limits, dependencies, scheduled maintenance, performance data, schedules, notification media choices, and the logical structure and schedule of escalation.

The relative utility of availability statistics illustrates the importance of a complete configuration. If the dependencies between network component availability and the collection of monitoring data are not correctly configured and maintained, the resulting availability statistics will be meaningless. Component outages will be recorded not because the components are inoperative, but because they are not available to be monitored.

OPERATIONAL CONSIDERATIONS

As with most computer systems, the successful design of monitoring systems is highly dependent on the recognition that human operators can determine the success or failure of a project. Here are a few important factors to consider:

Alarm Fatigue
Getting paged in the middle of the night to respond to false alarms is wearing and reduces responsiveness to real problems. There are a number of ways to minimize the critical alarms requiring employees to be paged during and outside of business hours. These include: 1) automatic alarm verification; 2) automatic restarts (event handlers); 3) priority-based notifications; 4) escalation intervals adjusted for time of day; 5) "follow-the-sun" techniques, and 6) disabling of alarm channels during planned maintenance. Alarms should be enabled selectively: not all monitoring channels require alarms, and many should result in email-only notifications.

Calibration Approach
The set points for many alarms are affected by network propagation rates and typical variations of these rates. Set points should be determined by setting the alarms high and calibrating down rather than the other way around. This will allow the system to detect hard failures while avoiding nuisance alarms. This is particularly important during the initial days of installation.

Web Access
The monitoring system must permit responding personnel to have remote web access to minimize inconvenience associated with after hours support.

NOC Outsourcing
A well-designed monitoring system which generates a lot of real alarm traffic can exhaust a busy staff during periods of rapid change or significant infrastructure expansion. Consider keeping staff fresher during these times by outsourcing non-business hour level one alarm response to a managed service provider (MSP). Alarm response can be brought back in-house when system stability and associated costs allow.

Change Management
Manage the change in expectations for IT employee performance during the design process with good project management techniques.

CONCLUSION

Information technology organizations, working with experienced architects to apply sound principles to the design and deployment of monitoring systems, achieve significantly greater business benefits from monitoring, including higher application availability, higher employee productivity, and lower IT capital and operating costs. Whether you are using expensive monitoring software or open source software tools, the system's value will be determined by the quality of the design and implementation plan, and the degree to which your organization's business processes use the output – not the expense or features of the software purchased.

ABOUT GROUNDWORK

GroundWork Open Source Solutions, Inc. provides open source-based IT infrastructure management solutions such as network and systems monitoring, service desk management and IT dashboards. GroundWork's solutions enable IT management to leverage the flexibility and low cost of open source tools to achieve enterprise-level availability, performance and operational efficiency for a fraction of the cost of commercial software.

Contact us
510.899.7700
www.itgroundwork.com
info@itgroundwork.com

GroundWork
Open Source Solutions, Inc.
2200 Powell Street, Suite 350
Emeryville, CA 94608