

Reducing Costs With MATtool

Introduction

The Monitoring and Administration Tool (MATtool) suite is a collection of tools for managing large groups of UNIX machines with dramatically less effort and greater ease of use than many existing technologies. It simplifies daily administrative functions and permits management of one or hundreds of servers from an intuitive interface. This document describes how MATtool reduces the administrative effort of managing UNIX servers.

Intended Audience

MATtool is a powerful tool for improving the productivity of UNIX system administrators. This document is intended for:

- IT Managers. It illustrates how MATtool saves time and allows mundane tasks to be delegated to less senior personnel.
- System Administrators. It shows administrators how to operate on massive groups of servers quickly and effectively.

Cluster Management

In many businesses managing the IT infrastructure is a constant ongoing task. In today's markets labor costs, turnover and budget cuts place more pressure on IT staff to get more out of existing systems. Growing numbers of machines and operating systems to support, place increased burdens on existing IT staff. MATtool provides a unified tool for the management of both heterogeneous and homogeneous compute environments.

The following sections will examine:

- The monitoring functions that alert administrators of problem hosts
- Resource Management and the simplification and delegation provided by MATtool
- Logical and role-centric host grouping, and the administrative savings it provides
- Scheduling tasks
- Moving data, patches, and binaries between hosts

Resource Monitoring

MATtool comes preconfigured with a series of resource monitors. During installation MATtool automatically discovers the resources on a host and builds a monitoring configuration. Some of the current monitored resources include:

- Disk space on local drives
- Swap space
- Run queue

- Concurrent logins
- CPU usage
- Network interface activity
- HTTP service
- SMTP (Email) service
- FTP service

Alerts are also provided. When a service is unavailable or a resource is low, warning and error alerts can be sent. The built-in alerts include:

- Email warning and error messages
- SNMP v1 traps
- Syslog messages

Administrators can also build in their own corrective actions to automatically correct known issues.

Administrators can quickly identify problem hosts with a glance at the console. Hosts with service problems or resource shortages are displayed differently than hosts that are ok.

Historic reports are also provided by the monitoring functions. The data gathered when monitoring a service or resource is logged. Historic graphs highlight peak usage times, aiding administrators in tracking faults.

Resource Administration

MATtool insulates administrators by provides configuration tools for the common UNIX services, hiding OS differences. Costs can be reduced because mundane tasks can be delegated to less experienced people through user, host, and task access controls.

Some of the configuration tools provided are:

- NIS (Yellow Pages)
- DNS (bind 4.5)
- NFS servers and clients
- Email lists
- Passwords and groups
- Cron
- Syslog
- Inetd
- And many others

Managing Hundreds of Hosts

Compute centers are increasingly using low profile rack-mount servers. The higher densities mean more systems to manage. With MATtool cluster size is irrelevant. Ten, or ten thousand hosts can be managed with MATtool. A host-centric architecture lets MATtool scale to massive levels. Each host is an autonomous object. This section discusses how MATtool manages hundreds of hosts.

MATtool provides an arbitrary host-grouping scheme. Hosts can be grouped by function, location, type, or any user defined grouping. Hosts can appear in multiple groups. A machine in the web server group can also appear in the Boston machines group. Administrators can create views for various roles and responsibilities. Viewing only those hosts an administrator has responsibility for automatically filters the administrative noise.

MATtool leverages the host groups to provide convenient group status display. The status of the hosts in the group is shown by the group status. The host group also allows operations to be run against all hosts in the group. MATtool uses technology that distributes commands to MAT hosts in parallel. By selecting hosts and groups in the console an administrator can operate on 100's of hosts as if it were a single host.

A visual representation of the hosts and groups allows the administrator to quickly identify problem hosts. Hosts can take four states:

- Okay - The host is running
- Warning - The host is running, but there is a resource or service problem
- Error - The host is running, but there is a serious resource or service problem
- Dead - The host is not responding

The host state is propagated to the groups the host is a member of. An administrator can see when there is a problem with any of the hosts in his view. Problem hosts are quickly identified allowing administrators to quickly react to failures, minimizing downtime.

Scheduling Tasks

MATtool comes with it's own scheduler (MATd). The scheduler is user to run periodic processes. These can be monitoring jobs, backup or replication jobs, as well as user defined tasks. It has several features, which set it apart from other schedulers like cron such as:

- Time based or interval based run times. For example run at 12:09PM every Friday, or run every 17 minutes
- Run as any user. Invoke application maintenance processes as a user other than root

- Store the output of every task in its own file and never worry about it growing too large. Logs can never grow more than a user-defined length.
- Specify a maximum run time.

The MATd scheduler provides worry free task execution. It handles run-away logs, improper privileges, and process failure automatically. A simple UI eases administrators through the available features.

Moving Datasets and Binaries

MATtool has a data replication tool. It can be used to move any type of file from a replication server to other MAT hosts. New executables, patches, or datasets can be pushed from the replication server to the other hosts in the cluster. Data is replicated in parallel to eight servers at a time. Hosts with slow network links will not slow replication to others.

In compute clusters, distributing common datasets to compute servers reduces the load on NFS servers. Large compute clusters can cripple NFS servers because of the large datasets that need to be distributed. MATtool allows the creation of replication tasks, which will consistently replicate the data from one host to many.

In web farms updates should be staged, so that links and content can be verified before the data is replicated to the web servers. Once the content is verified, MATtool can handle the task of updating the web servers.

The MATtool replication tool provides an integrated, easy to use tool for moving large quantities of data between hosts. A simple user interface with online help makes creating and scheduling new administrative jobs easy.

Conclusion

MATtool saves time. Running commands on hundreds of hosts simultaneously saves time. Simple user interfaces and solution-oriented tools simplify tasks for less experienced personnel.

MATtool has been evolving over the past five years. User feedback helps MATtool add new features to address real business issues. Reducing the administrative effort has been and will continue to be forefront in the design of MATtool. Even the upgrade of the tool is handled with ease. An upgrade of over 200 hosts only takes a half an hour, and is done directly from the console.