

# Advanced Health Monitoring of Computer Clusters

Presenter: Caleb Morse

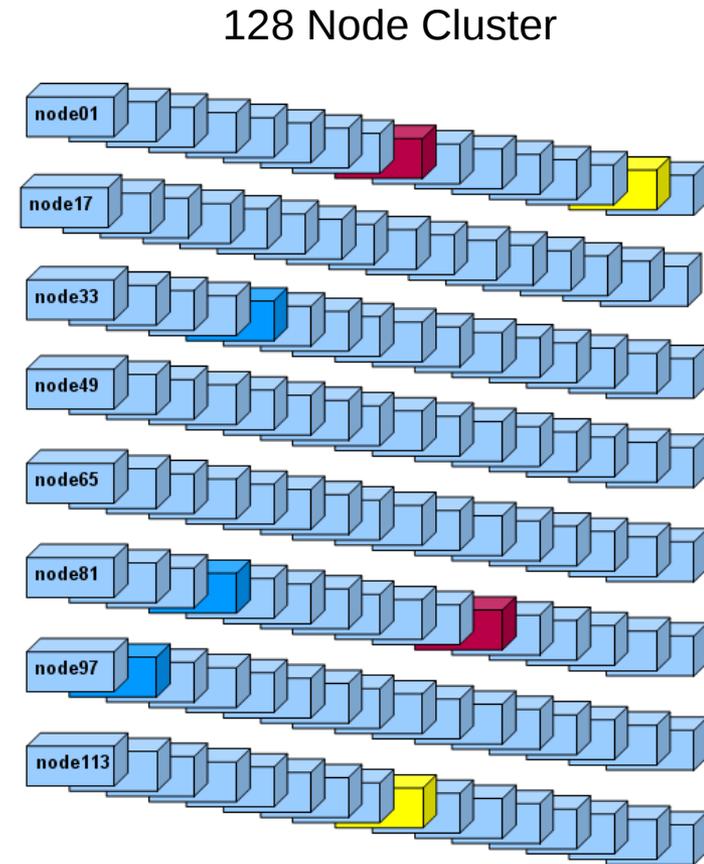
Team: Dylan Merrigan, Sherry Salas

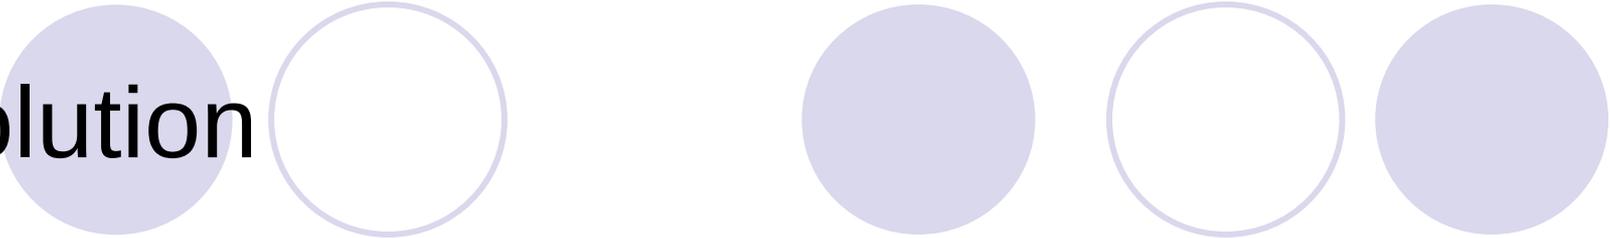
Mentors: Susan Coulter, Andree Jacobson, Kevin Tegtmeier  
LANL ISTI/IAS Cluster/Network Management Summer Institute

August 2, 2007

# Problem

- As a cluster grows in size it becomes much more difficult to monitor
- Mean time to failure increases dramatically with a large cluster





# Solution

- Identify system faults
- Recognize performance degradation
- Prompt notifications
- Low overhead on compute nodes
- Affordable



# Nagios

- Service checks
- Distributed monitoring and failover
- Custom plug-ins
- Multiple alert states
- Intelligently schedules service checks

# Nagios Gui

**Nagios**  
 Last Updated: Wed Jul 25 11:11:12 MDT 2007  
 Updated every 90 seconds  
 Nagios® - [www.nagios.org](http://www.nagios.org)  
 Logged in as nagiosadmin

**General**  
 Home  
 Documentation

**Monitoring**  
 Tactical Overview  
 Service Detail  
 Host Detail  
 Hostgroup Overview  
 Hostgroup Summary  
 Hostgroup Grid  
 Servicegroup Overview  
 Servicegroup Summary  
 Servicegroup Grid  
 Status Map  
 3-D Status Map  
 Service Problems  
 Host Problems  
 Network Outages  
 Show Host:

**Reporting**  
 Trends  
 Availability  
 Alert Histogram  
 Alert History  
 Alert Summary  
 Notifications  
 Event Log

**Configuration**  
 View Config

**Current Network Status**  
 Last Updated: Wed Jul 25 11:11:12 MDT 2007  
 Updated every 90 seconds  
 Nagios® - [www.nagios.org](http://www.nagios.org)  
 Logged in as nagiosadmin  
[View History For all hosts](#)  
[View Notifications For All Hosts](#)  
[View Host Status Detail For All Hosts](#)

**Host Status Totals**

Up	Down	Unreachable	Pending
3	1	0	0
<a href="#">All Problems</a>		<a href="#">All Types</a>	
1		10	

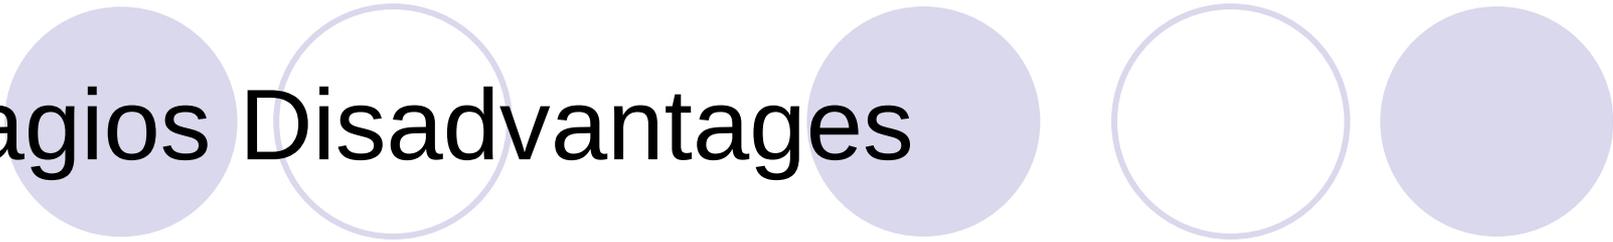
**Service Status Totals**

OK	Warning	Unknown	Critical	Pending
75	0	6	1	0
<a href="#">All Problems</a>		<a href="#">All Types</a>		
7		82		

**Service Status Details For All Hosts**

Host	Service	Status	Last Check	Duration	Attempt	Status Information
head	Current Load	OK	2007-07-25 11:10:20	0d 14h 7m 31s	1/4	OK - load average: 0.08, 0.18, 0.08
	Current Users	OK	2007-07-25 11:10:20	1d 20h 50m 24s	1/4	USERS OK - 8 users currently logged in
	Disk Space	OK	2007-07-25 11:10:20	1d 20h 50m 10s	1/4	DISK OK - free space: / 179169 MB (82% inode=99%); /var/run 1005 MB (99% inode=99%); /var/lock 1005 MB (100% inode=99%); /proc/bus/usb 1005 MB (99% inode=98%); /dev 1005 MB (99% inode=98%); /dev/shm 1005 MB (100% inode=99%); /lib/modules/2.6.20-15-generic/volatile 967 MB (96% inode=99%); /net/10.201.3.29/handouts 30322 MB (82% i
	Tests bind service	OK	2007-07-25 11:10:20	6d 20h 57m 19s	1/4	DNS OK - 0.009 seconds response time (head.azule.com. 3600 IN A 192.168.1.1)
	Total Processes	OK	2007-07-25 11:10:20	6d 2h 18m 12s	1/4	PROCS OK: 222 processes
	httpd	OK	2007-07-25 11:10:22	6d 20h 40m 14s	1/4	HTTP OK HTTP/1.1 200 OK - 2398 bytes in 0.005 seconds
	httest	OK	2007-07-25 11:10:20	5d 8h 50m 0s	1/4	Infiniband operating at 4 X.
	mysqld	OK	2007-07-25 11:10:20	6d 20h 56m 50s	1/4	Uptime: 161452 Threads: 11 Questions: 231563 Slow queries: 0 Opens: 2424 Flush tables: 1 Open tables: 63 Queries per second avg: 1.434
	nfs	OK	2007-07-25 11:10:20	0d 16h 29m 31s	1/4	NTP OK: Offset -8.016471863 secs
	sensors	OK	2007-07-25 11:10:20	1d 20h 50m 30s	1/4	CPU1 -99 F CPU2 +109 F
h1	Current Load	UNKNOWN	2007-07-25 11:10:20	0d 0h 3m 52s	1/4	Remote command execution failed: ssh: connect to host 192.168.1.101 port 22: No route to host
	Current Users	UNKNOWN	2007-07-25 11:10:20	0d 0h 3m 52s	1/4	Remote command execution failed: ssh: connect to host 192.168.1.101 port 22: No route to host
	DNS	UNKNOWN	2007-07-25 11:10:20	0d 0h 3m 52s	1/4	Remote command execution failed: ssh: connect to host 192.168.1.101 port 22: No route to host
	Disk Space	UNKNOWN	2007-07-25 11:10:20	0d 0h 0m 52s	1/4	Remote command execution failed: ssh: connect to host 192.168.1.101 port 22: No route to host
	Total Processes	OK	2007-07-25 11:10:22	2d 2h 55m 43s	1/4	PROCS OK: 181 processes
	httest	UNKNOWN	2007-07-25 11:10:20	0d 0h 3m 52s	1/4	Remote command execution failed: ssh: connect to host 192.168.1.101 port 22: No route to host
	nfs	CRITICAL	2007-07-25 11:10:20	0d 0h 0m 52s	1/4	CRITICAL - Plugin timed out while executing system call
sensors	UNKNOWN	2007-07-25 11:10:20	0d 0h 3m 52s	1/4	Remote command execution failed: ssh: connect to host 192.168.1.101 port 22: No route to host	
h2	Current Load	OK	2007-07-25 11:10:20	1d 19h 3m 30s	1/4	OK - load average: 0.00, 0.00, 0.00
	Current Users	OK	2007-07-25 11:10:20	1d 19h 3m 22s	1/4	USERS OK - 1 users currently logged in
	DNS	OK	2007-07-25 11:10:20	1d 19h 3m 7s	1/4	DNS OK - 0.006 seconds response time (n2.azule.com. 3600 IN A 192.168.1.102)
	Disk Space	OK	2007-07-25 11:10:20	1d 19h 3m 5s	1/4	DISK OK - free space: / 216467 MB (99% inode=99%); /var/run 1005 MB (99% inode=99%); /var/lock 1005 MB (100% inode=99%); /proc/bus/usb 1005 MB (99% inode=98%); /dev 1005 MB (99% inode=98%); /dev/shm 1005 MB (100% inode=99%); /lib/modules/2.6.20-16-generic/volatile 967 MB (96% inode=99%); /home 179169 MB (82% inode=99%);

# Nagios Disadvantages

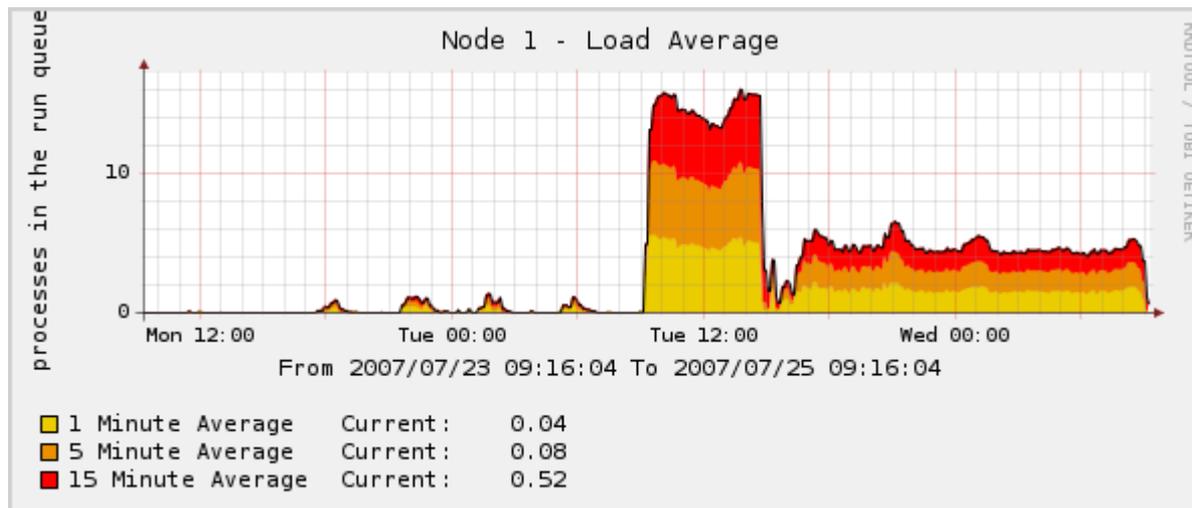


- Configuration
- Documentation
- High overhead
- Scalability

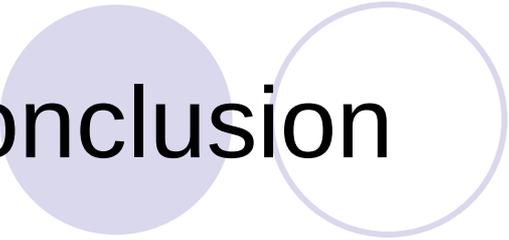
# Verify Nagios

- Generate performance degradation
- Independent testing scripts
- Cacti graphs and Nagios notifications

Cacti Load Average Graph



# Conclusion



- Scalability
- Robustness
- Extensibility
- Manageability
- Portability
- Overhead

