


Host and Service Monitoring at SLAC

Alf Wachsmann


Stanford Linear Accelerator Center

alfw@slac.stanford.edu




History

- “Monitoring” at SLAC:
 - Does not really exist
 - “ranger”
 - formerly known as “patrol”
 - modified at DESY: “scout”
 - local checks and fixes on Unix machines
 - emails to users or administrators about problems
 - BaBar started some monitoring using Ganglia (pretty graphs but no alarming)
- Monitoring at Fermi Lab: NGOP
- Monitoring at CERN: LHC Era Monitoring (Lemon)




Monitoring

- Why:
 - Health status at current time
 - Alarming in case of problems (ideally: fix the problem)
 - Long term trend analysis
- What:
 - Systems being alive and healthy
 - Services are running and functional
 - Service level agreements are met
- How:
 - Run a probe against or on a device
 - Gather data in central place
 - Display data and allow "data mining"




What to Monitor

- Networked devices:
network gear, printers, computers, appliances, UPS,
thermal elements
- Services provided by those devices:
OS, disk/memory space, CPU
NFS, AFS, printing, DHCP, electrical power,
temperature reading
- Service levels provided by those services:
50% free /tmp space, 90% batch computing,
95% AFS capacity, temperature $\leq 25^{\circ}\text{C}$




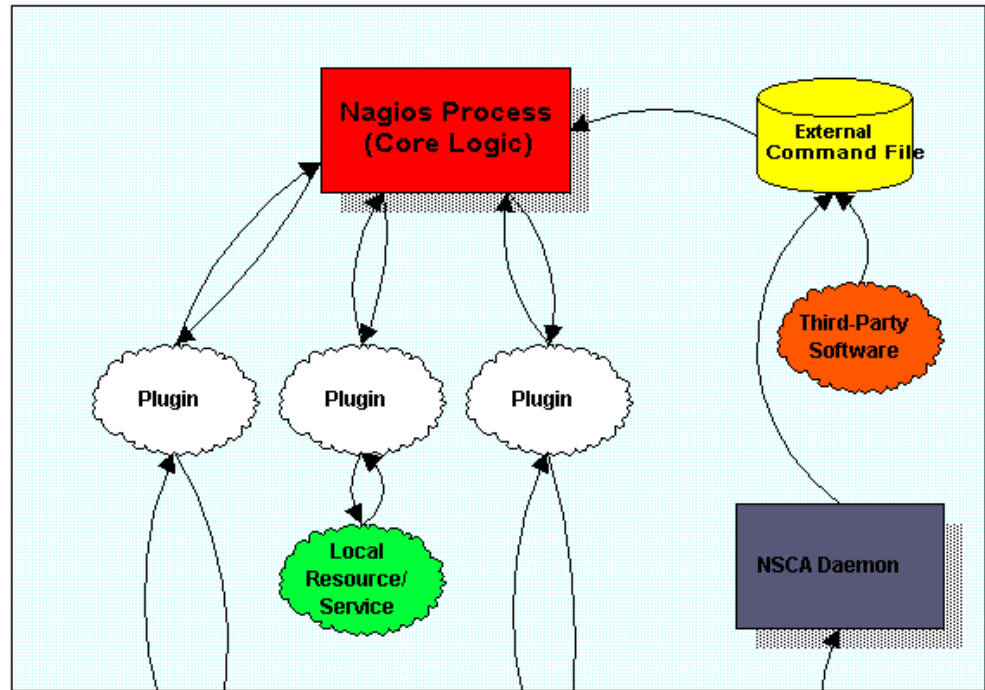
Nagios

- Open-source framework for Monitoring
<http://www.nagios.org/>
- Scheduler for data collection; alarming engine
- Not part of Nagios core:
 - plug-ins for measurements ("probes")
 - plug-ins for data storage and graphing
- Extensive collection of both as contributed source

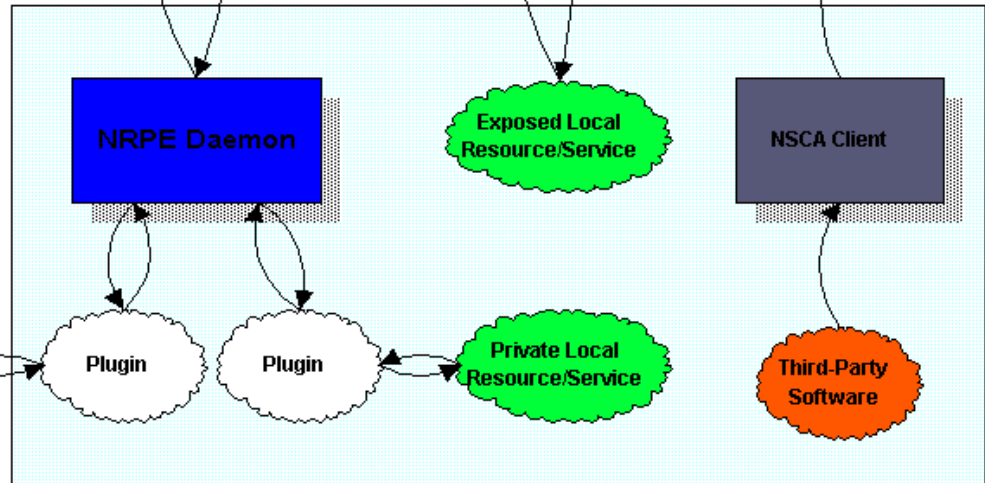


Nagios' Architecture...


Monitoring Host



Remote Host #1

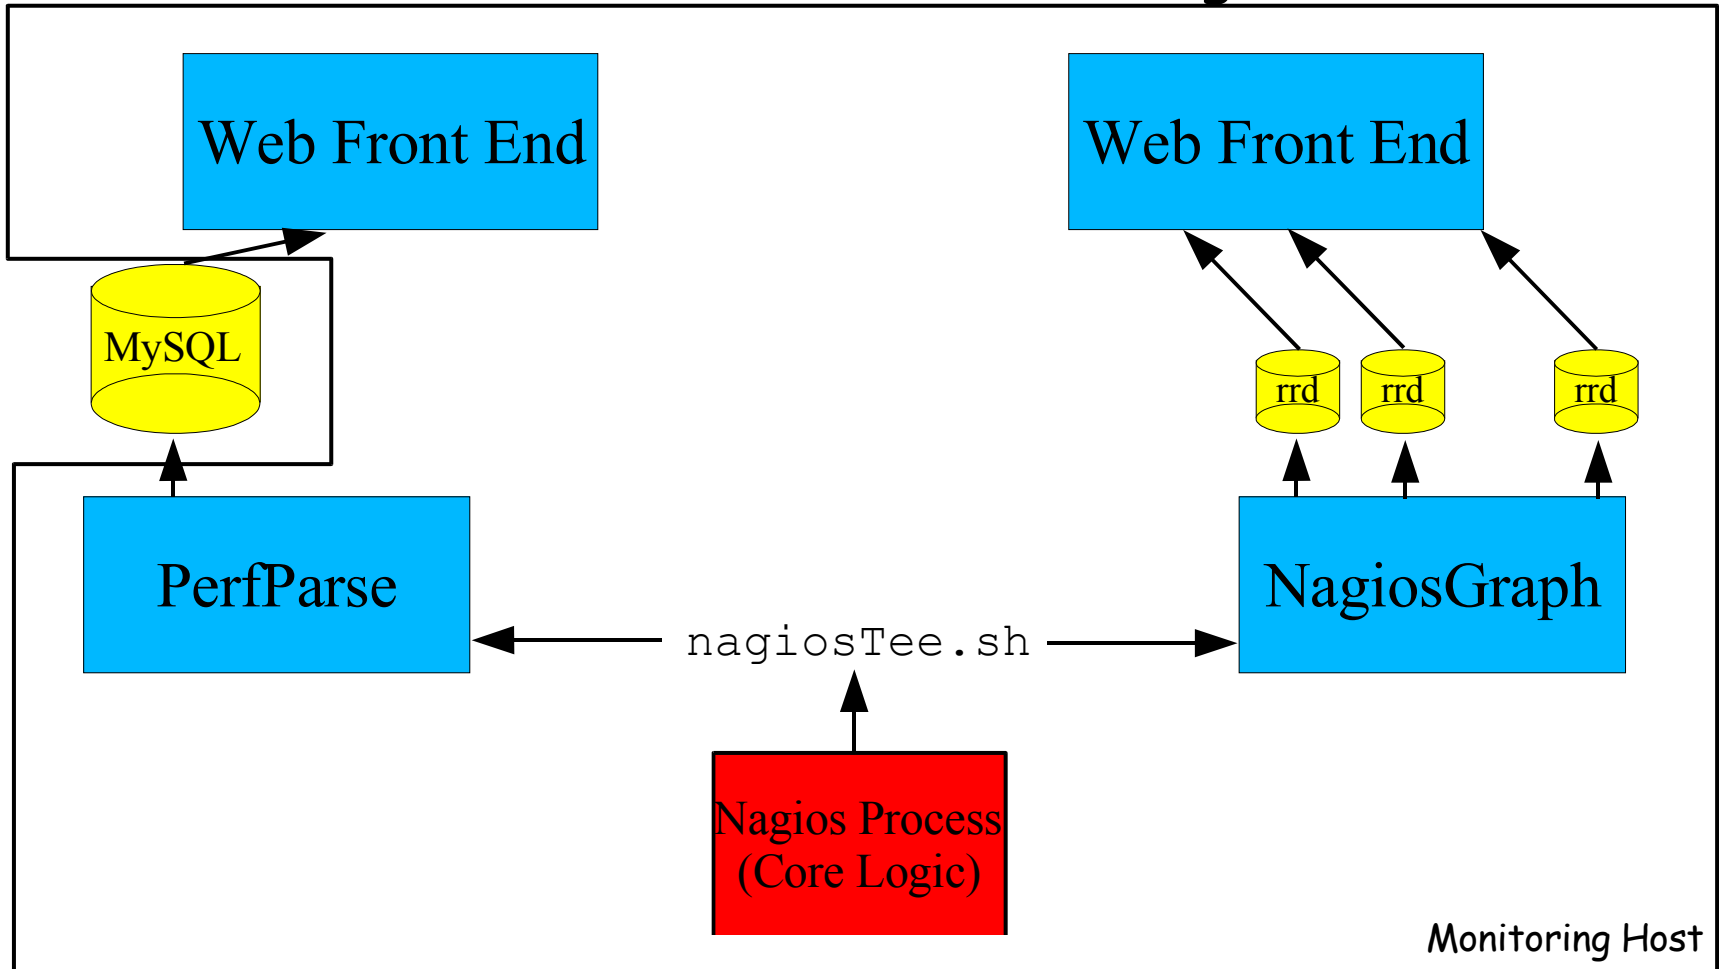



Remote Host #2



...Nagios' Architecture

Performance Data Processing





Internal Event Handler

- Alarming:
 - Built into Nagios
 - Visual
 - Email
 - Pager
 - Execute external program

Current Network Status
Last Updated: Tue Apr 5 10:37:02 PDT 2005
Updated every 90 seconds
Nagios® - www.nagios.org
Logged in as ?

Host Status Totals


Up	Down	Unreachable	Pending
926	0	0	0

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
1338	15	9	94	0

Status Summary For All Host Groups

Host Group	Host Status Totals	Service Status Totals
Printers (Printers)	172 UP	256 OK 12 WARNING 37 CRITICAL
SCS Printers (SCS.Printers)	6 UP	12 OK
Windows (Windows)	1 UP	23 OK 1 WARNING
AFS (afs)	14 UP	61 OK
interactive (Interactive)	42 UP	197 OK 1 WARNING 9 UNKNOWN 45 CRITICAL
Printing (lpd)	2 UP	10 OK
Monitoring (monitoring)	1 UP	8 OK
Network Monitoring (netmon)	11 UP	39 OK 1 WARNING 4 CRITICAL
nomas (nomas)	512 UP	504 OK 8 CRITICAL
orlows (orlows)	94 UP	94 OK
Sensors (sensors)	8 UP	8 OK
AWG Servers (sulkvs)	5 UP	20 OK
switch (switch)	77 UP	77 OK
Toasters (toasters)	16 UP	32 OK



Data Storage...

- Handler for Performance Data from Services
 - Nagios can have only one
 - I wrote my own that "tees" data to multiple ones: `nagiosTee.sh`

```
#!/bin/sh
```

```
# add these lines to your Nagios checkcommands.cfg file:
```

```
# define command {  
#   command_name   process-service-perfdata  
#   command_line   /path/to/nagiosTee.sh "$HOSTNAME$" "$HOSTALIAS$" "$HOSTADDRESS$" "$HOSTADDRESS2$"  
# }
```

```
HOSTNAME=$1
```

```
HOSTALIAS=$2
```

```
HOSTADDRESS=$3
```

```
# etc. etc.
```


```
# Now call the real perfdata handlers:
```

```
# Nagiosgraph:
```

```
/opt/nagios/nagiosgraph/insert.pl "$LASTCHECK\\t$HOSTNAME\\t$SERVICEDESC\\t$OUTPUT\\t$HOSTADDRESS2"
```

```
# PerfParse:
```

```
$/bin/perfparse_nagios_pipe_command.pl $/var/perfparse.pipe "$TIMET" "$HOSTNAME"
```





...Data Storage...

- NagiosGraph
 - one rrd file for each service check on each host
 - older data gets coarser

Host ↑	Service ↑	Status ↑	Last Check ↑	Duration ↑	Attempt ↑	Status Information
mailstore	PING	OK	04-05-2005 15:10:18	34d 3h 20m 37s	1/3	PING OK - Packet loss = 0%, RTA = 0.19 ms
	SNMP	OK	04-05-2005 15:12:06	35d 0h 40m 1s	1/3	mailstore OK - vol0_DiskUse=13,DiskReadB=3573129216,DiskWriteB=3683528704,NetRcvdKB=423053100,
surrey01	PING	OK	04-05-2005 15:12:17	5d 9h 53m 24s	1/3	PING OK - Packet loss = 0%, RTA = 0.54 ms
	SNMP	OK	04-05-2005 15:12:48	35d 0h 41m 11s	1/3	surrey01 OK - vol0_DiskUse=6,vol1_DiskUse=71,vol2_DiskUse=43,DiskReadB=1044803584,DiskWriteB=2596773888,Net
surrey02	PING	OK	04-05-2005 15:14:55	0d 17h 31m 19s	1/3	PING OK - Packet loss = 0%, RTA = 0.19 ms
	SNMP	OK	04-05-2005 15:12:34	35d 0h 41m 11s	1/3	surrey02 OK - vol0_DiskUse=77,DiskReadB=2679300096,DiskWriteB=4240404480,NetRcvdKB=1332526,Ne

(host=surrey05&service=SNMP&graph=118800)





...Data Storage

- PerfParse
 - MySQL database
 - full resolution for entire data collection
 - basic "data mining" web interface
 - easy to write scripts with SQL queries for more

No problem

Reboot!

PerfParse Metric Analysis v0.105.6

Return to main menu
Select new Metric

Host: **afs06** Service: **AFSfs** Metric: **idleThreads**

Scale: User: [v] User Min: 90.000000 Max: 125.000000

Period: Relative Absolute
Abs From: 2005-03-19 10:27:53 Abs To: 2005-03-23 10:27:53

Output: Histogram [v] Size: Normal [v]

Plot: Value Smooth* Warning Critical Stan. Dev* (* Where Used)

Host 'afs06' Service 'AFSfs' Metric 'idleThreads'

125.000
120.000
115.000
110.000
105.000
100.000
95.000
90.000

12 18 00 06 12 18 00 06 12 18 00 06 12 18 00 06
2005-03-20 2005-03-21 2005-03-22 2005-03-23

v0.105.6 Value Warn Crit


Draw Graph

Advanced Options.

Graph Legend: Host 'afs06' Service 'AFSfs' Metric 'idleThreads'

Smooth Plot: Gaussian
Average: Width (pixels): 14.0 (Large = Slow, see FAQ)


Save graph for future reference: Save with Title: Host 'afs06' Service 'AFSfs' Metric 'idleThread:'



Local Data Measurements

- Nagios master pulls (requests) data
- Most probes work locally on a machine
- Each probe decides whether measurement is OK, WARNING, CRITICAL, UNKNOWN
- Each probe can (but does not have to) return "performance data":


```
% ./check_users -w 5 -c 10  
USERS OK = 3 users currently logged in |  
users=3;5;10;0
```



Remote Data Measurements


- Several mechanisms to remotely execute a probe:
Nagios Remote Plugin Executor (NRPE) at SLAC
- NRPE executes probe on a remote machine on behalf of Nagios master:

```
% ./check_nrpe -H noric10 -c check_users  
USERS OK - 37 users currently logged in  
|users=37;110;150;0
```
- Also works with Windows (haven't tried Mac yet)
- Can run over SSL
- Can accept parameters from Nagios master (insecure!)




"Passive" Data Measurements

- External programs can write measurements into file which Nagios checks regularly ("push")
- Client-Server facility for remotely pushed data (Nagios Service Check Acceptor; NSCA)
- Enables hierarchical monitoring




Own Data Measurements

- Probes for Unix, Windows, Mac OS X, SNMP, HP's JetDirect etc. already exist
- Plug-in repository: <http://www.nagiosexchange.org/>
- Simple to write your own:
 - `AFS::MonitorPerl` module
 - `check_AFScm.pl`, `check_ubik.pl`, `check_AFSfs.pl`
 - Thermal sensors connected via terminal server
 - `check_thermal.pl`
 - More general SNMP plug-in
 - Put OID and name in configuration file
 - `Net::SNMP` to do the query
 - Output in Nagios format (incl. performance data)
- Will probably have to write my own Windows probes




Own Web Front-End

- BaBar uses Ganglia to monitor some of their computers and services
- Effort independent from SCS
- Presentation of same data seems more pleasing than Nagios'
- Project with 2 Stanford graduate students:
 - Nagios-to-Ganglia data converter which allows Ganglia web front-end to display Nagios data
- It is in production now
- Only some minor things are not working




Exporting to the Grid

- SLAC participates in the Open Science Grid (OSG)
- Web services interface to Nagios exists: NagiosWS
- No web services interface to PerfParse MySQL DB
- Summer student will tackle this




Problems

- SLAC or Nagios have some problems
 - Lot of probes are not usable in (large scale) production
 - Works on Windows but need to write our own probes
 - Like to have action that simply reboots a machine
 - DB support for configuration was ripped out in version 2.x
 - Web interface in version 1.x does not scale very well
 - Much better in version 2.x (still beta)
 - Still browser rendering performance problems



Scalability

- Nagios master machine at SLAC:
 - Sun V20z: dual 1.8GHz Opteron, 2GB main memory
 - Sun StorEdge 6120 Array "T4": hardware RAID 5, 2Gb/s FC, ~1.6TB usable disk space
 - Qlogic QLA2312 Fibre Channel Adapter
 - 64bit Fedora Core 3
 - XFS for MySQL DB and RRDtool files partitions
- Currently 407 servers with 1096 services configured
- Test with all ~2000 SLAC batch machines and 4 services on each showed no problems on master



References

- Nagios:
<http://www.nagios.org/>
- NRPE and other plug-ins (click on "Categories" tab)
<http://www.nagiosexchange.org/>
- NagiosGraph
<http://nagiosgraph.sourceforge.net/>
- PerfParse
<http://perfparse.sourceforge.net/>
- NagiosWS
<http://www.i-xs.de/nagiosws/>