

# Experiences monitoring grid services at NIKHEF

*Ronald Starink*

- **NIKHEF part of Dutch T1**
  - ~ 400 cores, 90 TB storage, ~ 150 hosts
  - Will grow ~ 10 x
  - Site without active monitoring, only Ganglia
- **Motivation**
  - Recognized need for site monitoring
  - Planned to look at Nagios
  - Opportunity to get started!
  - Hesitation: investment of time on short notice
- **Expectation**
  - Let's see what it does...
  - Gain experience with Nagios

- **Initial setup:**
  - Dedicated host for monitoring
  - Remote probes only (SAM)
  - No gLite
  - Only secure web server
- **Extended setup:**
  - New VM as gLite 3.0 UI
  - Remote probes
  - Local probes
    - CE, RB, BDII, site BDII, LFC, SE DPM, classic SE, MON
  - Only secure web server
- **Future:**
  - Add local probes to dedicated host
  - Use notifications, perhaps event handlers

- **Not difficult to setup**
  - Some manual actions
  - Early adopter: some small issues
  - Configuration script complex
- **Very useful!**
  - Almost immediate feedback when services are failing
    - Used on daily basis
  - Lots of tests (SAM and local)
    - Not always clear what they test. Documentation?
- **Some issues**
  - Permanently failing tests: SAM + local probes
  - Occasionally failing tests, spontaneous recovery
  - Proxy nearly expired: many tests failing?

- Nagios**
- 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:
- Comments
- Downtime
- Process Info
- Performance Info
- Scheduling Queue
- Reporting
  - Trends

**Current Network Status**  
 Last Updated: Tue Sep 11 09:22:08 CEST 2007  
 Updated every 90 seconds  
 Nagios@ - [www.nagios.org](http://www.nagios.org)  
 Logged in as  
 /O=dutchgrid/O=users/O=nikhef/CN=Ronald Starink

[View History For all hosts](#)  
[View Notifications For All Hosts](#)  
[View Host Status Detail For All Hosts](#)

**Display Filters:**  
 Host Status Types: All  
 Host Properties: Any  
 Service Status Types: All Problems  
 Service Properties: Any

**Host Status Totals**

Up	Down	Unreachable	Pending
9	0	0	0
<b>All Problems</b>		<b>All Types</b>	
0		9	

**Service Status Totals**

Ok	Warning	Unknown	Critical	Pending
142	0	0	3	5
<b>All Problems</b>		<b>All Types</b>		
3		150		

**Service Status Details For All Hosts**

Host	Service	Status	Last Check	Duration	Attempt	Status Information
<a href="#">siteinfo03.nikhef.nl</a>	<a href="#">org.wlcq.aggregate-status.local</a>	CRITICAL	09-11-2007 07:33:00	11d 21h 49m 8s	1/1	(Return code of 127 is out of bounds - plugin may be missing)
<a href="#">tbn20.nikhef.nl</a>	<a href="#">CE-sft-apel-Alice-remote</a>	CRITICAL	09-05-2007 11:23:35	20d 16h 30m 59s	1/1	SAM status: error
	<a href="#">org.wlcq.aggregate-status.sam</a>	CRITICAL	08-30-2007 10:56:30	15d 14h 41m 48s	1/1	CRITICAL - Used 25 metrics for calculation

3 Matching Service Entries Displayed

- **Not difficult to setup**
  - Manual configuration
  - Help from mailing list
- **Very useful!**
  - Good overview of service status
  - Fast feedback
  - Extensible
  - Jeff: *“For me it's already worth the investment”*
- **Not yet “production quality”, but close**
  - Some permanently failing tests
  - Does not yet feel 100% stable
- **Documentation on tests**
  - What is tested?
  - What does a failure mean?

- **Future:**
  - Part of infrastructure
  - Add more tests, perhaps contribute tests
  - Use alarms