

# NGOP

*J.Fromm*

*K.Genser*

*T.Levshina*

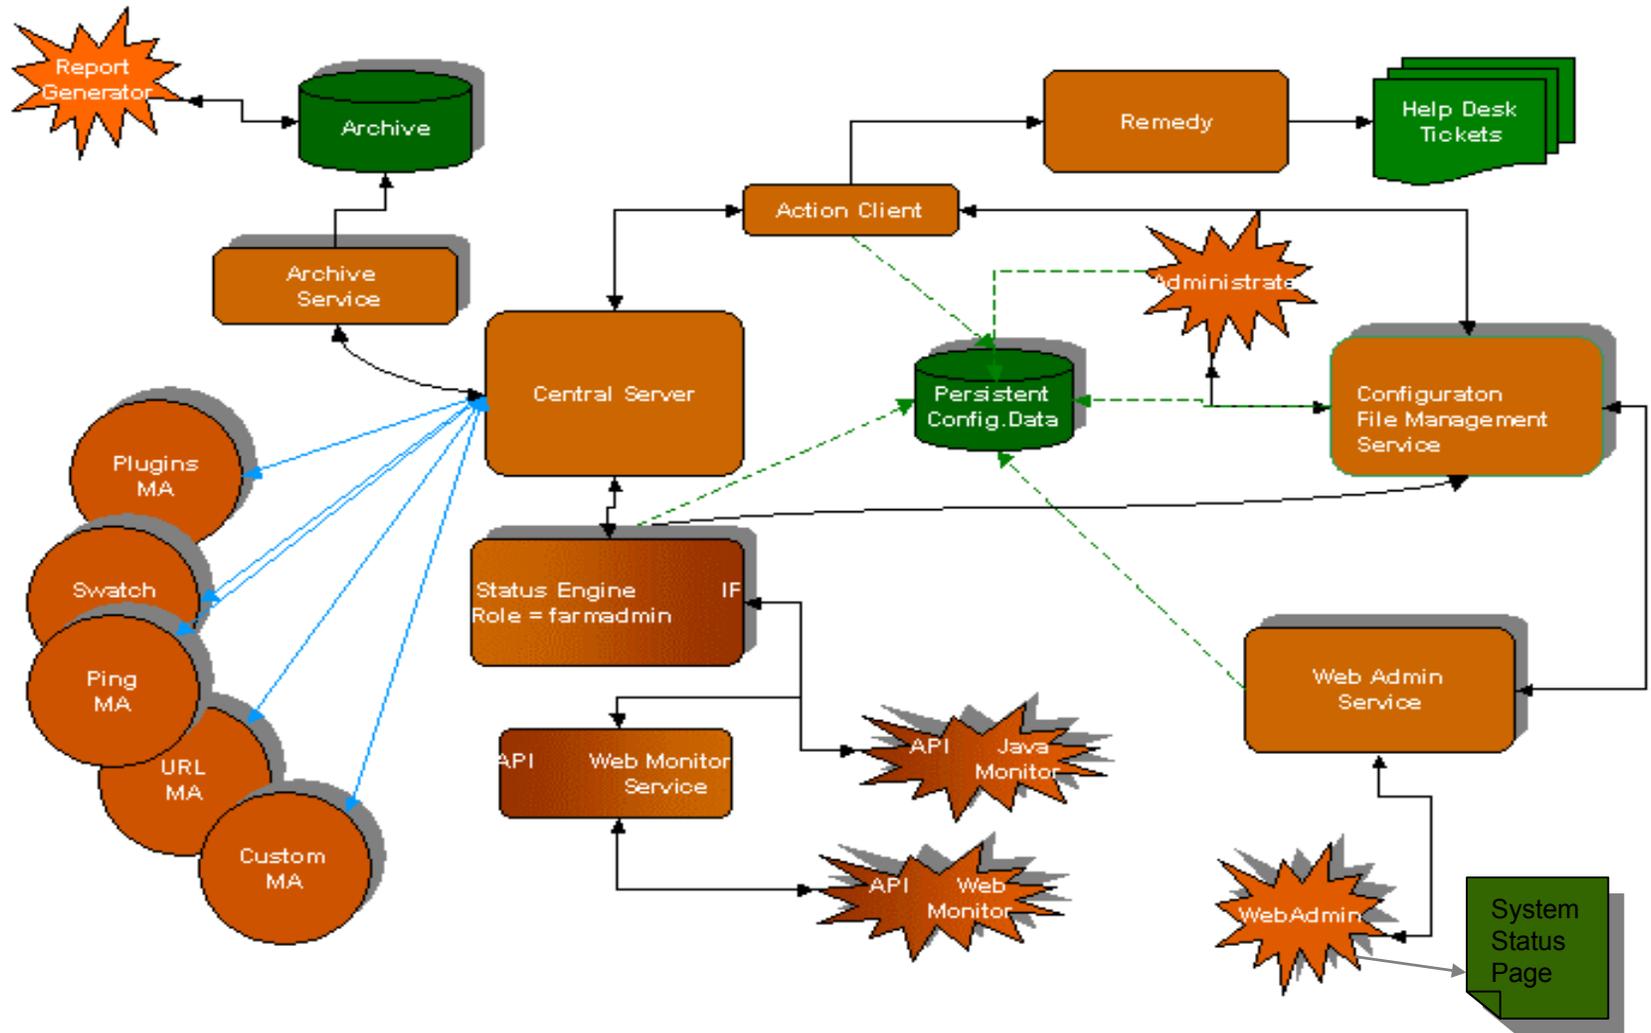
*M.Mengel*

*V.Podstavkov*

# What is NGOP and who is using it?

- What:
  - A Distributed Monitoring System that scales to the anticipated requirements for Run II (up to 10,000 nodes during next 5 years)
  - Provides active monitoring of software and hardware
  - Provides customizable service-level reporting
  - Facilitates early error detection and problem prevention
  - Provides persistent storage of collected data
  - Executes corrective actions and sending notifications
  - Offers a framework to create Monitoring Agents for monitoring the overall state of computers and software that are running on them.
- Who:
  - System administrators
  - Software administrators
  - Help Desk and computer center personnel
  - Management
  - Developers (the most curious ones)
  - End users

# NGOP Architecture



# NGOP Central Services

- NGOP Central Server (NCS):
  - collects messages from multiple monitoring agents
  - provides clients with requested information
  - forwards requests to Action Server to perform action
  - forwards all events to Archive Service
- Configuration Files Management Service (CFMS)
  - provides a central repository for all configuration files.
  - performs configuration sanity check
  - provides clients with component subscription list
  - allows dynamic reconfiguration
  - notifies clients about new configuration
- Action Server:
  - gets configuration information from the CFMS
  - gets action requests from the NCS
  - verifies user authorization to request the actions
  - verifies that monitored object associated with an action is not marked as “known bad”
  - performs actions
  - notifies the NCS about success/failure of performed actions

# Configuration Language

The NGOP configuration language provides a framework for creating monitoring tools.

NGOP configuration language

- written in XML
- allows the creation of hierarchies of monitored objects
- describes rules to determine the status of the object
- defines when and what kind of actions should be performed
- uses expansion mechanism that allows the replication of a particular fragment of an XML document
- uses conditions simplified handling of various fragments of XML that are relevant for a particular “role”

# Monitoring Agents (I)

- Monitoring Agent (MA) is process that monitors the characteristics of a particular monitored object and report a status to the NCS.
- MA can monitor multiple objects.
- MA can perform local actions or request NCS to perform centralized actions.
- NGOP provides a framework for creation of the MAs: either by using the MA API or the PlugIns Agent.
- PlugIns Agent:
  - runs on the local node
  - allows the monitoring of software or hardware components utilizing existing scripts or executables (plug-ins)
  - plug-ins should be able to measure and print some quantitative characteristics of the monitored objects.
  - uses template configuration file

# Monitoring Agents (II)

- Ping Agent:
  - runs on the central node, pinging remote nodes
  - sends ICMP packets to nodes listed in its configuration file.
  - performs route discovery and has an ability to distinguish failure to ping the node from the failure to ping the switch, as well as discovery of simultaneous multiple failures.
  - determines the boot time of a node as well as it's cpu load if rstatd daemon is running on remote node
- Swatch Agent
  - runs on the local node
  - watches a log file for lines matching a regular expression
- URL Agent
  - runs on the central node
  - scans given URL's for reachability and content

# Status Engine, Rules and Roles

- The Status Engine is the component that collects selected information from the NCS and processes it according to the specific rules.
- Multiple Status Engines can be running simultaneously each configured in such a way that reflects the interests of one particular group of people (role).
- Rules define the status of the monitored object
  - A Generic Rule sets the monitored object status based on the event received from the NCS.
  - A Dependent Rule sets the monitored element status based on the event received from the NCS and the status of each dependent monitored object in some group.
- Roles define what subset of the configuration will be seen by a particular group of users and what rules will be used to define the status of the monitored objects
- A full python API is provided allowing users to retrieve information about a particular monitored object. Web and Java Monitors are using API as well.

# Snapshots (Web GUI)

This screenshot shows the Netscape browser window titled "Netscape: NGOP Monitor". The address bar displays the URL `https://ngopcli.fnal.gov/cgi-bin/web_gui/web_gui.fcgi/`. The page content includes a navigation menu with links for News, Downloads, Software, Hardware, Developers, Help, Search, and Shop. Below the menu, there are links for "NGOP Web Admin" (Role: admin) and "Settings/Monitor". The main content area shows the date "Wed Mar 19 11:22:03 2003" and the status "[no actions pending]". A tree view on the left shows "allFermi" selected. The main display area is titled "Display for allFermi" and contains a grid of 15 network diagrams, each representing a different system: bphysics, cdf, cdcaf, cms, core\_afs, d0, farms, fnalu, ktev, mass\_storage, ninos, niscomp, ods, ppd, sdss, services, theory, and www. Each diagram shows a hierarchical network structure with various nodes and connections.

This screenshot shows the Netscape browser window titled "Netscape: NGOP Monitor" at a later time, "Wed Mar 19 11:26:26 2003". The address bar shows the same URL. The page content is similar to the first screenshot, but the main display area is titled "Display for OSHealth.fcdcode1". On the left, a tree view shows a more detailed hierarchy, including "OSHealth.CDFCAF" and "OSHealth.CDFCAF.Central", with "OSHealth.fcdcode1" selected. The main display area contains a grid of 15 system status icons, each with a label and a "fcdcode1" identifier. The icons represent various system components: /home, /root, /tmp, /var, automount, cpuLoad, cronD, inetd, klogd, procNum, rstatd, sshd, swap, syslogd, and xatpd. Each icon shows a different status, such as a folder, a person, or a graph, indicating the health of the system.

# Snapshots (Java Gui)

Ack	Date	ID	Type	Status	Sev_Level	State	
<input type="checkbox"/>	2002-12-05 12:25:29.620	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:25:25.680	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:24:12.320	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:23:42.040	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:23:41.060	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:23:27.070	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:11:34.650	nfsd.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:11:02.690	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:10:54.130	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:10:22.210	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:09:44.250	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 12:09:37.420	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 10:51:40.180	cpuLoad.d0bbin.OSHealth.d0bbin	sysUsage	0 Good	6 Bad	1	cpuLoad:26.2
<input type="checkbox"/>	2002-12-05 10:29:01.050	cpuLoad.d0bbin.OSHealth.d0bbin	sysUsage	0 Good	6 Bad	1	cpuLoad:26.9
<input type="checkbox"/>	2002-12-05 08:27:48.830	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 08:27:47.100	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 08:27:46.530	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 08:23:31.510	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 08:23:31.300	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-04 20:28:01.530	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-04 20:23:57.170	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-04 20:23:41.810	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-04 20:23:33.430	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-04 18:22:44.900	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-04 18:22:15.250	enstore_cntn.d0bbin.OSHealth.d0bbin	Daemon	0 Good	6 Bad	1	-
<input type="checkbox"/>	2002-12-05 15:00:43.860	memory.fnd0318.OSHealth.fnd0318	sysUsage	0 Good	5 Error	1	total:1027424
<input type="checkbox"/>	2002-12-05 14:57:59.300	memory.fnd056.OSHealth.fnd056	sysUsage	0 Good	5 Error	1	total:1027928

# NGOP Archiver

- Responsible for storing/retrieving messages generated by NGOP.
- Data stored in Oracle database
- Cleanup process runs daily – 14 days of data is available.
- Archive server caches messages from the NGOP Central Server. A separate process (Database Interface) periodically reads cached messages and puts them in Oracle.
- “Best effort” used to store messages. Some messages may be dropped.
- Web based interface

# Snapshots (Archiver)

File Edit View Go Bookmarks Tools Window Help

Back Forward Reload Stop http://miscomp.fnal.gov/misweb/cgi/misweb.pl?tables=MESSAGE\_LOG Search Print

Home Bookmarks The Ultimate White P... Members WebMail Connections BizJournal SmartUpdate Mkplace

This form allows you to query the MISCOMP Database using the tables:  
MESSAGE\_LOG

You can pick the columns you want printed in your output, and pick the conditions to use for limiting your query. You are allowed to pick any number of columns for printing, but only 3 search conditions from each of the table. Query output will be printed in html tables.

Output Columns for MESSAGE\_LOG

Description  Event\_Value  Severity\_Level  
 Description\_Id  Format\_Version  State  
 Event\_Name  Monitored\_Object\_Id  Timestamp  
 Event\_Type  Monitoring\_Agent\_Id  Unix\_Timestamp

Query Options for MESSAGE\_LOG

Monitored\_Object\_Id like %fncdf180%  
 Choose One like  
 Choose One like

Do you want your query options combined with  AND  OR logic applied?

Sort Order Options

Timestamp Descending  
 Sort Order Ascending  
 Sort Order Ascending

How many rows do you want to see per page? 20  
 What is the maximum number of rows you want to see? 240 Select Output Type

Query Action Run Edit Make Clear Form

For help contact [miscomp@fnal.gov](mailto:miscomp@fnal.gov) MISW

MISWEB Query Interface

Document: Done (0.344 secs)

File Edit View Go Bookmarks Tools Window Help

Back Forward Reload Stop http://miscomp.fnal.gov/misweb/cgi/misweb.pl Search Print

Home Bookmarks The Ultimate White P... Members WebMail Connections BizJournal SmartUpdate Mkplace

\*\*\*\*\*  
 \* You are pointing to the procpd1 database.  
 \* Try the various production query pages,  
 \* available under the [MISCOMP Home Page](#)  
 \*\*\*\*\*

## MISCOMP Web Query Interface

### MESSAGE\_LOG

Description	Monitored_Object_Id	Timestamp	Monitoring_Agent_Id
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	Pingfncdf171-202.
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	Pingfncdf171-202.
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	Pingfncdf171-202.
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	Pingfncdf171-202.
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	Pingfncdf171-202.
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	Pingfncdf171-202.
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	Pingfncdf171-202.
None	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	LinuxHealthAgent
Average_cpu_load_is_between_4_and_8_during_last_15_min	cpuLoad.fncdf180.OSHealth.fncdf180	06-DEC-02	Pingfncdf171-202.

Document: Done (1.149 secs)

# WEB Admin tool, Remedy

- Web Admin Tool can mark any monitored object as known to be out of service, so this object will be excluded from determination of the status of the dependent monitored objects
  - Schedules maintenance in advance
  - Provides multiple maintenance intervals
  - Provides “cron” like maintenance intervals
  - Shows hierarchy of clusters/nodes, and system/elements
  - Provides search for particular host/clusters
  - Provides secure access for authorized users
  - Keeps change log
- NGOP is interfacing Remedy Help Desk using Remedy API to generate help desk tickets.

# Snapshots (Web Admin Tool)

Cluster/System view - Microsoft Internet Explorer

Address: https://ngopcl.fnal.gov/cgi-bin/webadmin

Fast Search (Cluster or Host name)

Changes History Known Status System Status

Last modification: user kschu on Fri Mar 14 13:53:20 2003 from hamshack.fnal.gov

Cluster/Hosts	Clusters/Systems
Cluster: Worker001-016	
frdfca001 8to17by5	Change Status: good
frdfca002 8to17by5	Change Status: good
frdfca003 8to17by5	Change Status: good
frdfca004 8to17by5	Change Status: good
frdfca005 8to17by5	Change Status: good
frdfca006 8to17by5	Change Status: good
frdfca007 8to17by5	Change Status: good
frdfca008 8to17by5	Change Status: good

Cluster/System view - Microsoft Internet Explorer

Address: https://ngopcl.fnal.gov/cgi-bin/tstadmin

Fast Search (Cluster or Host name)

Changes History Known Status System Status

Last modification: user tlevshin on Thu Mar 13 15:38:18 2003 from hppc.fnal.gov

Status for: d0enmvr26a

Service Type: 0by0

Start Date: 01-09-2003 09:10

End Date: [ ]

Duration: -Select-duration-

Good Bad Repair Maintenance

Cron: Min Hour Day Month DoW Durtn

Off-line for now

You can  the changes or

# Scope of deployment

- Monitoring a total of 1420 nodes
- Number of Monitored Objects ~ 32,000
- Number of agents ~ 2,500
- Number of Status Engines 6
- Average rate of events per day ~ 3,000
- Two dedicated computers:
  - ngopsrv
    - Central Server
    - CFMS
    - Action Server
    - Ping Agents
    - URL Agents
  - ngopcli
    - Status Engines
    - Web Admin Tool
    - Web Service

# Implementation Details

- Written primarily in Python (some modules in C, NGOP Monitor in Java)
  - Compatible with python 2.1
  - Java 1.4.0 and higher
  - Python code ( ~18,000 lines), C code (~ 350 lines), Java (~ 3,000 lines)
- Uses XML (and partially MATHML) for all configuration files. DTD files are provided with distribution.
  - Central configuration ( ~ 8,000 lines)
  - Central Agents (URL, Ping) configuration ( ~ 8,000 lines)
- Uses Oracle Database for event logging
- Product availability:
  - Monitoring Agents are available on Linux, Irix, Solaris
    - “PlugIns” Agent was ported to Windows
  - NGOP Central Services, Web Admin Tool run on Linux
  - NGOP Web GUI is available via any Web Browser, NGOP Java Monitor runs on Linux, Windows and Sun

# Who else is using it and how you can use it too?

- Working installation (beta – release) in IN2P3 Lyon (P. Olivero)
  - 779 hosts
  - 7 roles
  - 40 Applications
  - 9 Printers queues
  - 42 drives-status
- NGOP version v2\_1 is in Fermi Tools, could be download via anonymous ftp
- More info:
  - <http://www-isd.fnal.gov/ngop>
    - Documentations
    - Tutorials
  - Email: [ngop-team@fnal.gov](mailto:ngop-team@fnal.gov)

# Summary

- A comprehensive framework was created to fulfill monitoring needs of system administrators, operators and end users.
- A structured framework was provided to collect events, alarms and actions.
- NGOP Service has already proven itself in helping to increase the systems uptime and efficiency.
- NGOP interface to the Fermilab Remedy Help Desk system provides means for possible future complete automation of the notification process.
- Comprehensive documentation is provided.
- Creating configuration and rules is quite complicated and time consuming procedure. It requires knowledge of XML and NGOP configuration language. The tools that shield end users from these do not exist.