

# Adaptive Grid Scheduling and Resource Management

Next Generation Networking  
Multi-Service Networks 2004

Aleksandar Lazarević, Dr Lionel Sacks  
University College London



# The Grid

---

- What is Grid?
  - De-centralised, distributed, dynamic, cross-domain
- Grid vs. Cluster
  - Accessibility, Time- and Space-sharing, Administration, Flexibility, Reliability, Cost
- What Grid isn't?
  - New batch system, new distributed storage
- What can Grid become?

# Harnessing the Grid

---

- Information Capture vs. Complexity
  - Who really needs to know what?
  - What is worth remembering?
- Scheduling Problem
  - Holy grail: Deadline & Cost based
  - Matching jobs to resources requires predicting future performance
- Administration & Security Burden
  - Increasingly complex and time consuming

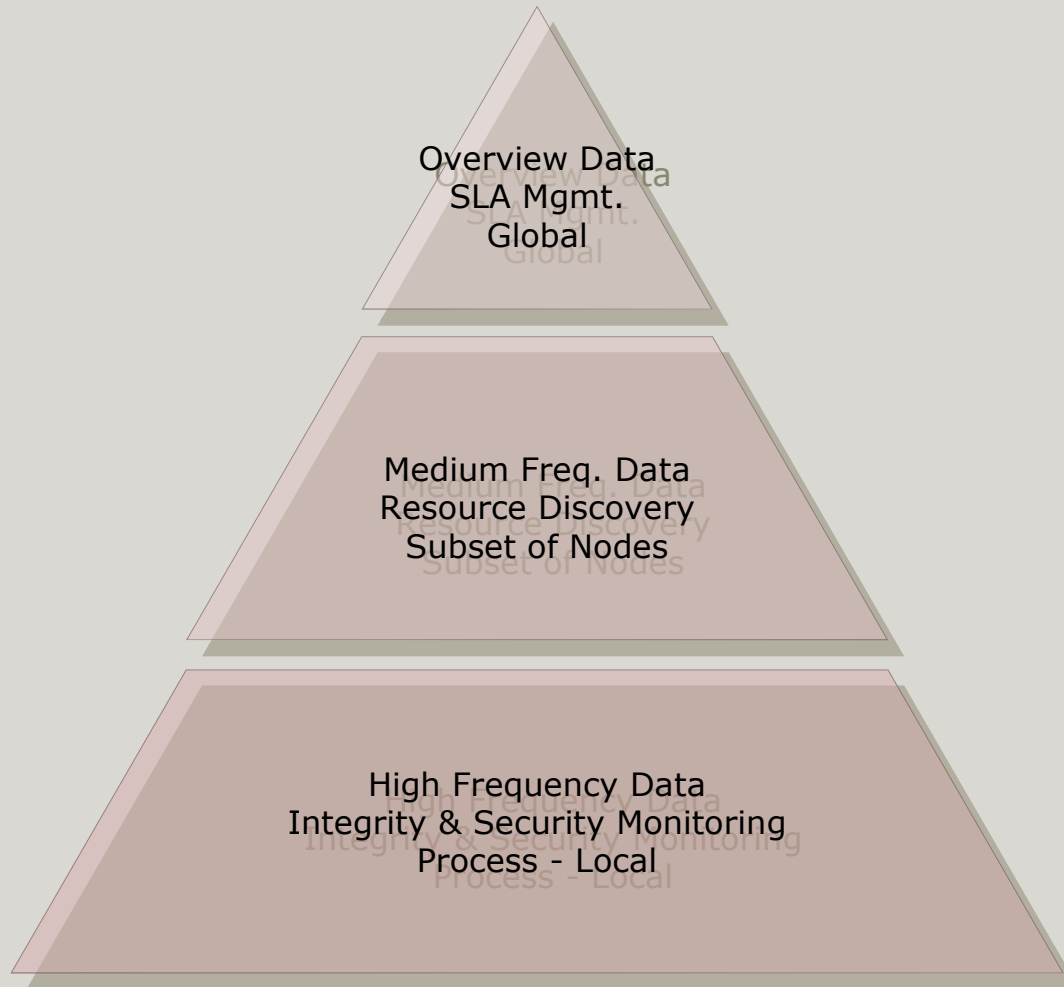
# SO-GRM Approach

---

- Multi-tier Information Dissemination
- Self-organising Resource Discovery
- Predictive User Orientated Scheduling
- Autonomous Security Monitoring & Enforcement
- Policy-based Management
- Modest Resource Requirements

# Information Dissemination

---



# SO-GRM Approach

---

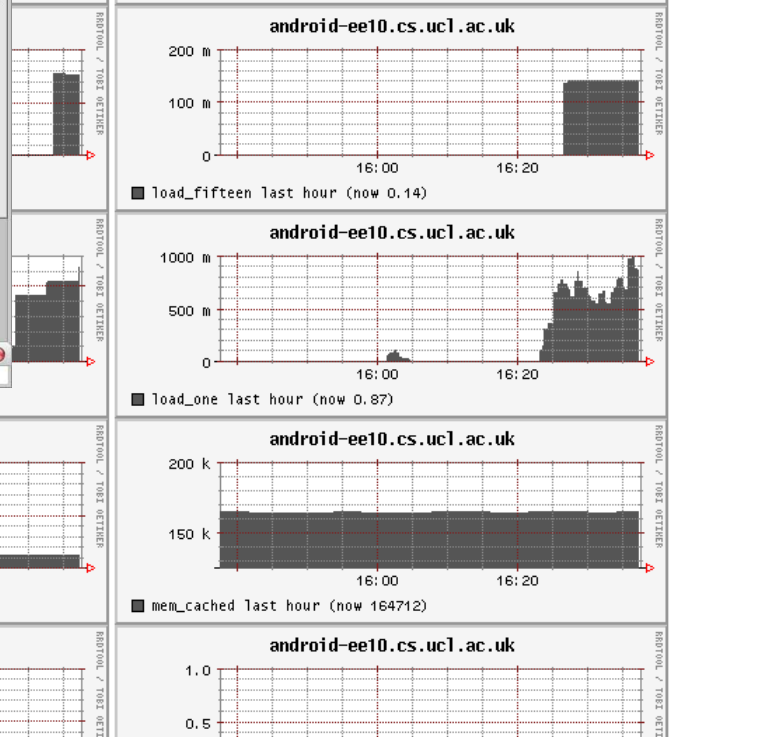
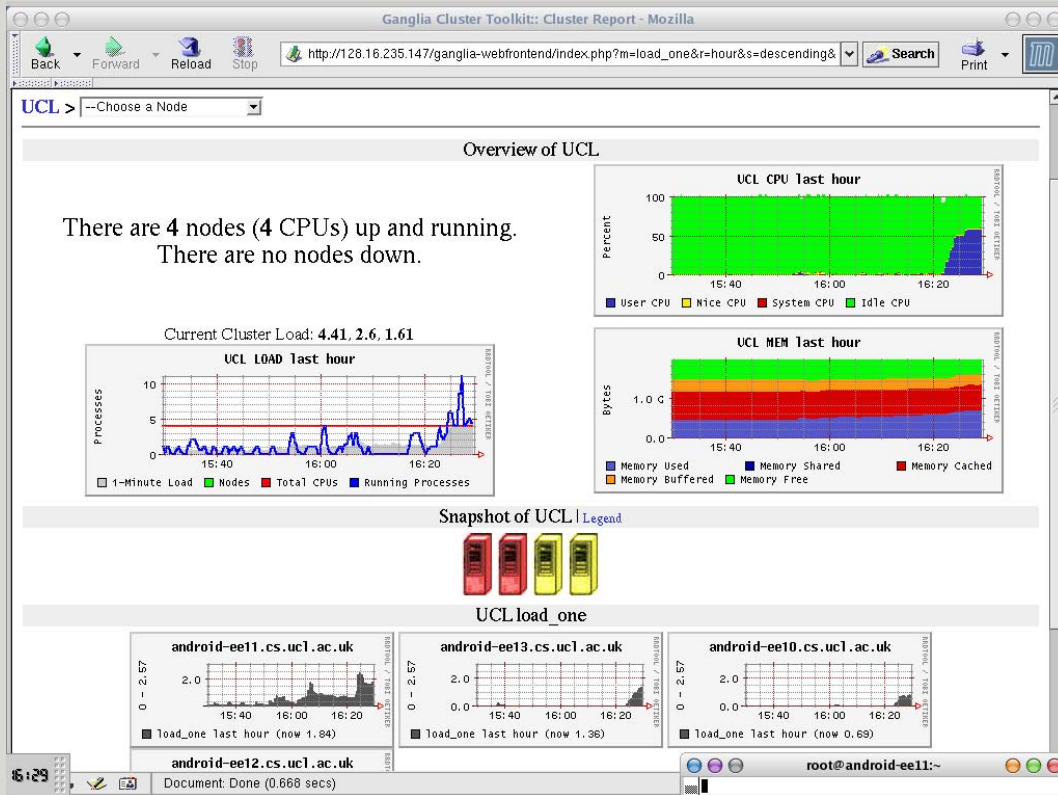
- Multi-tier Information Dissemination
- Self-organising Resource Discovery
- Predictive User Orientated Scheduling
- Autonomous Security Monitoring & Enforcement
- Policy-based Management
- Modest Resource Requirements

# Measurements

---

- Based on Ganglia, support for multicast
- Extended to monitor Grid submitted jobs
  - Missing support in middleware for unique job ID and local machine PID
- High-res data for resource discovery
- Low-res data integrated into MDS
- Ganglia Round-Robin Database used for statistical mining

# Ganglia Monitoring





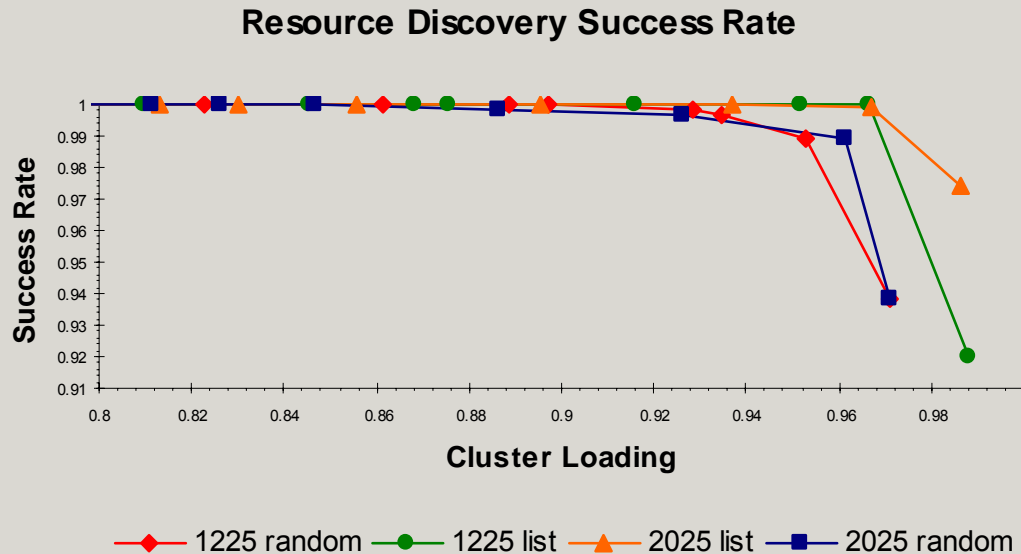
# Measurements

---

- Based on Ganglia, support for multicast
- Extended to monitor Grid submitted jobs
  - Missing support in middleware for unique job ID and local machine PID
- High-res data for resource discovery
- Low-res data integrated into MDS
- Ganglia Round-Robin Database used for statistical mining

# Resource Discovery

- Agent using queries and advertisement
- Queries on any number of metrics
- Develops a small-worlds metric specific overlay topology
- Scales well, has good accuracy



# Scheduling

---

- Our lightweight context rules out:
  - Application recompilation or re-linking
  - Resource Utilisation Prototyping
  - Resource Loading Instrumentation
  - Resource Intensive Predictions
- Develop a fast prediction based on statistics of previous runs
- Gradually improve confidence levels

# Scheduling

---

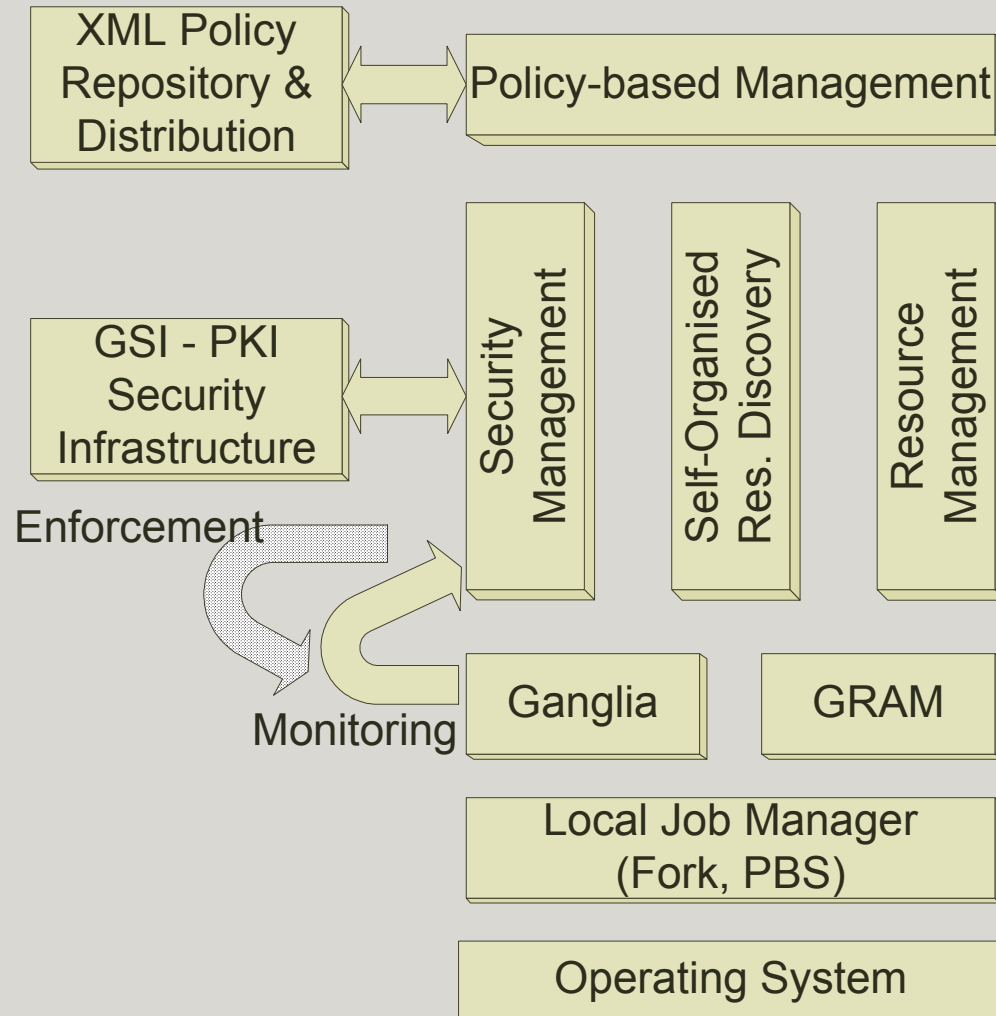
- Exploit time and space locality of jobs and target machines
- Process monitoring information close to source
- Evaluate methods for effective refinement of schedules and compact representation:
  - Stochastic intervals
  - Statistical models
  - Wavelets

# Order Out Of Chaos

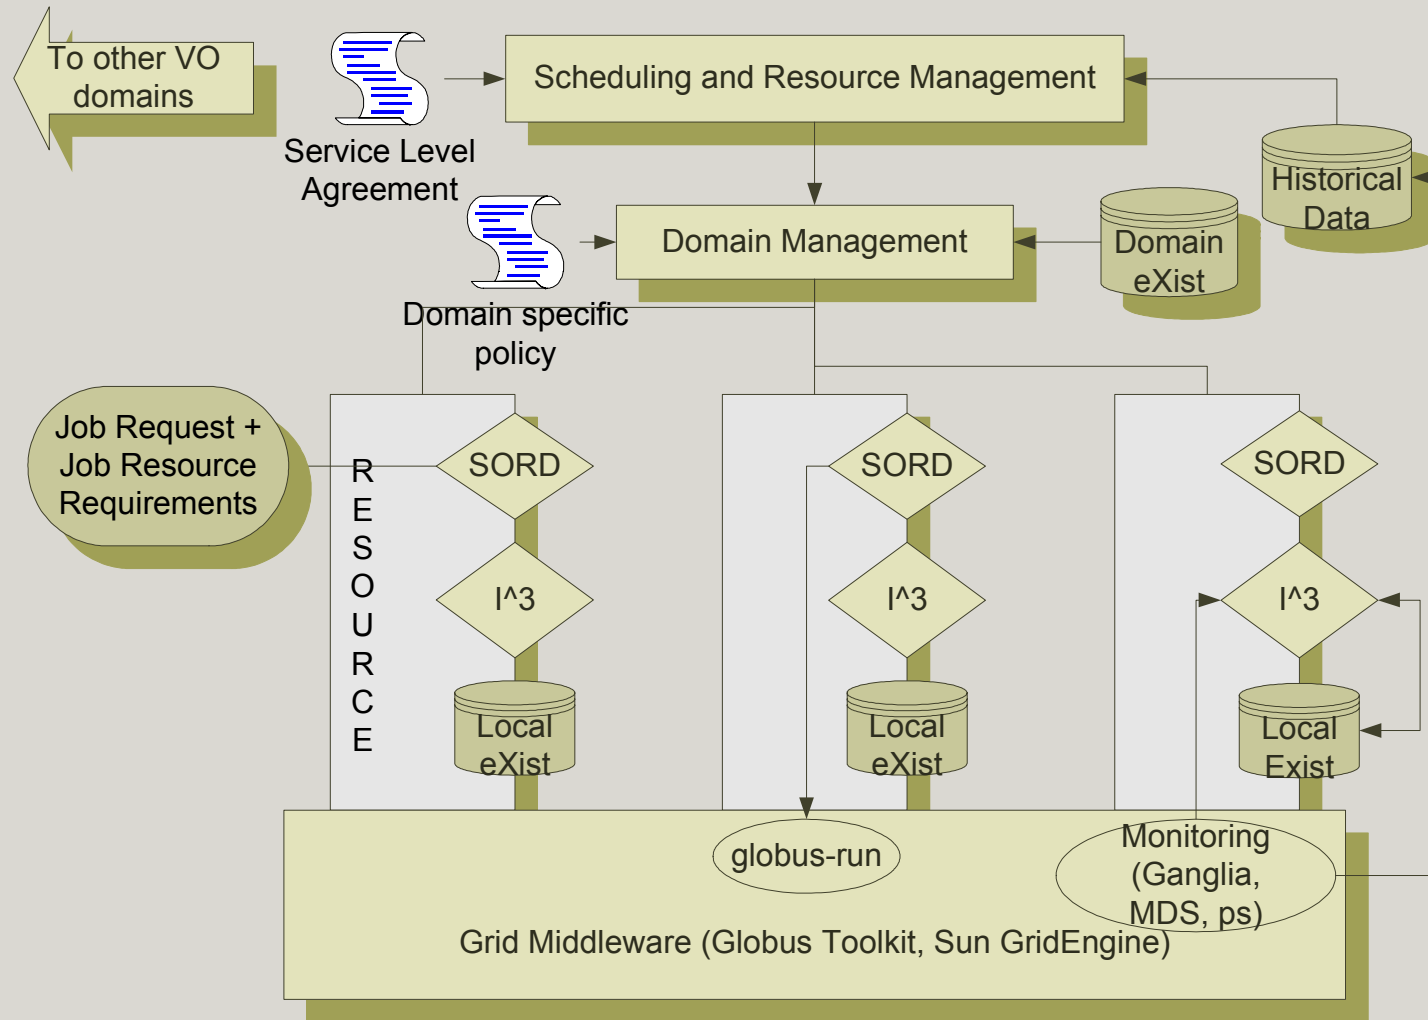
---

- SLA Management layer with overall picture
- Consumers will accept variable – if it can be somehow quantified
- Ensure reliable and predictable output from a fluctuating commodity system
- Fuse probabilistic & deterministic resources to offer QoS
- Relieve operators of tedious admin work

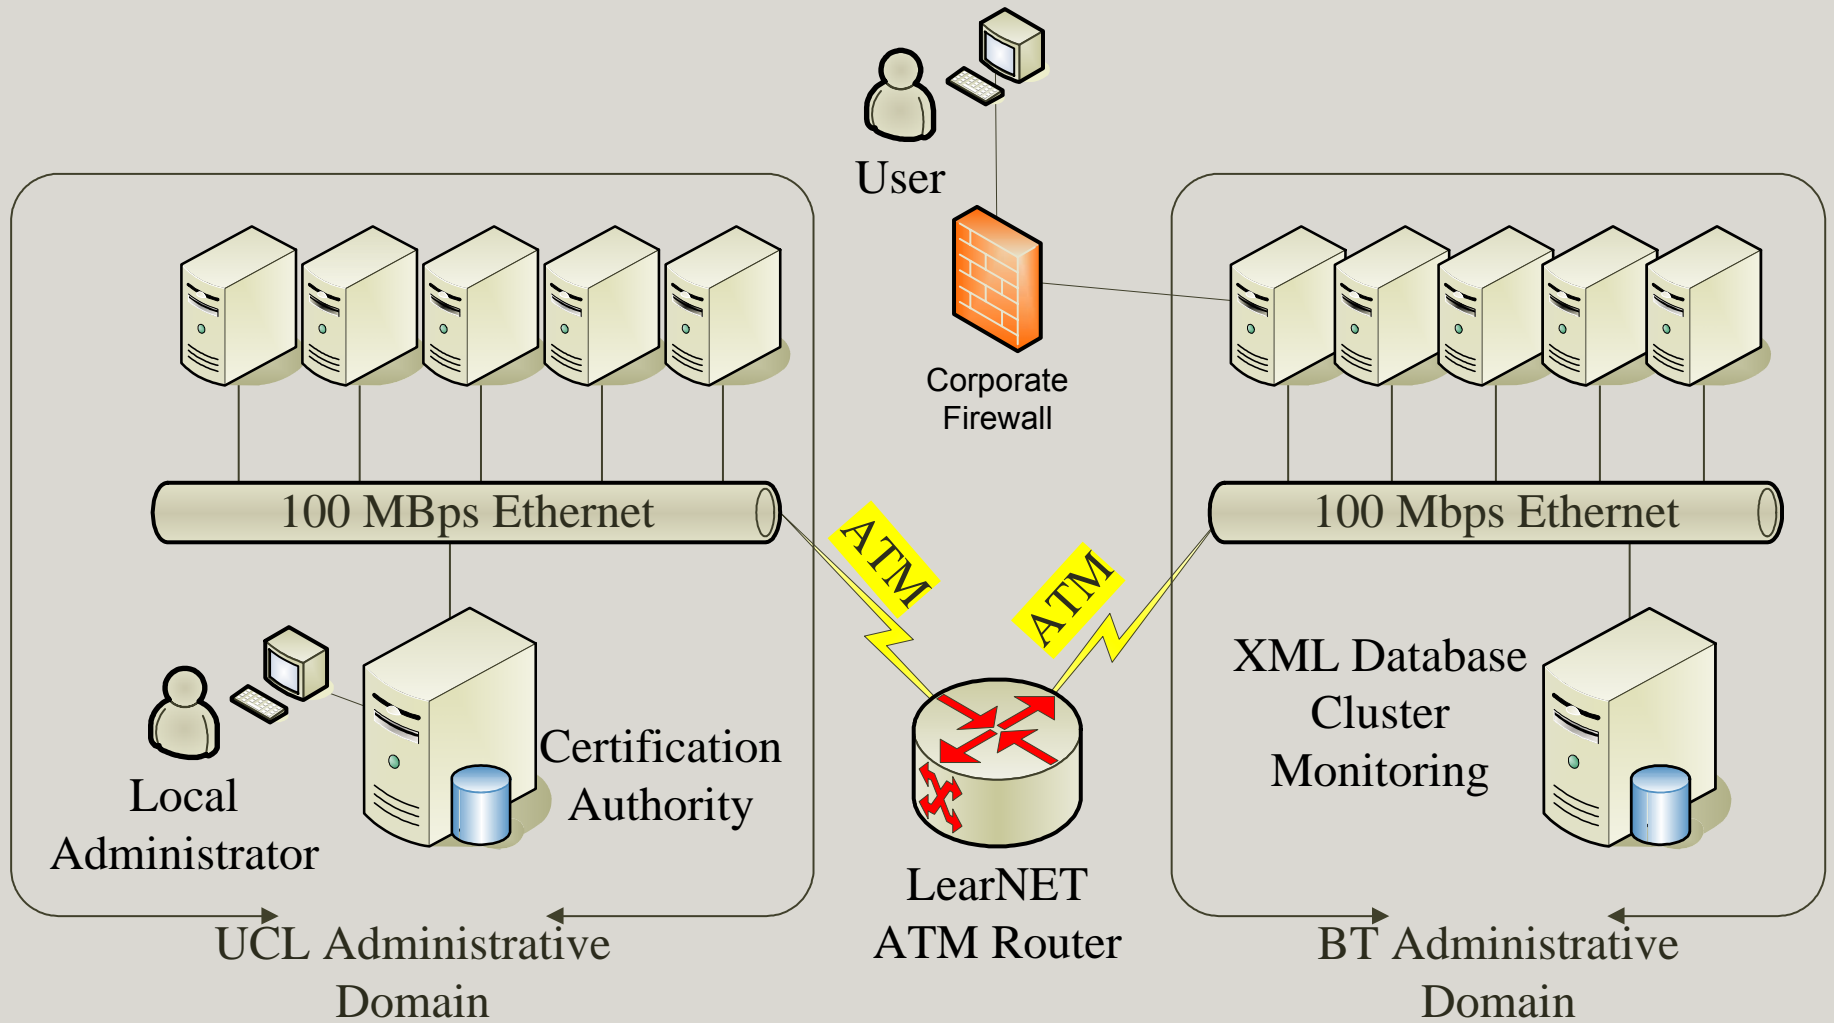
# Implementation – Node Level



# Implementation – VO Level



# Testing – Network Layout






# Testing – Node Monitoring

Ganglia Cluster Toolkit: Host Report - Mozilla

http://128.16.235.147/ganglia-webfrontend/index.php?c=UCL&h=android-ee11.cs.ucl.ac.uk

UCL > android-ee11.cs.ucl.ac.uk

### android-ee11.cs.ucl.ac.uk Overview

 This node is up and running

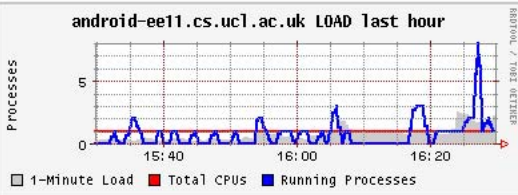
#### Time and String Metrics

Name	Value
boottime	Mon, 22 Mar 2004 16:03:48 +0000
gexec	OFF
gmond_started	Tue, 23 Mar 2004 14:22:34 +0000
ip	128.16.235.147
machine_type	x86
os_name	Linux
os_release	2.4.18-14
reported	Tue, 23 Mar 2004 16:29:31 +0000
sys_clock	Tue, 23 Mar 2004 14:22:34 +0000
uptime	1 day, 0:25

#### Constant Metrics

Name	Value
cpu_idle	99.9 %
cpu_num	1
cpu_speed	2088 MHz
mem_total	513952 KB
mtu	1500 B
swap_total	1052248 KB

#### android-ee11.cs.ucl.ac.uk LOAD last hour

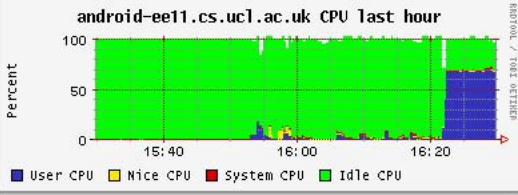


Proceses

15:40 16:00 16:20

1-Minute Load Total CPUs Running Processes

#### android-ee11.cs.ucl.ac.uk CPU last hour

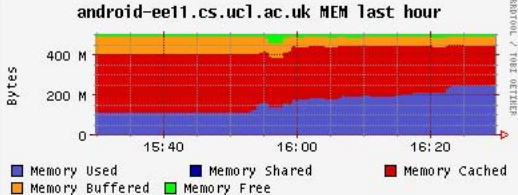


Percent

15:40 16:00 16:20

User CPU Nice CPU System CPU Idle CPU

#### android-ee11.cs.ucl.ac.uk MEM last hour



Bytes

15:40 16:00 16:20

Memory Used Memory Shared Memory Cached Memory Buffered Memory Free

16:30 Document: Done (2.19 secs) root@android-ee11:~

# Testing – VO Monitoring

The screenshot displays the Android Studio environment with several windows open:

- XML Editor:** Shows an XML snippet for a resource: 

```
<resource>
<name>physical-mem</name>
<operator>MoreEqual</operator>
<value>50</value>
<type>integer</type>
</resource>
```
- System Monitors:** Four panels for devices android-ee10, android-ee11, android-ee12, and android-ee13. Each shows CPU usage (0%, 1%, 1%, 0% respectively) and memory usage (69%, 68%, 0%, 98% respectively).
- Logcat (sogrm@android-ee10):** Shows the start of a test run: 

```
[sogrm@android-ee10 sogrm]$ ./startRD
128.16.235.83
10101
128.16.235.147
10101
Zero sized Pool created into eXist
```
- Logcat (sogrm@android-ee11):** Shows query results and match findings: 

```
releasing query results
sordProtocol_T::run()->eventtype got: Query
releasing query results
releasing query results
releasing query results
releasing query results
releasing query results
Match found for: USERid:USER0 USERid:USER0
Match found for: Executable:/home/globus/gridload Executable:/home/globus/gridload
Match found for: mipsAvail:4054 mipsAvail:500
Match found for: physical-mem:7 physical-mem:1
```
- Logcat (sogrm@android-ee13):** Shows similar query results and match findings: 

```
releasing query results
releasing query results
Match found for: USERid:USER0 USERid:USER0
Match found for: Executable:/home/globus/gridload Executable:/home/globus/gridload
Match found for: mipsAvail:4063 mipsAvail:500
Match found for: physical-mem:170 physical-mem:1
```
- Logcat (sogrm@android-ee12):** Shows a stack trace for a parsing error: 

```
at org.apache.crimson.parser.Parser2.parse(Parser2.java:333)
at org.apache.crimson.parser.XMLReaderImpl.parse(XMLReaderImpl.java:448)
at org.jdom.input.SAXBuilder.build(SAXBuilder.java:287)
at org.jdom.input.SAXBuilder.build(SAXBuilder.java:617)
at org.jdom.input.SAXBuilder.build(SAXBuilder.java:599)
at domainStore.getResource(domainStore.java:135)
at sordProtocol_T.getMyAvailability(sordProtocol_T.java:411)
at sordProtocol_T.run(sordProtocol_T.java:43)
Root cause: org.xml.sax.SAXParseException: Document root element is missing.
at org.apache.crimson.parser.Parser2.fatal(Parser2.java:3339)
at org.apache.crimson.parser.Parser2.fatal(Parser2.java:3327)
at org.apache.crimson.parser.Parser2.parseInternal(Parser2.java:635)
at org.apache.crimson.parser.Parser2.parse(Parser2.java:333)
at org.apache.crimson.parser.XMLReaderImpl.parse(XMLReaderImpl.java:448)
```
- Terminal (root@android-ee11):** Shows a command: 

```
[root@android-ee11 root]# xwd -root -out "aleks/dump3,xwd"
```

# Q & A

---