Project Acronym	**Fed4FIRE**
Project Title	**Federation for FIRE**
Instrument	**Large scale integrating project (IP)**
Call identifier	**FP7-ICT-2011-8**
Project number	**318389**
Project website	**www.fed4fire.eu**

D8.8 – First Report on First Level Support (FLS)

Work package	WP8
Task	T8.1
Due date	30/11/2014
Submission date	22/12/2014
Deliverable lead	Dai Davies (DANTE)
Version	1.0
Authors	Dai Davies (DANTE) Akil Radhakrishnan (DANTE)
Reviewers	David Margery (INRIA), Brecht Vermeulen (iMinds)

Abstract	This document describes the experience gained in the first 9 months of Operations of FLS. It analyses the operational data gathered during this period and summarises the lessons learned.
Keywords	First Level Support, Trouble ticket, Operations, Maintenance

Nature of the deliverable	R	Report	X
	P	Prototype	
	D	Demonstrator	
	O	Other	
Dissemination level	PU	Public	X
	PP	Restricted to other programme participants (including the Commission)	
	RE	Restricted to a group specified by the consortium (including the Commission)	
	CO	Confidential, only for members of the consortium (including the Commission)	

Disclaimer

The Fed4FIRE project received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no FP7-ICT-318389.

Executive Summary

This report analyses the activity of the First level Support (FLS) function of Fed4FIRE during its initial nine months of operation. FLS was defined at the start of the project and input its initial requirements to the Architectural work of Fed4FIRE. The requirements led to the development of a Dashboard that displays a standardised set of operational alarms from each of the test-beds in the federation. Using this dashboard, the FLS team monitors and manages faults detected in the test-beds in co-operation with experts in the individual test-beds. Incidents are logged by FLS in a Trouble Ticket System (TTS) which acts as a central repository for Fed4FIRE. It was originally envisaged that the FLS team would also assist experimenters using the federation. In practice the activity of assisting experimenters has been implemented using an e-mail list observed by technical experts form the test-beds and the project developers.

FLS collects operational data from the dashboard, from activity on the e-mail list supporting experimenters, and from analysis of Trouble tickets and summarises it in a monthly operations report which is distributed within the project.

Experience of using the dashboard has led to requests for enhancements to its capabilities, some of which have already been implemented. A particular change, which was implemented in August 2014, was the automatic generation of an e-mail to FLS when an incident is detected. The FLS service only operates for 8 hours a day during working days, whereas the e-mail notification is constantly available. This has led to the generation of significantly more operational data regarding incidents.

This deliverable analyses the operational data collected by FLS, including data from e-mail incident notifications for August and September 2014. It looks at trends within the data and summarises the lessons learned during the first 9 months of operation.

Acronyms and Abbreviations

CET	Central European Time
FLS	First Level Support
SME	Subject Matter Expert
TTS	Trouble Ticket System

Table of Contents

1 Introduction.

Fed4FIRE is a federation of independent test-beds who have all implemented a common set of standards for experiment definition, monitoring and control. This enables experimenters to utilise multiple test-beds, of differing technologies, to create more complex experiments closer to real-life situations. Historically, prior to the federation the test-beds had relatively close contact with their own experimenters. The creation of the federation means that experimenters will not generally have a close relationship with the test-beds that they use. It also introduces operational issues related to the interaction between test-beds as well as the operation of the additional functionality provided by the federation itself.

All of these factors require a new approach to operational support for experimenters. The First Level Support Function (FLS) acts as a single, common point for the detection and management of operational incidents occurring on test-beds within the federation. Using a standard Trouble Ticket System (TTS), it logs these incidents and analyses them on a monthly basis.

This deliverable reviews the experience of operating FLS since January 2014. It analyses the data gathered as an assessment of the operational performance of the federation. On the basis of the lessons learned it details further improvements that can be made to FLS.

2 Background to First Level Support (FLS)

First Level Support (FLS) provides the fault logging tools and resolution processes that underpin the operational support model of the federation. Figure (i) illustrates the positioning of FLS and its interactions with other functions within the federation.

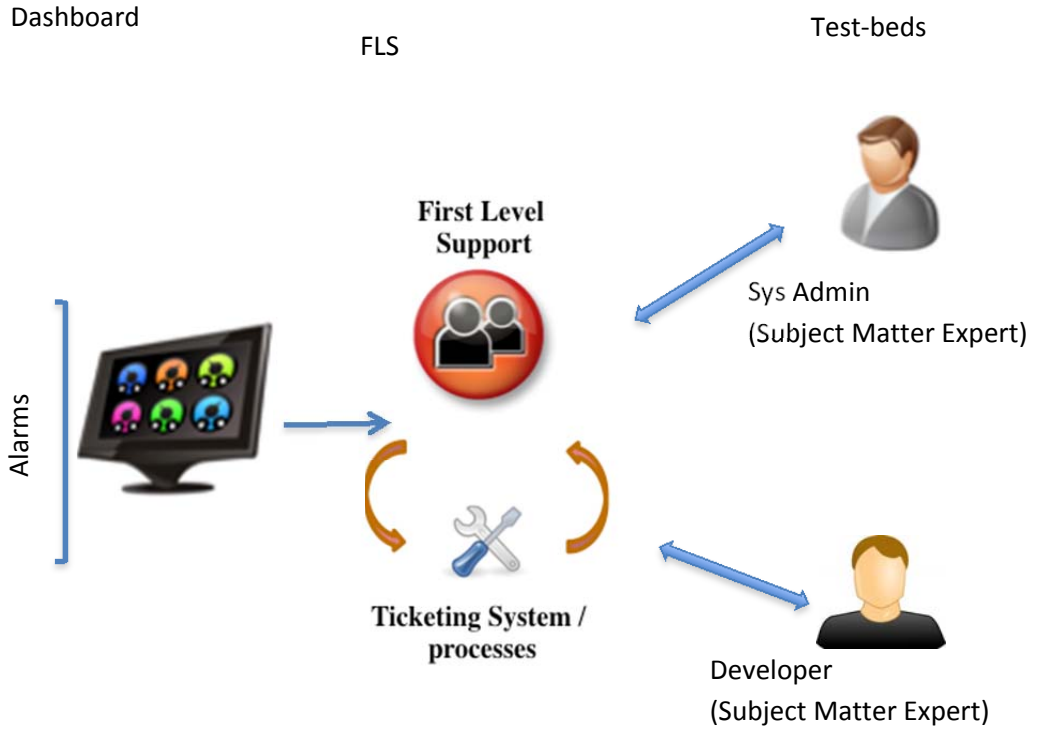

Figure (i) FLS and the Fed4FIRE federation

The following key elements are briefly described:

- **The Dashboard**. This provides a visual indication of the state of health of the test-beds within the federation. It is based on a polling function that regularly tests a number of key elements of each test-bed and displays the results in visual format for FLS operators. Figure (ii) shows a snapshot of the dashboard. The dashboard, which was based on the requirements set out in Deliverable D8.1, provides a Red/Amber/Green indication of alarms status.
- **Trouble Ticket System (TTS).** The TTS is the database where all incidents are recorded. It represents the repository of incidents dealt with by FLS.
- **Standard Operational Processes**. Support operations are driven by a number of practical considerations. In particular, there is a strong need for consistency in the way problems are treated. It is also the case that a given incident will often be handled by more than one individual. Both of these reasons require a very standardised approach to incident management.

- **Subject Matter Experts (SME's)**. FLS staff members are generalists. Their main skill-set is in managing and resolving incidents. They do not have an in-depth knowledge of the functioning of all the test-beds in the federation. In order to deal with test-bed specific issues an expert, familiar with the functioning of that particular test-bed, is needed. These 'Subject Matter Experts' liaise with FLS staff to manage the resolution of incidents. In order to provide a second line of expertise an escalation level of SME is also organised. This is provided by the operational group of the co-ordinating partner. It is referred to as the ***default SME.*** Problems that are not resolved in a timely manner by a test-bed SME are escalated to the default SME.

2.1 The Dashboard Function

The dashboard is a key tool for FLS to observe the behaviour of test-beds within the federation. For each test-bed it shows that status of a common set of key status indicators. Figure (ii) below shows a screen shot of the FLS dashboard. The indicators are:

- **Ping Latency.** This tests the reachability of the test-bed from the Internet. It is expressed as a latency time. A test-bed, which is not reachable, is effectively disconnected from the federation.
- **Get Version Status.** This is the result of a Get version Status call to the Aggregate Manager Software of the test-bed.
- **Free Resources.** This is an indication of the availability of resources for experimentation. As such, it is not necessarily a critical indicator as far as incidents are concerned. There will be occasions when there are no free resources. It is nevertheless an aid in terms of diagnosing problems.
- **Internal Status.** This is an aggregate alarm based on the status of individual, test-bed specific, alarms.

Testbed Name	Ping latency (ms)	GetVersion Status	Free Resources	Internal testbed monitoring status	Last check internal status
BonFIRE	31.17	N/A	N/A	ok	2014-11-26 14:21:13+00
C-Lab	52.63	ok	111	ok	2014-11-26 14:20:06+00
FUSECO	15.78	ok	4	ok	2014-11-26 14:19:04+00
Koren	313.06	ok	3	N/A	N/A
NETMODE	66.32	ok	20	ok	2014-11-26 14:14:48+00
NITOS Broker	72.1	ok	38	ok	2014-11-26 14:25:02+00
NITOS SFAWrap	30.35	ok	121	ok	2014-11-26 14:25:02+00
Norbit	N/A	N/A	N/A	ok	2014-11-26 14:20:15+00
Ofelia (Bristol openflow)	17	ok	69	ok	2014-11-26 14:15:02+00
Ofelia (Bristol vtam)	16.95	ok	2	ok	2014-11-26 14:15:02+00
Ofelia (i2CAT openflow)	16.9	ok	49	ok	2014-11-26 14:15:02+00
Ofelia (i2CAT vtam)	17.86	ok	6	ok	2014-11-26 14:15:02+00
Planetlab Europe	30.3	ok	273	ok	2014-11-26 14:15:04+00
SmartSantander	59.04	ok	0	ok	2014-11-26 14:20:01+00
Virtual Wall 1	0.11	ok	63	N/A	N/A
Virtual Wall 2	0.12	ok	45	ok	2014-11-26 14:19:50+00
Virtual Wall 2 (openflow)	0.63	ok	2	ok	2014-11-26 14:19:50+00
w-iLab.t 2	4.61	ok	7	ok	2014-11-26 14:19:39+00

Calendar with testbed maintenances and reservations

Figure (ii) The FLS dashboard screen shot showing link to Maintenance Calendar

FLS provides service during working hours (9am-5pm CET) on working days. As the activity is supported by DANTE, based in the UK, this translates into 8am-4pm UK time. Outside this period incidents will not be detected. FLS started full operational service in February 2014. The service introduction was preceded by operation, in a pilot mode, for the month January 2014 so that the

Operations staff and Test-bed SME's could gain 'live' experience of operating the service. During this period the FLS staff have observed the dashboard and raised trouble tickets when incidents are detected as a result of changes to the dashboard.

An important issue for the individual test-beds, the federation and FLS is scheduled maintenance. With any IT system there is a need for scheduled maintenance, but it is important that maintenance is visible, in advance, both to experimenters and FLS, otherwise experiments will fail and FLS will detect the consequences of scheduled maintenance as incidents rather than managing this incident as maintenance. A maintenance calendar is therefore associated with the dashboard. It is accessible from a link below the dashboard alarm status display. FLS will receive e-mail notification of maintenance. This leads to the auto-creation creation of a maintenance trouble ticket. FLS will also manage the entry of scheduled maintenance in the calendar.

2.2 Experimenters' List

It was originally planned that FLS would also deal with problems experienced by experimenters using the federation. The FLS operators, however, do not have the detailed understanding or operational knowledge of the test-beds in the federation. Although it might have been possible to acquire this experience, it was decided that specific issues raised by experimenters were best dealt with by the experts from the individual test-beds. As a consequence, operational issues resulting from experimentation have been handled via an Experimenters' email list where experimenters post their issues.

The list is monitored by Subject Matter Experts (SME's) from the participating test-beds who deal with the queries that arise. Although FLS is not actively involved in the resolution of experimenters' issues it, nevertheless, monitors the list and provides a monthly analysis of the issues dealt with. The experimenter's list is analysed in terms of:

- The test-bed or software to which the query relates. Most queries relate to individual test-beds but some relate to software used in the federation and an occasional query relates to multiple test-beds
- The category of query, e.g. whether it is a general question, a difficulty with an actual experiment or some other category.

This is an important input to the federation in terms of its operational performance. Given that, actually, most of the issues raised relate to problems establishing experiments or more general queries about capabilities, it is of relevance to the development of the federation as a whole.

2.3 E-Mail Notification of Incident Reports

The dashboard is the primary source of incident reports for FLS. It provides a visual overview of the functioning of the federation by displaying the status of a set of key functions. There were, however, some limitations to the effectiveness of a dashboard. When multiple alarms occurred in adjacent fields they could be difficult to spot. Transient incidents would leave no lasting record. From August 2014 onwards, the visual indication was supplemented by an e-mail generated when a new incident was detected. This has three effects:

1. The e-mails are generated on a 7*24 hour basis whereas the FLS service only functions between 9am and 5pm CET on working days. This gives rise to a huge increase in the number

of incidents detected. Logically this increase should be a factor of 4. In practice it is closer to 7. An analysis of the incidents reported by e-mails is contained in section 3.4.

2. Each separate change of state from normal to 'incident state' generates an e-mail. A major failure of a test-bed will typically generate multiple incident e-mails. Whereas a human operator will generally be able to correlate these, it is much more difficult for an automatic process to do so.
3. As currently implemented, there is no matching e-mail created when an incident clears. This means that transient faults or 'false positives' will be registered without any real means to detect this.

Because of these limitations, the e-mail notification is currently only used as an adjunct to the visual detection of incidents from the dashboard by FLS. It, nevertheless, generates a very significant amount of operational information about the performance of the federation. This is analysed further in Section 3.4.

3 Operational Experience to Date

The activities of FLS are recorded in a monthly summary report, which is distributed within Fed4FIRE. The report covers the following areas of First Level Support:

1. An analysis of trouble tickets generated during the month by test-bed and by type of incident.
2. A count of tickets where a problem has resolved itself within two working days and no problem resolution been registered.
3. A count of tickets that are escalated to the default SME because of a lack of responsiveness by a test-bed SME.
4. An analysis of activity on the Experimenters' list on a monthly and a cumulative basis.
5. From August 2014 onwards, an analysis of the automatically generated e-mail incident reports.

Points 1, 4 and 5 are fairly self-evident. Points 2 and 3 require further clarification. When a trouble ticket is raised by FLS it is referred to the SME of the relevant test-bed. In some instances, a fault will be rectified before there is a response from the test-bed SME. The trouble ticket will, however, only be closed when such a response is received indicating the diagnosis of the problem. In order to avoid having trouble tickets remaining open for faults that are no longer present, but where no fault diagnosis is received from the test-bed SME, these tickets are automatically closed two working days after they have been opened. The count referred to by Point 2 is a monthly count of these types of tickets and represents an indicator of the service responsiveness of the federation. As explained in Section 2 there is an escalation path from the test-bed SME to an SME function provided by the co-ordinating partner. The count in point 3 registers the number of times the escalation path is used where there has been no initial response from the test-bed SME. It also represents an indicator of the service responsiveness of the federation.

Tickets should be closed with a fault diagnosis. The test-bed SME should be the function for dealing with a fault on a particular test-bed. Capturing the occasions when these processes are not followed provides an indication of how well operational support is functioning within the federation; the lower the number of such tickets the better operational support is functioning. This point is discussed further in section 3.2 below (Operational Discipline)

This deliverable summarises the trends from these reports over the period February to September 2014. A typical monthly report is attached at Appendix A.

3.1 Analysis of Trouble Tickets

A Trouble Ticket is generated by FLS whenever an incident is detected on the dashboard. The same trouble ticket may cover more than one incident if there is a clear relationship between the incidents. As an example, the Virtual Wall at iMinds actually houses a number of test-beds; a network access incident affecting all the test-beds would be covered by a single trouble ticket even though it would generate multiple dashboard alarms. Figure (iii) below shows the count of trouble tickets opened, on a monthly basis. Figure (iv) analyses these tickets further by category of incident.

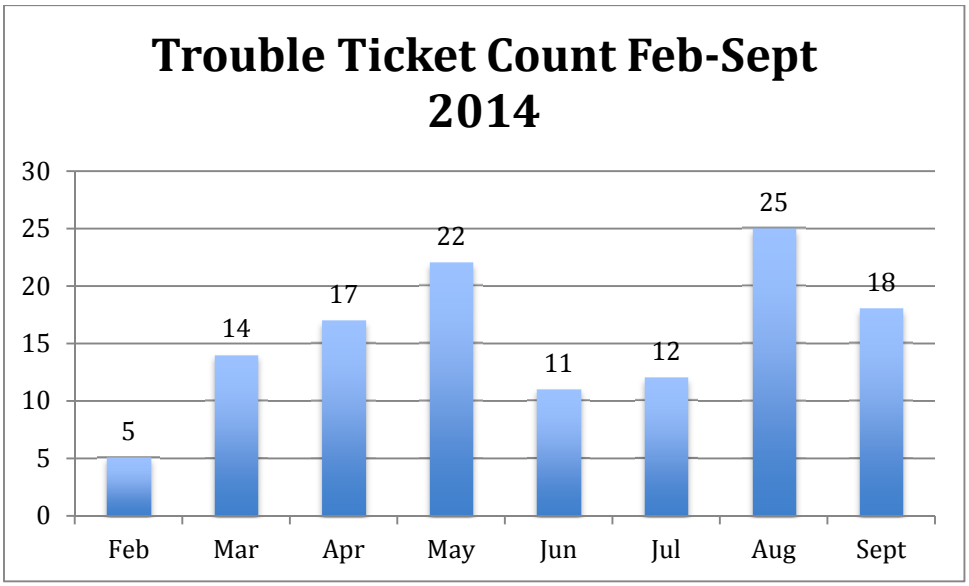

Figure (iii) FLS Trouble Ticket Count February – September 2014

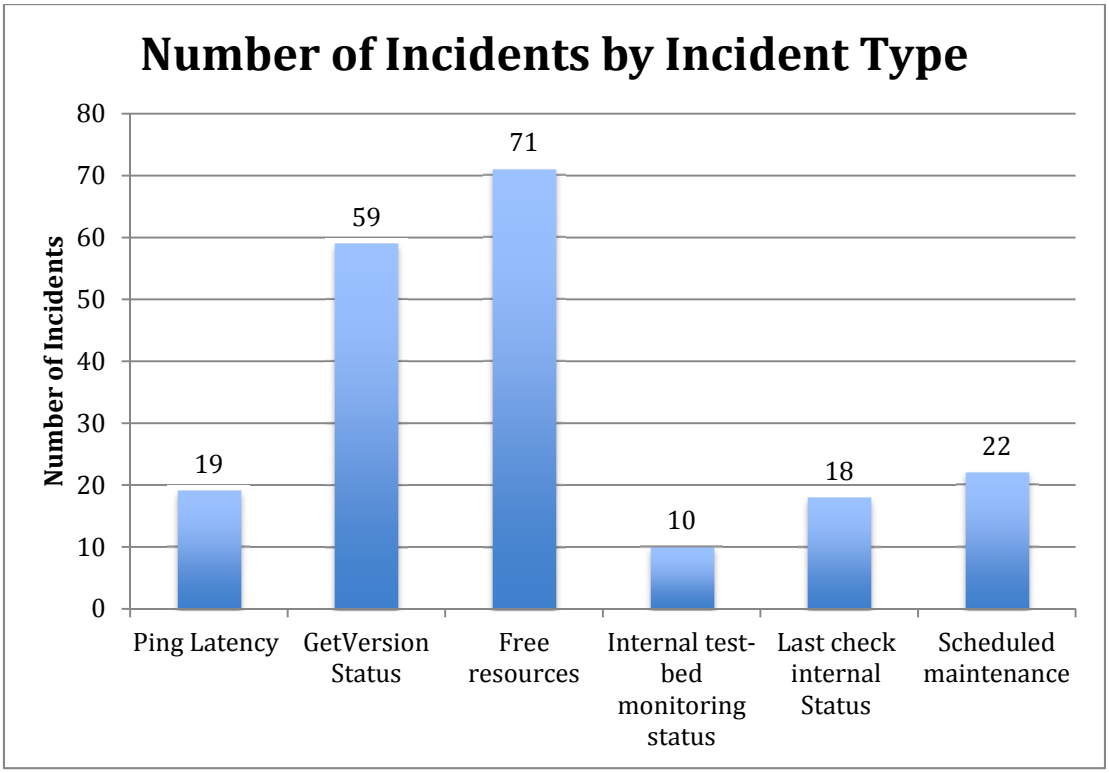

Figure (iv) Analysis of Trouble tickets by type of Incident

Two comments are important. The column free resources may or may not represent an operational issue. It records the state of usage of the test-bed. In isolation it is arguable whether it represents a fault. It is more relevant to experimenters as a reason behind experimental problems. Scheduled maintenance can also not be seen as a fault. As noted in section 2, management of maintenance is an important element of operational management within the federation.

3.2 Operational Discipline

In any federation it is important that common attention to operations applies across the federation. Members of the federation are, typically, not directly accountable to experimenters in a way that a single test-bed would be to its users. FLS measures three indicators which provide a view on operational discipline within the federation. These are the count of tickets where they are closed with an 'incomplete code,' the count of tickets that have been escalated and the count of scheduled maintenance tickets. Unscheduled maintenance can be difficult for FLS to diagnose. It is therefore not generally recorded as such in trouble tickets. Nevertheless, apparent incidents of unscheduled maintenance are publicised in the monthly service reports. Figure (v) below plots the monthly development of these trouble ticket counts. Low counts of Tickets with incomplete code and low counts of escalated tickets are signs of good operational discipline. It is also important that maintenance is scheduled.

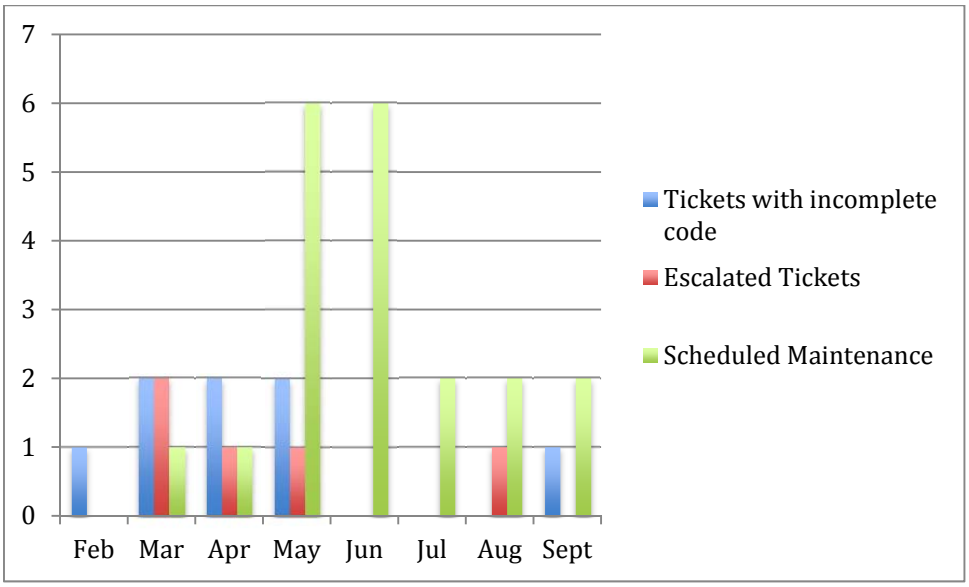

Figure(v) Indicators of Operational Discipline

In general this shows quite high levels of operational discipline within the federation with good attention to reporting maintenance in advance.

3.3 Experimenters' List Activity

As described in Section 2.2, operational issues experienced by experimenters are resolved via an e-mail distribution list that includes experts from the test-beds in the federation. Figures (vi) and (vii) analyse the activity on the experimenters' list since it started in January 2014. Queries on the list are categorized depending on whether they are requests for information (General Query About Capabilities) or issues occurring during the process of an experiment (Experimental Issue). The list is also occasionally used to broadcast information when there is no direct underlying query (Informational). There are a small number of queries about issues relating to the federation (Interaction between test-beds) and about network connections (Network Connectivity). The distribution by type of query is shown in Figure (vi). It can be seen that 90% of the queries relate to

requests for information or specific issues associated with running an experiment. The monthly distribution of queries over the period is shown in Figure (vii).

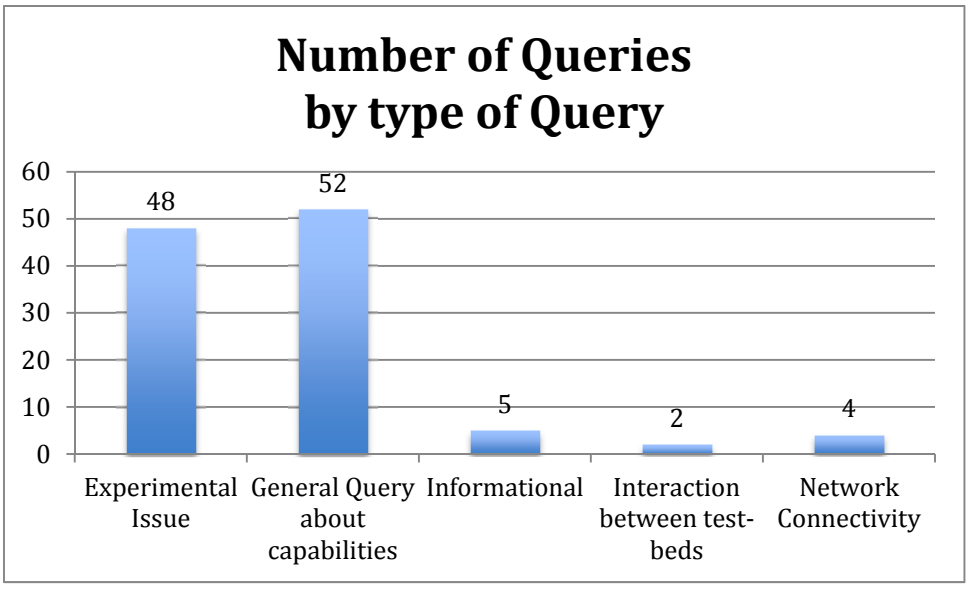

Figure (vi) Analysis of Experimenters' list by type of query.

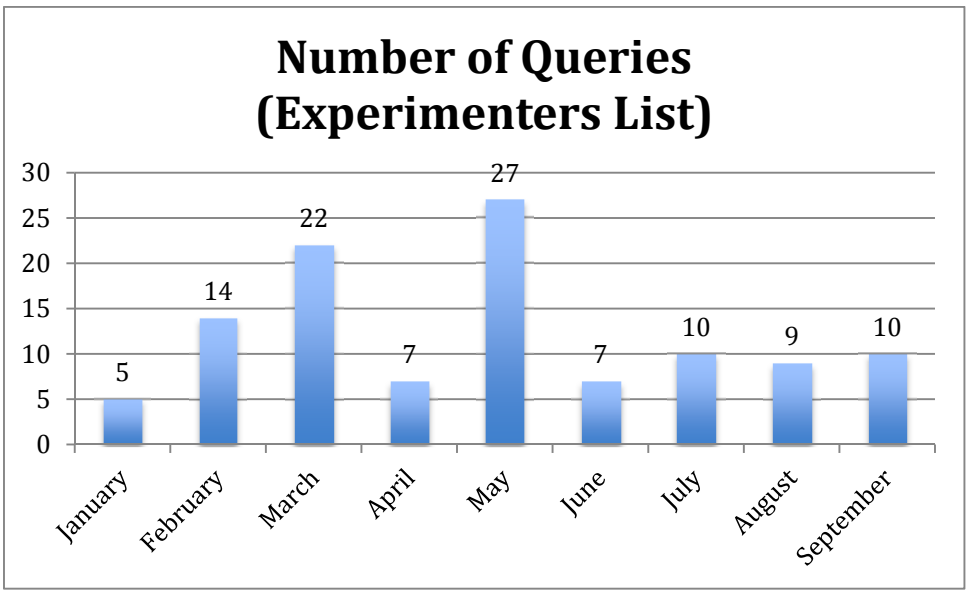

Figure (vii) Monthly Activity Experimenters' List January- September 2014.

The queries generally relate to the functionality of an individual test-bed, although some relate to software packages, particularly jFED (the basic software framework used by the federation) or other issues including, more recently, network connectivity. Table 1 below shows the distribution of queries according to test-bed/type of query and category of query.

In order to give a more detailed insight into the type of queries that are raised, Table 2 below shows the breakdown of activity on the experimenters' list for September 2014, as presented in the monthly operational report.

Test-bed/Query	Experimental Issue	General Query about capabilities	Informational	Interaction between test-beds	Network Connectivity	Grand Total
Bonfire	14	19				33
Bonfire, Amazon				1		1
Emulab Log in	1					1
INRIA	5	1				6
jFED	5	2	3			10
Multiple		1				1
NEPI	2					2
OpenFlow		2				2
Perform LTE		1				1
PlanetLab	5	4				9
Virtual Wall	7	9	1		1	18
Virtual Wall, Bonfire	2	9		1	3	15
W-Ilabt	7	4	1			12
Grand Total	48	52	5	2	4	111

Table 1 Summary of Experimenter's list activity January-September 2014

Date	Query/Test-Bed(s)	Question Type	Detail	Comment
16-Sep-14	W-Ilabt	Experimental Issue	Mobile Node reboot	Unintended reboot
17-Sep-14	Bonfire	Experimental Issue	Elasticity service in bonfire	Operational Issues with elasticity service.
20-Sep-14	W-Ilabt	Experimental Issue	Mobile Node deadlocked	Intermittent issue
23-Sep-14	W-Ilabt	Experimental Issue	Robot Control failure	Reset
23-Sep-14	W-Ilabt	Experimental Issue	Robot Control failure	Debugged
24-Sep-14	jFED	General Query about capabilities	Establish network connection	GÉANT VLAN transition
24-Sep-14	jFED	Experimental Issue	Slice Sharing	
24-Sep-14	jFED	Experimental Issue	Network Connection issue	
29-Sep-14	jFED	Experimental Issue	Network Connection issue	Addressing - VLAN, Port and IP issues

29-Sep-14	Virtual Wall	Experimental Issue	Unable to create VM's	Potential bug

Table 2 Experimenters' List Activity September 2014.

3.4 Analysis of e-Mail Notification of Incidents

E-mail notification of incidents has been operational since August 2014. As noted in section 2.3 e-mails are only generated on incident detection so they cannot be used to automate the process of Trouble Ticket management. The data gathered from analysing the e-mail notifications is nevertheless a useful addition to understanding the operational performance of the federation. Data has only been gathered for two months. This is not yet sufficient to be able to detect trends. Analysis does however show some interesting points

3.4.1 Distribution of Incidents Over Time

The automatic identification of alerts by e-mail has led to a very significant increase in the number of incidents detected. During August and September there were 592 incidents (274 in August and 318 in September). These figures include 'Free resources' alerts. It is arguable whether 'Free resources' is, itself, an incident that would be regarded by FLS as a fault. Taken in isolation, it is an indication of test-bed loading but it can also be triggered by a fault. Eliminating 'Free Resources' still leaves 308 incidents for the two months (130 in August and 178 in September). This compares with a typical, monthly trouble ticket count of 20. A trouble ticket may cover more than a single incident report. FLS operators will correlate multiple incidents and generate a single trouble ticket. There is still a significant difference between the count of trouble tickets and the number of incidents reported automatically by e-mail.

FLS operates for a limited period of the week, Mondays to Fridays between 9am CET and 5pm CET. The boundaries are not quite as rigid as this implies. Incidents occurring prior to the start of FLS observation and still present during operational hours will be detected. The FLS operational period represents approximately 25% of the total week. Assuming that incidents are uniformly distributed in time this would suggest that a similar proportion of incidents would occur during the FLS operational time. Figure (viii) shows that, in fact, rather more incidents occur during the operational period (46% of total incidents as opposed to a predicted 25%).

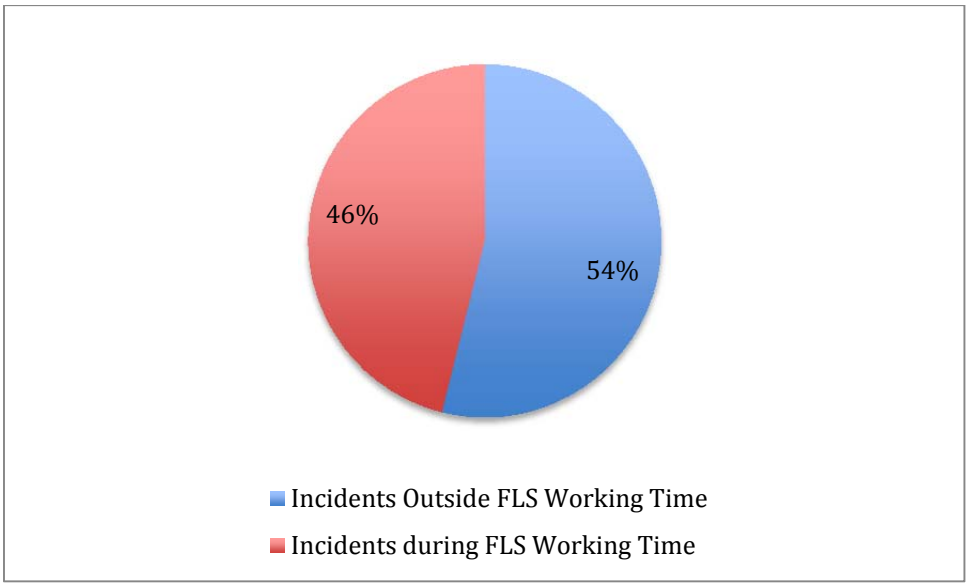

Figure (viii) Distribution of e-mail incidents over time.

This is an aggregate analysis. The distribution of incidents over time, per-test bed varies quite considerably. Figure (ix) below shows the distribution on a per test-bed basis. It should be noted that there were no incidents recorded for Bonfire in the August- September period.

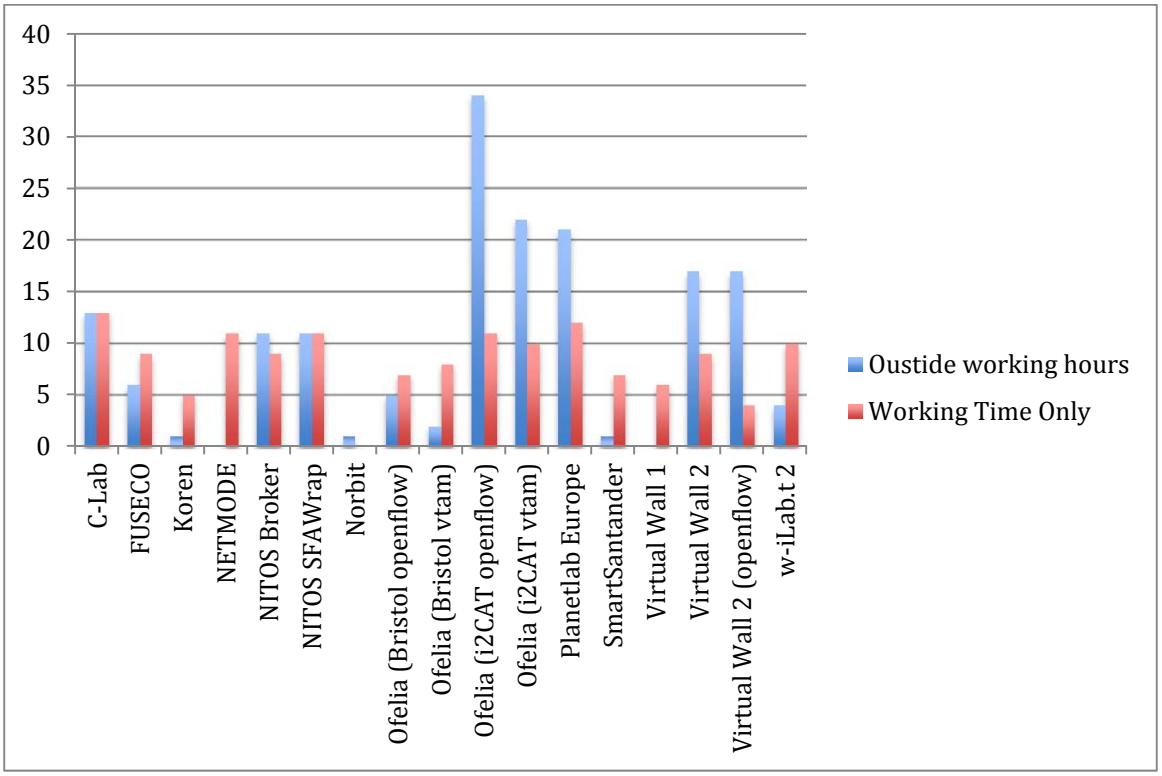

Figure (xi) e-Mail incidents by test-bed over Time

This is further analysed in table 3, which shows the distribution per test-bed according to time of day and day of week.

	Saturday/Sunday	Monday-Friday	

Test-bed		Inside FLS hours	Outside FLS hours	All Incidents
C-Lab	1	13	12	26
FUSECO		9	6	15
Koren		5	1	6
NETMODE		11		11
NITOS Broker	2	9	9	20
NITOS SFAWrap	6	11	5	22
Norbit			1	1
Ofelia (Bristol openflow)	1	7	4	12
Ofelia (Bristol vtam)		8	2	10
Ofelia (i2CAT openflow)	17	11	17	45
Ofelia (i2CAT vtam)	9	10	13	32
Planetlab Europe	8	12	13	33
SmartSantander		7	1	8
Virtual Wall 1		6		6
Virtual Wall 2	11	9	6	26
Virtual Wall 2 (openflow)	9	4	8	21
w-iLab.t 2		10	4	14
Total	64	142	102	308
% of Incidents	21%	46%	33%	
% of week	28%	24%	48%	

Table 3 Distribution of Incidents per test-bed over time

3.4.2 Distribution of Incidents by Incident Type.

The analysis in Section 3.4.1 suggests that there is quite a considerable variation in operational practice among the test-beds in the federation. Another relevant set of data to consider is the types of incident being detected. Table 4 shows this by test-bed and by type of incident, excluding 'Free Resources' for the reasons indicated above.

Test-bed	GetVersion Status	Internal testbed monitoring status	Last Check Internal status	Ping Latency	Total
C-Lab	10			16	26
FUSECO	6	1	1	7	15
Koren	5			1	6
NETMODE	6			5	11
NITOS Broker	7		1	12	20

NITOS SFAWrap	15		1	6	22
Norbit		1			1
Ofelia (Bristol openflow)	7			5	12
Ofelia (Bristol vtam)	6			4	10
Ofelia (i2CAT openflow)	40	1	1	3	45
Ofelia (i2CAT vtam)	28	1	1	2	32
Planetlab Europe	11	15	1	6	33
SmartSantander	5			3	8
Virtual Wall 1	4			2	6
Virtual Wall 2	6	18		2	26
Virtual Wall 2 (openflow)	2	18		1	21
w-iLab.t 2	5	2		7	14
Total	163	57	6	82	308

Table 4 Distribution of Incidents by incident type by test-bed

Again there is no discernable pattern in the distribution of incidents recorded, possibly reflecting the different technologies in use in the different test-beds but possibly also reflecting differences in operational practice.

The final comparison to be considered is the relationship between incidents detected via the dashboard by FLS operators and those reported via e-mail. Figure (x) below show this comparison.

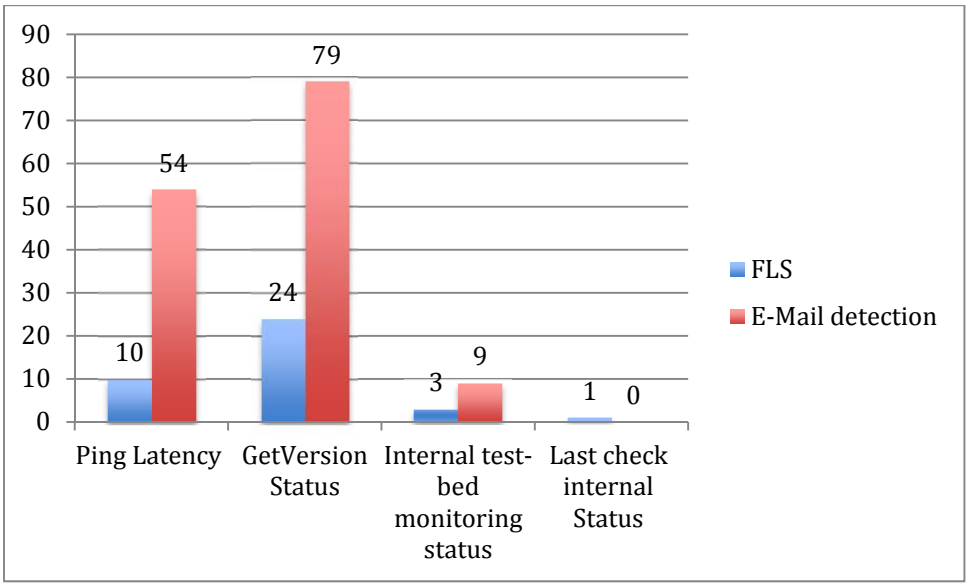

Figure (x) Comparison between FLS records and e-mail incident notifications

The data in figure (x) is taken from the trouble tickets recorded by FLS compared with the e-mails incident reports received during FLS operational hours. It shows that there are significant differences. There are several potential reasons for this:

- Once a trouble ticket has been opened by FLS, it is not closed until it is acknowledged by the appropriate SME. A second instance of the same fault would not normally give rise to the opening of a new ticket whilst there is an existing ticket covering the fault.
- E-mails are only currently generated when an incident is detected. An intermittent fault will, potentially, cause multiple e-mails to be generated. The status of the dashboard will remain at red and these additional incident reports will be missed. Without a matching 'Incident closed' e-mail it is impossible to know the duration of incidents and determine whether they are transient.
- Certain test-beds, notably Ofelia Bristol, Ofelia I2CAT and Virtual Wall have multiple test-bed instances as part of their capability. An overall failure in any of these test-beds will give rise to multiple e-mail incidents, one for each component part. An FLS operator will correlate these dashboard notifications and typically create a single incident report.

4 Lessons Learned

4.1 Dealing with Experimenters' issues.

It was originally planned that FLS would deal with problems raised by experimenters using the federation. The FLS operators, however, do not have the detailed understanding or operational knowledge of the test-beds in the federation. Although it would have been possible to acquire this experience, given the time and resources available, it was decided that specific issues raised by experimenters were best dealt with by the experts in the individual test-beds. As a consequence, operational issues resulting from experimentation have been handled via an Experimenters' list where experimenters post their issues. The list is monitored by Subject Matter Experts (SME's) from the participating test-beds who deal with the queries that arise. This has functioned reasonably effectively, although the level of activity is not high. It has the advantage that experimenters get direct answers from experts. It does however mean that there is no central logging of issues.

4.2 The need for operational discipline

Prior to the creation of the federation each test-bed operated its own support functions. Since there was quite a direct relationship between the test-bed and its users, this approach could be quite informal Deliverable D8.1 describes the arrangements that were in place prior to the start of Fed4FIRE. In a federation, a test-bed will not generally have any direct relationship with experimenters. This implies that there is a need for much greater formality in the operational management of the federation and consequently in the approach that individual test-beds take to operational management at least as far as the federation is concerned. There are two areas where this is important:

- Maintenance.
- Availability of expertise to resolve problems.

As described in Section 3.2 these measures of operational discipline (Unscheduled/scheduled maintenance and responsiveness of SME's) are logged on a monthly basis. The results have been generally positive.

4.3 Dashboard Functionality

A number of Dashboard Enhancements have been implemented. The original dashboard provided a visual indication of the operational status of each test-bed in the federation. A screen shot of the original dashboard is shown at Figure (ii) above. Practical experience of using the dashboard showed several limitations:

- Identification of incidents relies on an FLS operator recognizing a change in the state of the dashboard. If there is a short-term change of status and it is not observed a potential incident will be missed.
- When there are a number of incidents displayed in Red on the dashboard, and a new incident occurs which causes a cell adjacent to the existing red cells to turn red, it can be difficult to spot this change.
- The dashboard is a passive real-time description of the status of the various functions of the dashboard. It has no memory and gives only a visual indication of change of state.

In August 2014 additional features were introduced to address some of the practical concerns described above. These included:

- Flashing audible alerts introduced when an event is triggered. Until an FLS operator acknowledges an alarm, the alarm flash with an audio alert. This feature enabled the FLS operator to spot and react to any change in state of the dashboard immediately.
- Email alerts introduced to compliment the visual indication of alarms. While it enabled logging the transient faults it introduced other practical concerns as described in detail in section 2.3 and 3.4.

There remain a number of potential improvements to the dashboard that could still be made. Currently there is no 'comment' capability associated with the dashboard to allow operators to exchange informal comments associated with the status of a test-bed. These include points such as hand-over notes, informal observations etc. As the operators work in shifts, and more than one operator may deal with an incident, the easy availability of incident-specific notes is an important requirement. Although the TTS provides a formal repository for information relating to incidents, the dashboard is the principal tool used by FLS operators. A request for this enhancement has been submitted to the development team.

4.4 E-mail Notification of Incidents

E-mail notification of incidents offers the possibility of automating the process of Trouble Ticket management. By using a standard format e-mail it is possible to automatically generate Trouble Tickets. Currently e-mails are only generated when an incident is detected. There is no matching e-mail generated when a current incident is resolved. As e-mail notifications are generated throughout the entire week (i.e. 7*24) it would be possible to extend the coverage period of FLS but to achieve two additional functions would be necessary. Firstly, there would need to be a matching closure e-mail generated when an incident is resolved. Secondly all the emails would have to adhere to a standard format readable by the TTS. Implementing these functions would certainly enhance FLS. They have been requested from the development team. When implemented, it would then be possible to determine the usefulness of auto ticket generation.

5 Appendix A: Specimen Monthly (FLS) – August Monthly Report

5.1 Executive summary

This was the seventh month of operation for FLS. During it 25 trouble tickets were opened. This was a much higher level of activity than in recent months. They are analysed in Tables 1-3. In contrast, there was relatively little activity on the experimenters' list with only 9 issues. An analysis of the activity on the Experimenters' list, in August, is presented at Table 4 and a cumulative summary analysis of the first seven months activity on the experimenters' list is available as Table 5.

Changes have been made to the FLS dashboard, which affect the way in which incidents can be recorded. This new version of the dashboard has been used, in 'shadow' mode, in August. The figures shown in this report are based essentially on observations of the old dashboard. The differences between the two approaches are potentially dramatic. The numbers of incidents detected, when using the new dashboard in shadow mode, has increased by a factor of 15, compared with the average for the previous six months. This raises a number of operational questions that will be discussed at the Fed4FIRE plenary.

5.2 FLS Report August

The FLS service began working in operational mode at the beginning of February. August is its seventh full month of operation. During the month of August, 25 trouble tickets were opened. There were two periods of scheduled maintenance during the month. No unscheduled maintenance was detected. As noted in the July report, requested upgrades to the dashboard have been delivered by iMinds. The updates generate an automatic e-mail when an incident occurs. As a consequence, incidents which occur outside the current FLS service window, as well as transient incidents, are more likely to be detected. We have tested the revised dashboard in August but in shadow mode (i.e. we have not used the information to generate trouble tickets). The new dashboard has resulted in hugely more incidents being detected (Nearly 280 in August alone). These changes potentially alter the way in which faults are recognized and recorded. The data contained in this report is based on the original dashboard, although certain, very obvious, issues arising from the new dashboard data have been raised as trouble tickets and appear in the analysis. These are covered in Notes 4, 5 and 6.

Three graphs are contained in the report to continue the trend analysis. The presentation of these was slightly modified in the July report to show a rolling six-month view. Comments are invited on how this data can best be presented looking ahead.

Tables 1, 2 and 3 below summarise the position by test-bed and by type of alarm.

5.3 Trouble Ticket Analysis.

The following tables summarise the FLS activity by Trouble Ticket count, Trouble Ticket per test-bed and ticket cause.

Total number of tickets	25
Total number of tickets resolved	19 Note 1
Number of Dashboard tickets opened	23
Number of Scheduled maintenance tickets	2
Number of Incident reported via email	0
Number tickets closed with the Resolution code "Incomplete"(i.e. no response from the default SME but alarms cleared)	0
Number of tickets escalated to Default SME	1
Number of tickets on Hold	4 Note 2

Note 1: seventeen tickets from August plus two tickets from previous months were resolved.

Note 2: four tickets regarding the dashboard improvement request are put on hold. The new features were implemented by iMinds on a test instance of the dashboard. The FLS team have tested the features and few minor changes have been requested from iMinds.

Table 1 Tickets statistics for the period 01/08/2014 to 31/08/2014

Test-bed	Number of incidents based on Dashboard	Number of incidents reported via Email	Number of bugs reported via Email	Number of Service and Information request	Number of Scheduled Maintenance
Bonfire					1
C-Lab	2				
Fuseco	3				
Koren	1				
NETMODE	1				
NITOS Broker	2				
NITOS SFAWrap	2				
Norbit					
Ofelia (Bristol OpenFlow)	1				
Ofelia (Bristol vtam)	2				
Ofelia (i2CAT OpenFlow)	4 Note 3				
Ofelia (i2CAT vtam)					

Planetlab Europe	2				1
SmartSantander					
Virtual Wall 1					
Virtual Wall 2					
Virtual Wall 2 (OpenFlow)	1				
w-iLabt.t 2	2				
Total	23	0	0	0	2

Note 3: ticket FEDFIRE-139 and FEDFIRE-146 covered incident impacting both Ofelia i2CAT openflow and vtam testbeds. FEDFIRE-154 covered incident affecting Ofelia i2CAT vtam and FEDFIRE-160 covered incident affecting Ofelia i2CAT openflow.

Table 2 Breakdown of tickets based on the Test-bed

This table illustrates the type of alarms and the total number of dashboard events detected per testbed. Note that the total number of dashboard incidents detected may not be in line with total number of dashboard tickets, as normally multiple alarms for a testbed will be tracked from one ticket.

Test-bed	Ping Latency	GetVersion Status	Free resources	Internal test-bed monitoring status	Last check internal Status	Total number of Incidents
Bonfire						
C-Lab	1[Note 4]	1	1			3
Fuseco		1	1	1	1	4
Koren		1	1			2
NETMODE	1					1
NITOS Broker	1	1	1			3
NITOS SFAWrap	1	2	2			5
Norbit						
Ofelia (BristolOpenflow)			1			1
Ofelia (Bristol		1	2			3

vtam)						
Ofelia (i2CAT Openflow)		3[Note 5]	2			5
Ofelia (i2CAT vtam)		3[Note 5]	2			5
Planetlab Europe	1	1	1	1[Note 6]		4
SmartSantander						
Virtual Wall						
Virtual Wall 2						
Virtual Wall 2 (openflow)			1			1
w-iLabt.t 2	2	1	1			4
Total	7	15	16	2	1	41

Note 4: ticket FEDFIRE-140 covered multiple intermittent alarms on C-Lab for "Ping latency". (New test Dashboard)

Note 5: ticket FEDFIRE-139 covered multiple intermittent alarms on Ofelia i2cat (openflow and vtam) for "GetVersion Status". (New test Dashboard)

Note 6: ticket FEDFIRE-141 covered multiple intermittent alarms on Planetlab for "Internal status." (New test Dashboard)

Table 3 Dashboard incident break down

5.4 Trend Analysis

Three graphs have been added to the report to illustrate the trend analysis of the last 7 months.

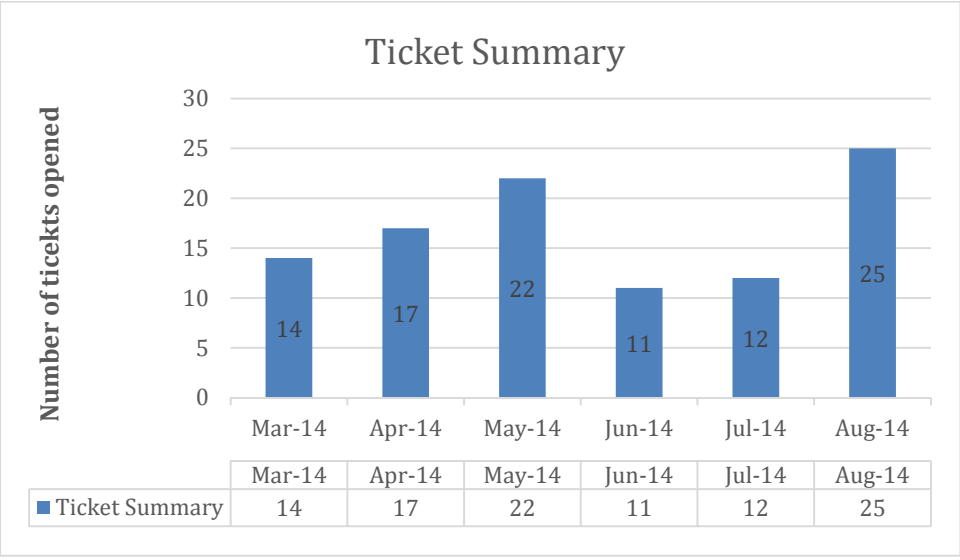

Figure 1 Number of tickets opened by FLS on a monthly basis.

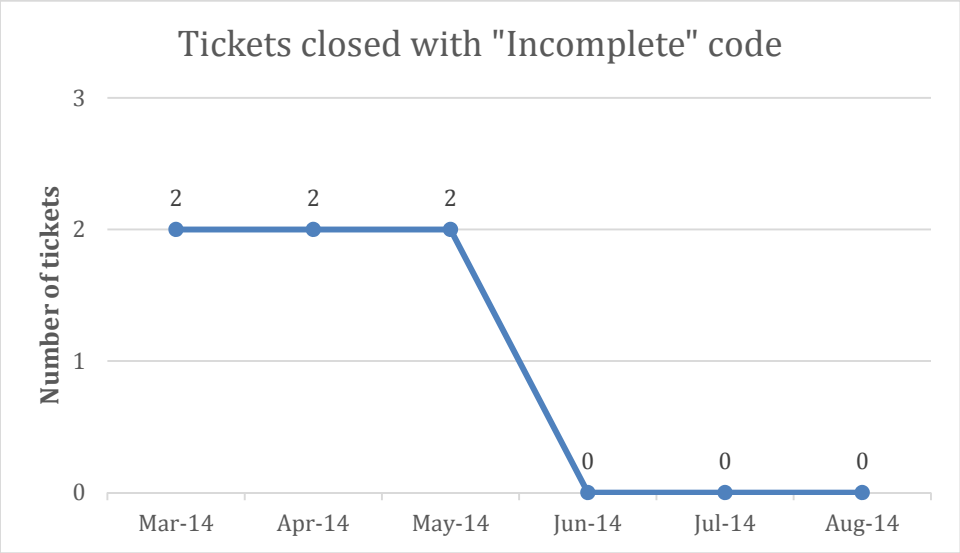

Figure 2 - Tickets closed with incomplete code

For a dashboard ticket assigned to SME, where there is no response for 2 working days and if the alarm has cleared, the ticket will be closed with “Incomplete” code. Figure 2 indicates the number of tickets closed per month as a result of no response from SME.

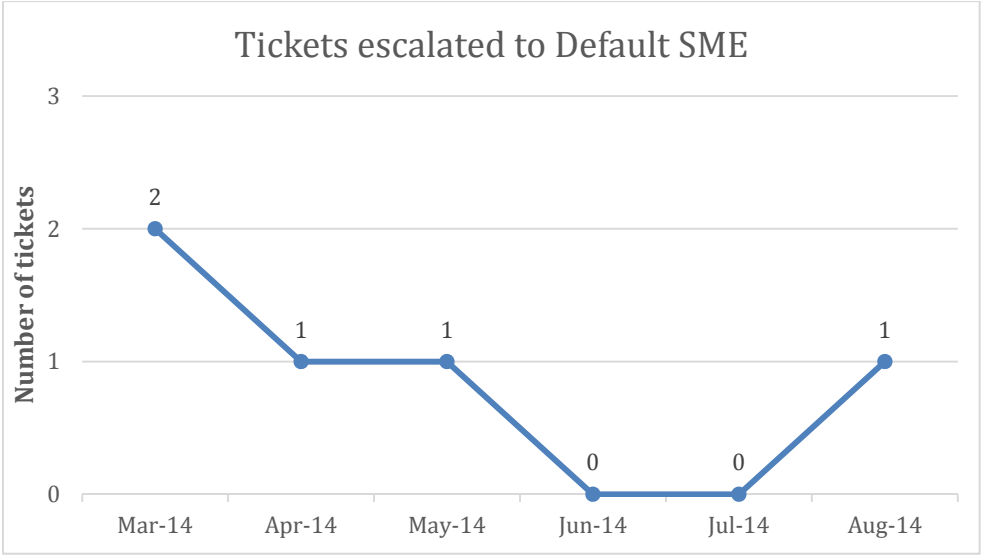

Figure 3 – Tickets Escalated to Default SME.

For an active dashboard alarm or an Incident, if a prompt response is not received from the corresponding test-bed SME, the ticket is escalated to "Default SME" (i.e. iMinds OPS). Figure 3 indicates the number of escalated tickets.

Date	Query/Test-Bed(s)	Question Type	Detail	Comment
06-Aug-14	OpenFlow	General Query about capabilities	Use of OpenFlow controller	Interaction with OpenFlow test-beds
06-Aug-14	Virtual Wall	Experimental Issue	IP addressing	Daemon Failure
08-Aug-14	OpenFlow	General Query about capabilities	Use of OpenFlow controller	Interaction with OpenFlow test-beds
12-Aug-14	Bonfire	General Query about capabilities	External Access	
16-Aug-14	Bonfire	General Query about capabilities	Storage of results	
18-Aug-14	INRIA	General Query about capabilities	Storage of results	
25-Aug-14	W-Ilabt	General Query about capabilities	Issue of interference in mobile space	Need to reserve all relevant nodes to ensure control of interference
26-Aug-14	W-Ilabt	General Query about capabilities	Deployment of other Wireless Cards	
27-Aug-14	W-Ilabt	Experimental Issue	Expiry of reservation	

Table 4 Experimenters list activity breakdown - August

Query/Test-bed(s)	Experimental Issue	General Query about capabilities	Informational	Interaction between test-beds	Network Connectivity	Grand Total
Bonfire	13	19				32
Bonfire, Amazon				1		1
Emulab Log in	1					1
Experiment	2					2
INRIA	5	1				6
jFED	2	1	3			6
Multiple		1				1
NEPI	2					2
OpenFlow		2				2
Perform LTE		1				1
PlanetLab	4	4				8
Virtual Wall	5	9	1		1	16
Virtual Wall, Bonfire	2	9		1	3	15
W-Ilabt	3	4	1			8
Grand Total	39	51	5	2	4	92

Table 5 Cumulative analysis of Experimenters' list queries (January-August)