SDC
STORAGE DEVELOPER CONFERENCE
SNIA ■ SANTA CLARA, 2014

Customer-Oriented Storage Performance Management

Dany Felzenszwalbe
Intel

Legal Notices

This presentation is for informational purposes only. INTEL MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.

About the Presenter

Felsenschwalbe

German to English translation

cliff swallow

Dany Felzenszwalbe

- Joined Intel in 2001
- Working in PDIT / EC
  - **P**roduct **D**evelopment **IT**
  - **E**ngineering **C**omputing
- Primary focus is NAS/NFS data and storage solutions

Agenda

- Intel Design NAS environment overview
- Performance monitoring challenges
- Performance management overview
- Reactive performance management
- Proactive performance management
- Lessons learnt
- Call for Action

Intel's Design NAS Environment

43 Petabytes of NFS capacity

- On > 1000 fileservers
- Multiple vendors / platforms
- In 44 locations

91% in the largest 10 data centers

- 58% in the largest 4 data centers
- From 1 to 190 fileservers in a site

More than 1000 projects

- About 30,000 Users (aka Customers)
- 20-30 million batch jobs running weekly

… and "everything" has to run "FAST"… !

Motivation

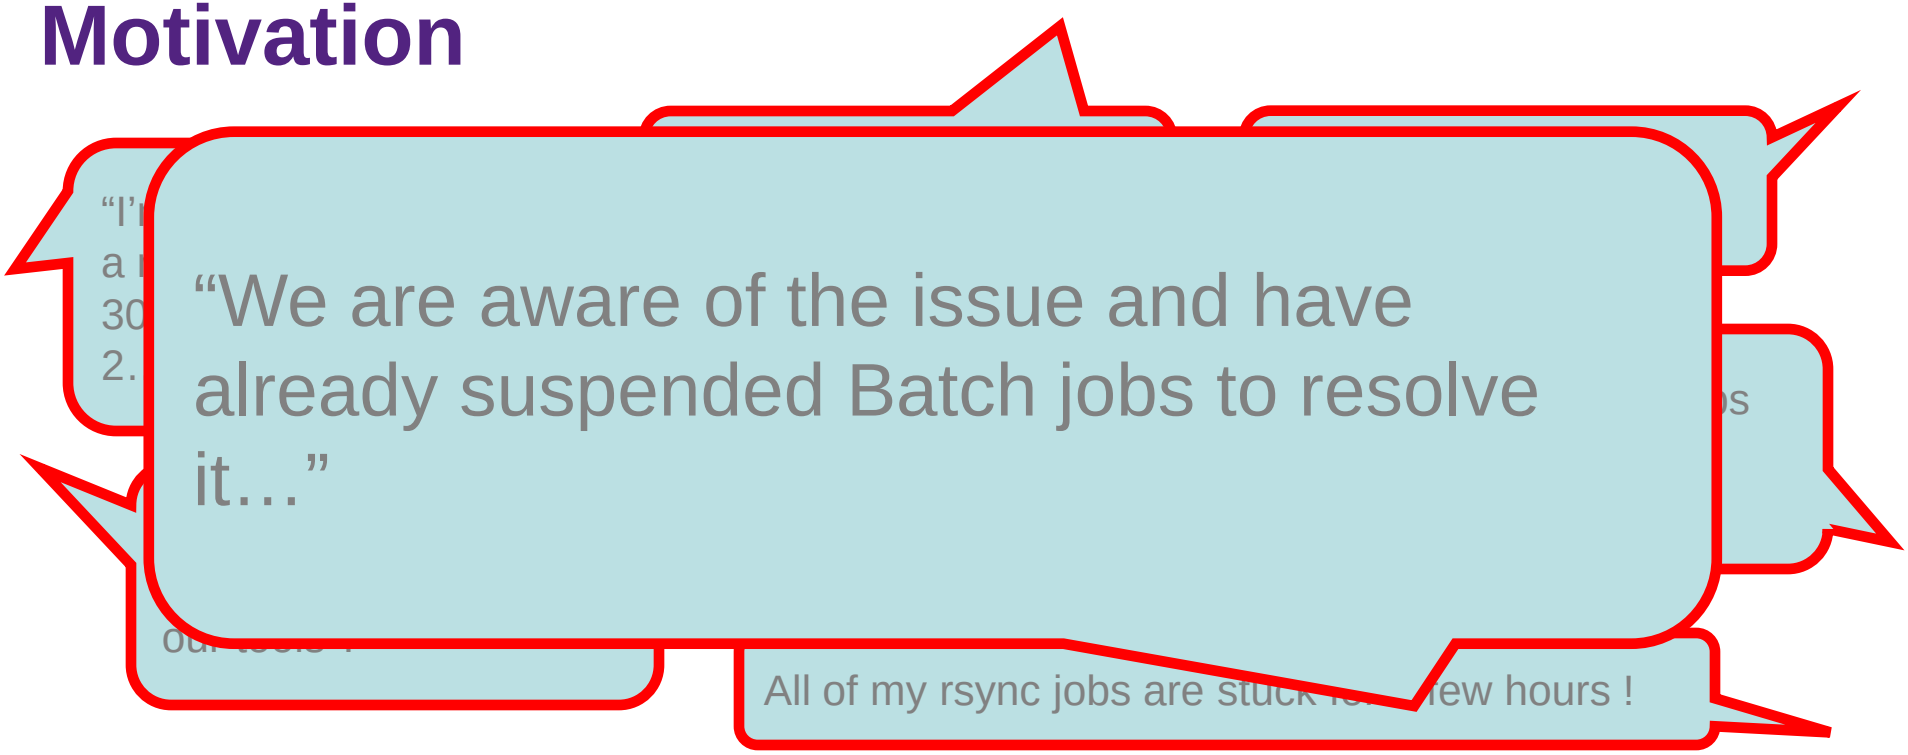

We want to be as proactive as possible in identifying and solving NFS performance issues

NFS Performance

- 4 performance (and cost) tiers are defined for NFS
- The design teams/activities are expected to use the appropriate tiers
- When performance is “bad” the tier “doesn’t matter”
- Understanding when it happens is challenging
- Severe NFS performance issues could impact a project’s TTM

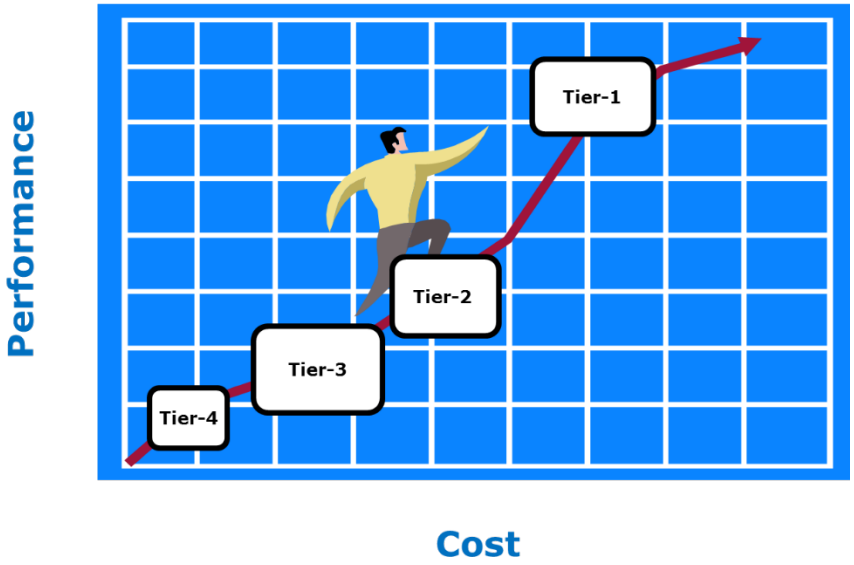

- Until 2007 we relied on customers as our Monitoring

Customer-Oriented Storage Performance Management

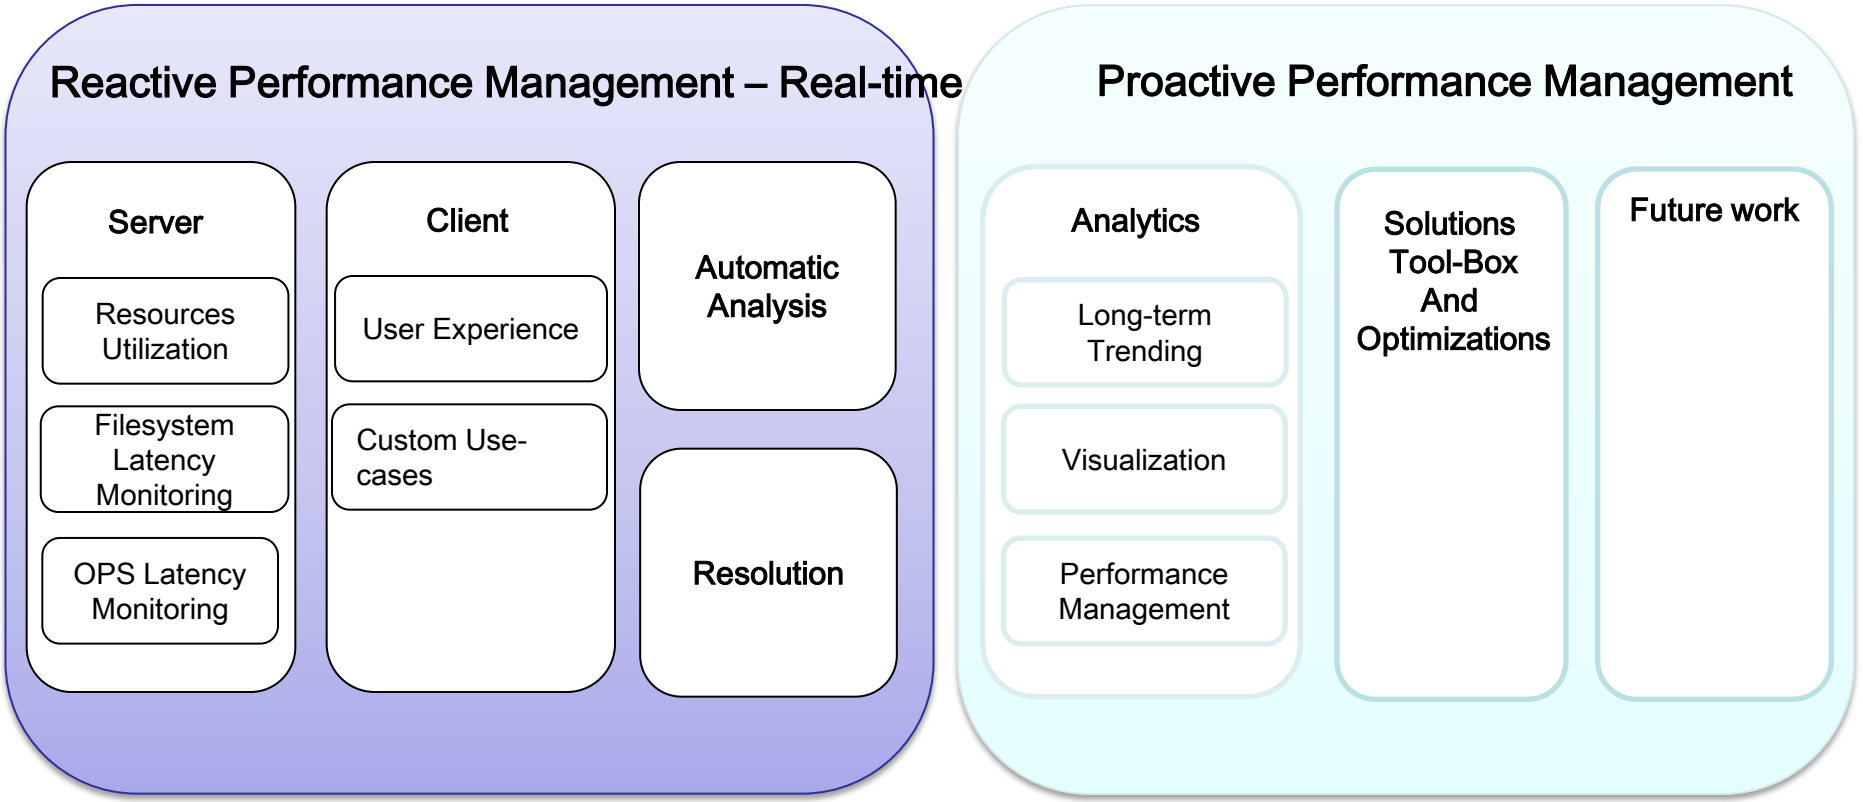

Resources Utilization Monitoring

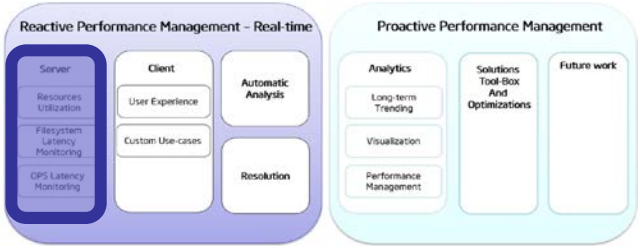

- Monitoring the fileservers' CPU utilization, disk-drives utilization, network collisions and such
- Used Cricket (Open Source from SourceForge) for tracking and as an alerting mechanism
- No correlation between high utilization and users reporting performance issues

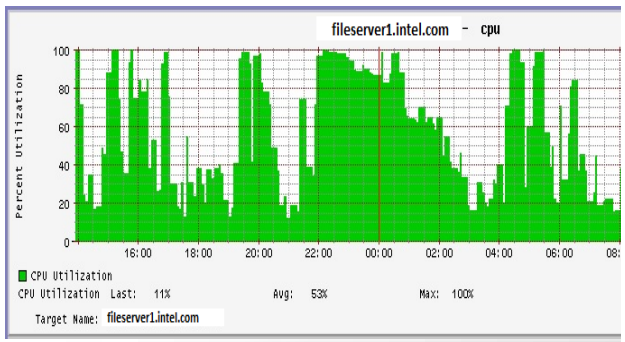

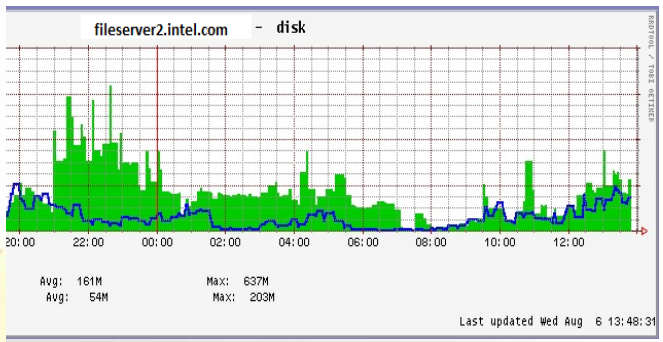

Client-Side User-Experience Monitoring

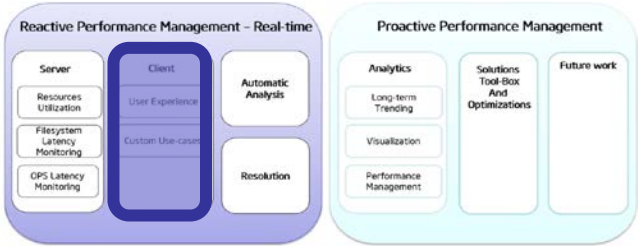

- Simulate user activity to identify problems
  - Slow mounts and reads
- Using home grown monitoring framework
  - 2 clients
  - Different subnets
  - Duration matters
  - Time of day
  - Server type/tier

Threshold is crossed

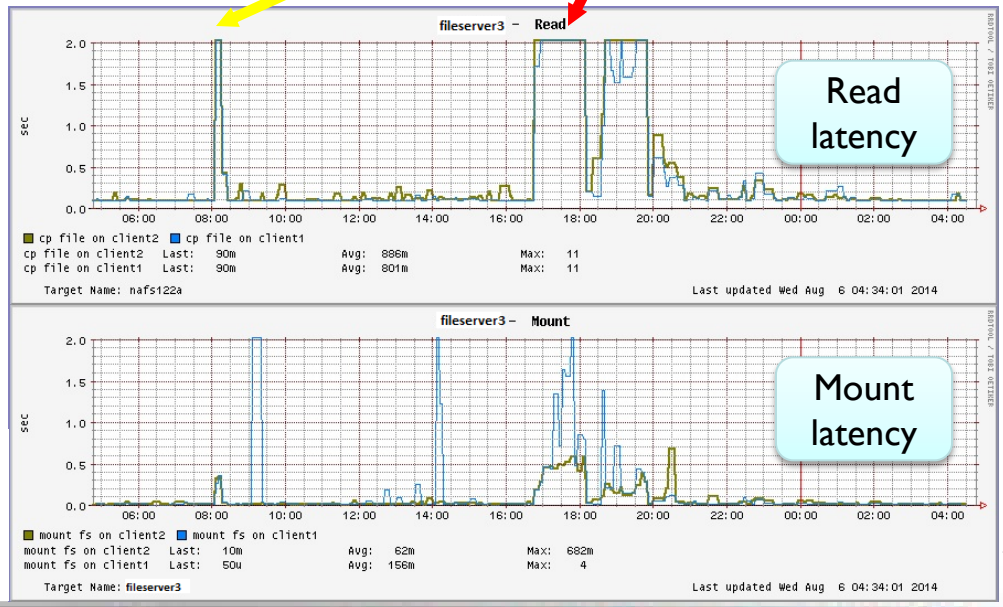

HARMLESS – event is 90 seconds long

CRICITICAL – event is 300 seconds long

FATAL – event is 600 seconds long

Custom Use-Case Performance Monitoring

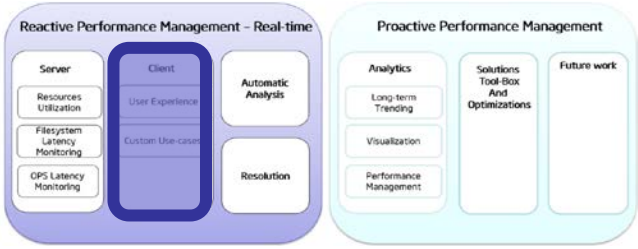

- In spite of the successful monitoring capabilities some undetected events remained painful
- Still needed to identify problems before customers
- Cloning with GIT and a similar tool were chosen and monitors were created

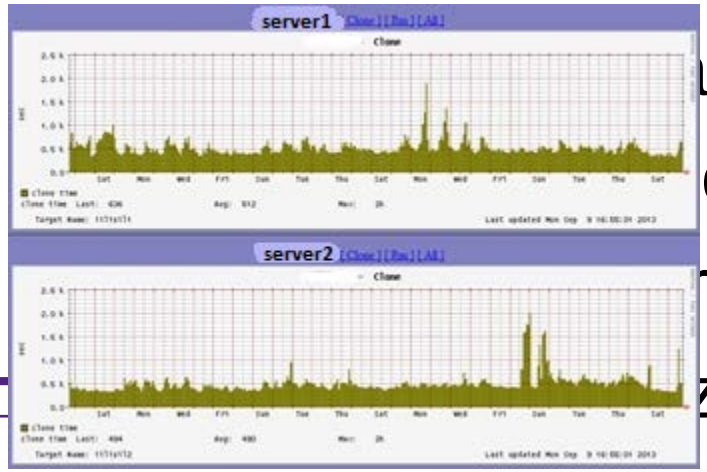

- ...ny use-cases
- ...custom...
- ...
- ...ze and ease implementation

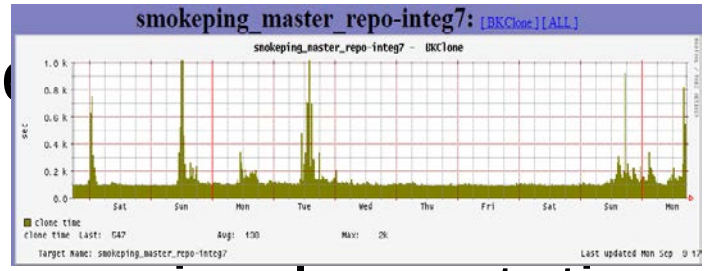

Latencies Monitoring

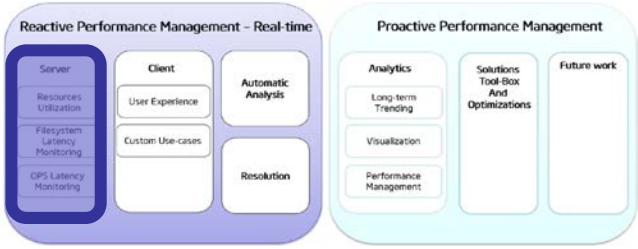

- Latency is a key metric that should provide visibility into customer impact
- Fileserver reported latencies should correlate with the user-experience monitoring and should also match other performance issues (“misses”)
- Different platforms offer different capabilities
- We have been experimenting with “disk” latencies and NFS operations latencies

NFS operations latencies

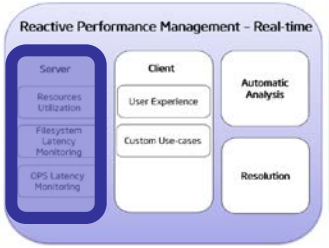

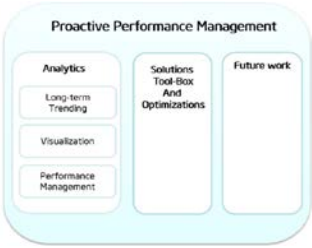

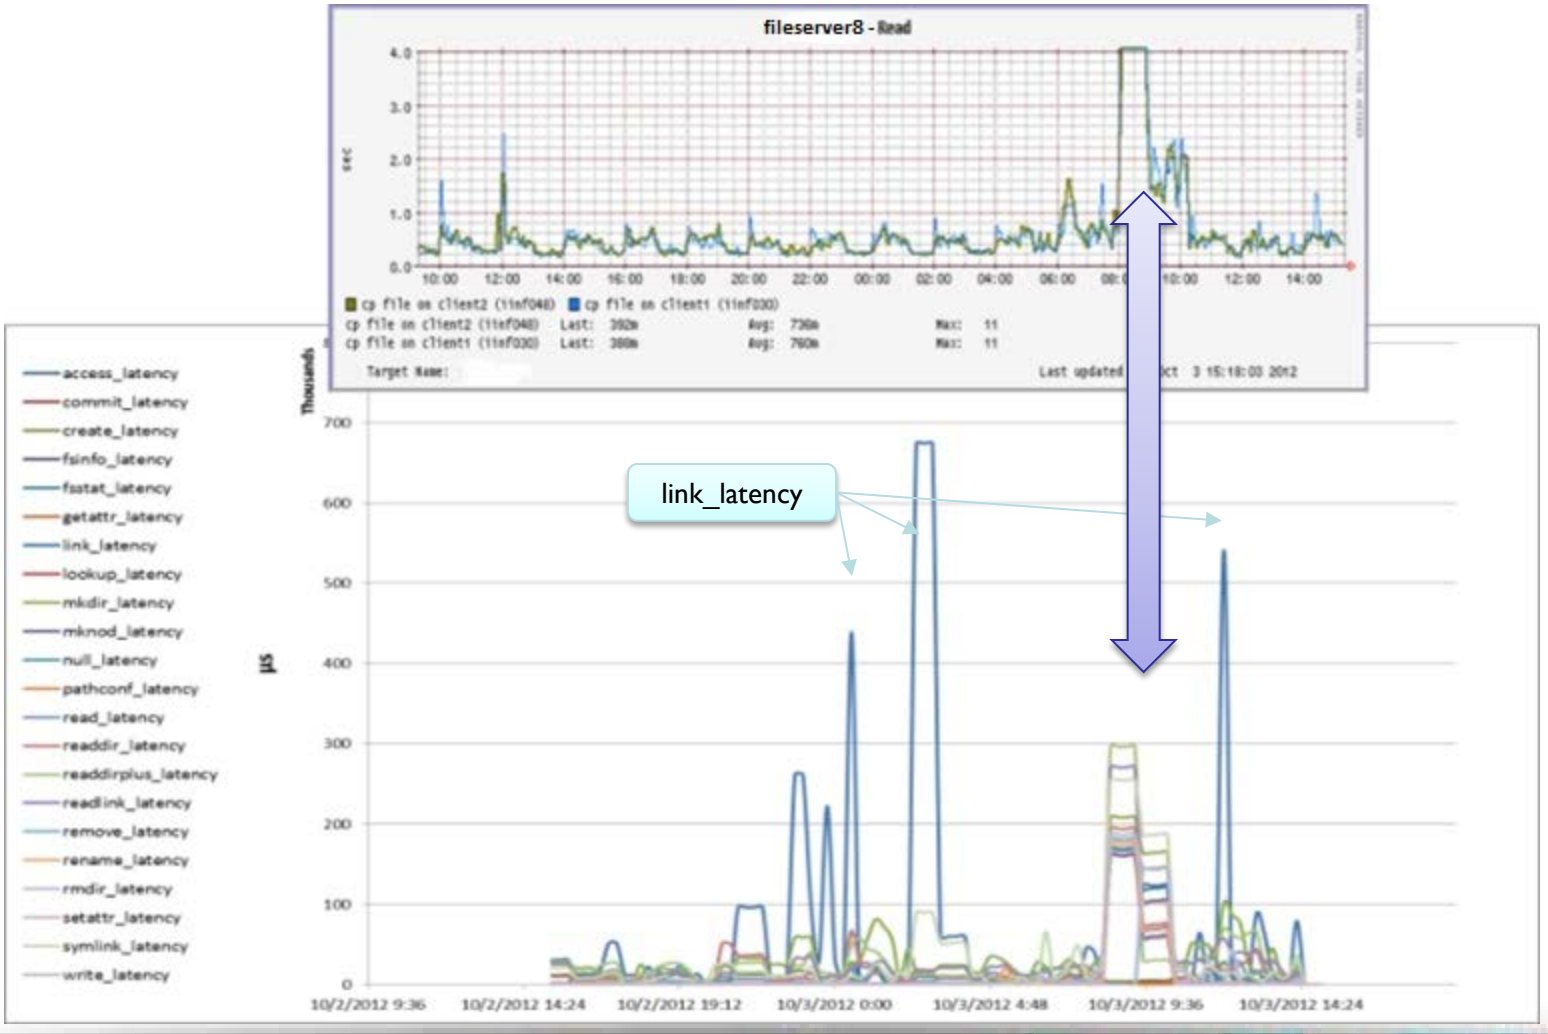

Analysis and Resolution

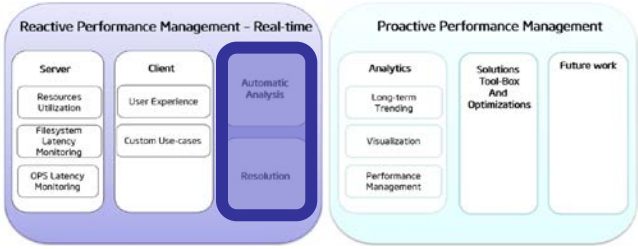

- We know that a NFS server is slow – now what ?

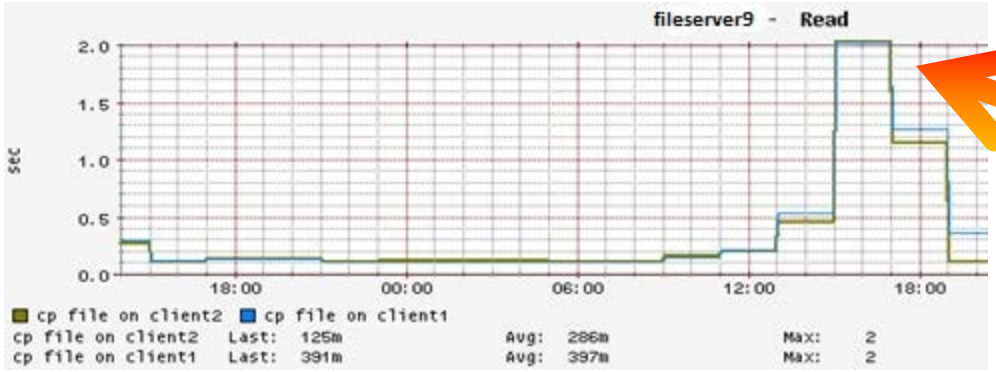

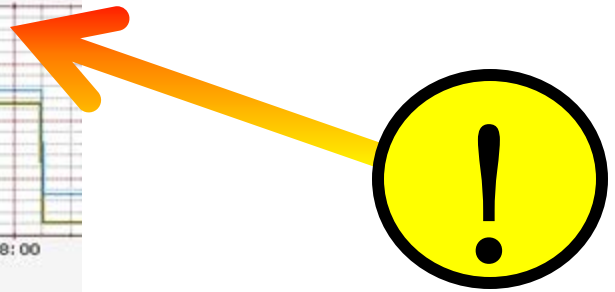

- We need to promptly clear the impact
- We need the same solutions for any platform

NFS slowness - FILESERVER10.intel.com (FATAL) - reason copy - sent by MONITORING_CLIENT2 - Message (HTML)

FILE | MESSAGE | INSERT | OPTIONS | FORMAT TEXT | REVIEW

Ignore | Junk | Delete | Reply | Reply All | Forward | Meeting | IM | More | temp | To Manager | Team Email | Move | Rules | OneNote | Actions | Assign Policy | Mark Unread | Categorize | Follow Up | Translate | Find | Related | Select | Zoom | Start Inking

Delete | Respond | Quick Steps | Move | Tags | Editing | Zoom | Ink

Sun 8/10/2014 10:43 PM

root

NFS slowness - FILESERVER10.intel.com (FATAL) - reason copy - sent by MONITORING_CLIENT2

To nfsadmin danyf

Retention Policy Mail Cloud - Inbox (60 days) Expires 10/9/2014

Fileserver name, severity (duration), copy/mount

Server model, tier and groups using it

Workload (OPS mix)

Top-hitter user with OPS mix and Batch jobs amounts and paths being used

```
Model   VENDOR_A MODEL_VV (tier 2)
Business groups with data on fileserver:
ABC
BCD

SUMMARY

                NFS operations
read                                    6589            39.92%
write                                     57             0.35%
lookup                                   916             5.55%
getattr                                 3303            20.01%
access                                  5640            34.17%

username1    OPS 10735      weight 3350.771914
(read=6546 write=0 lookup=734 getattr=1774 access=1670 create=0 remove=0 readdir=0 readdirplus=0 setattr=0 rmdir=0 symlink=0)
pool_1          1
pool_2          271
pool_3          405

disk6                   s                                       /nfs/site/disks/some_project_disk477                    8650
disk0                   s                                       /nfs/site/disks/some_project_disk059                    1625
disk0                   s                                       /nfs/site/disks/some_project_disk058                    155
disk0                   s                                       /nfs/site/disks/some_project_disk030                    151
disk0                   s                                       /nfs/site/disks/some_project_disk001                    82
disk0                   some_project_disk051                    /nfs/site/disks/some_project_disk051                    72
================================================================================================================================

username2    OPS 2173       weight 277.846433
(read=0 write=0 lookup=36 getattr=515 access=1618 create=0 remove=0 readdir=0 readdirplus=0 setattr=0 rmdir=0 symlink=0)
pool_1          262
pool_2          334
pool_3          193

disk0                   some_project_disk058                    /nfs/site/disks/some_project_disk058                    1181
disk0                   some_project_disk117                    /nfs/site/disks/some_project_disk117                    231
disk0                   some_project_disk051                    /nfs/site/disks/some_project_disk051                    222
disk0                   some_project_disk030                    /nfs/site/disks/some_project_disk030                    208
disk0                   some_project_disk059                    /nfs/site/disks/some_project_disk059                    192
disk0                   some_project_disk001                    /nfs/site/disks/some_project_disk001                    103
disk2                   some_project_disk071                    /nfs/site/disks/some_project_disk071                    36
================================================================================================================================
```

Resolution

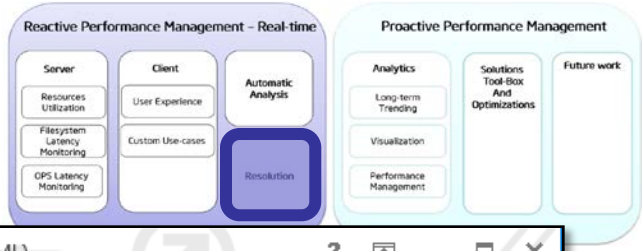

Your NB jobs were suspended due to high load on NFS areas - Message (HTML)

FILE | MESSAGE | INSERT | OPTIONS | FORMAT TEXT | REVIEW

Ignore | Junk | Delete | Reply | Reply All | Forward | Meeting | IM | More | temp | To Manager | Team Email | Move | Rules | OneNote | Actions | Assign Policy | Mark Unread | Categorize | Follow Up | Translate | Find | Related | Select | Zoom

Delete | Respond | Quick Steps | Move | Tags | Editing | Zoom

Sat 6/21/2014 4:35 AM

root

Your NB jobs were suspended due to high load on NFS areas

To

Cc

This item will expire in 6 days. To keep this item longer apply a different Retention Policy.

Hi **Jane Doe,**

About 1861 of your Netbatch jobs caused high load on the file server **FILESEVER12.intel.com**
In order to prevent slowness for other users and Netbatch jobs, we suspended your Netbatch jobs and resuming them in a way that prevents slowness
We suspect that the disk with high usage is **/nfs/site/disks/some_project_area01068**

Please note that this Policy was approved by Engineering Computing and customer groups.
Please consult your DA/RA and try to find a way to avoid these kinds of issues in the future.
We also suggest you consider submitting your jobs in small amounts (through a Netbatch feeder) to prevent causing environmental slowness.
Please contact us 24x7 of you have any question.

Have a nice day
Regards,
Computing Monitoring & Control Center

Automatic Resolution

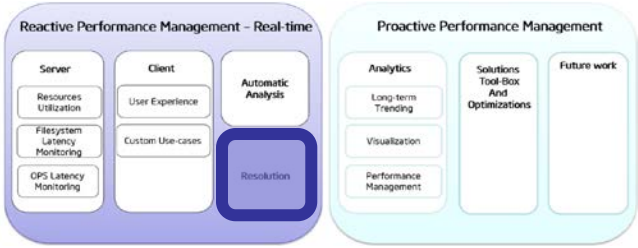

- Different approach – “autosuspend”
  - No human intervention required
  - Starts at the first sign of slowness
  - Hooks on “yana” and suspends jobs by clients’ traffic
- Tradeoff between **NFS Slowness** and **Batch jobs suspending**

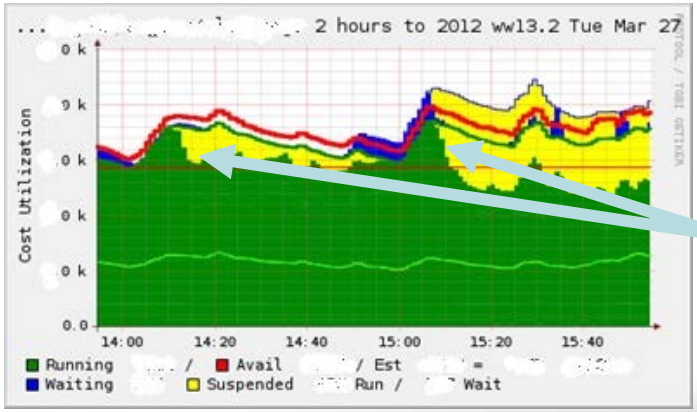

Rasta suspending jobs

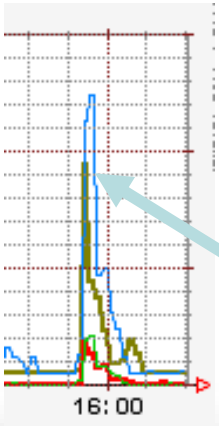

Autosuspend takes action after 90 seconds

Analytics

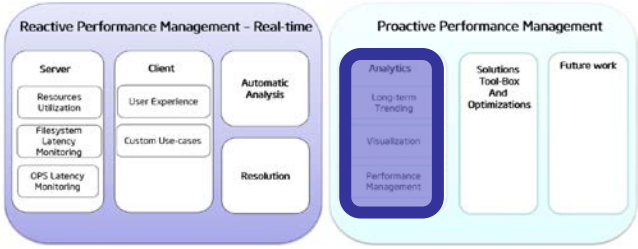

- Can't see the forest for the trees
  - Many sites, servers, admins, customers, jobs
  - Needed trending and visualization
  - Needed to focus on the “big” top-hitters

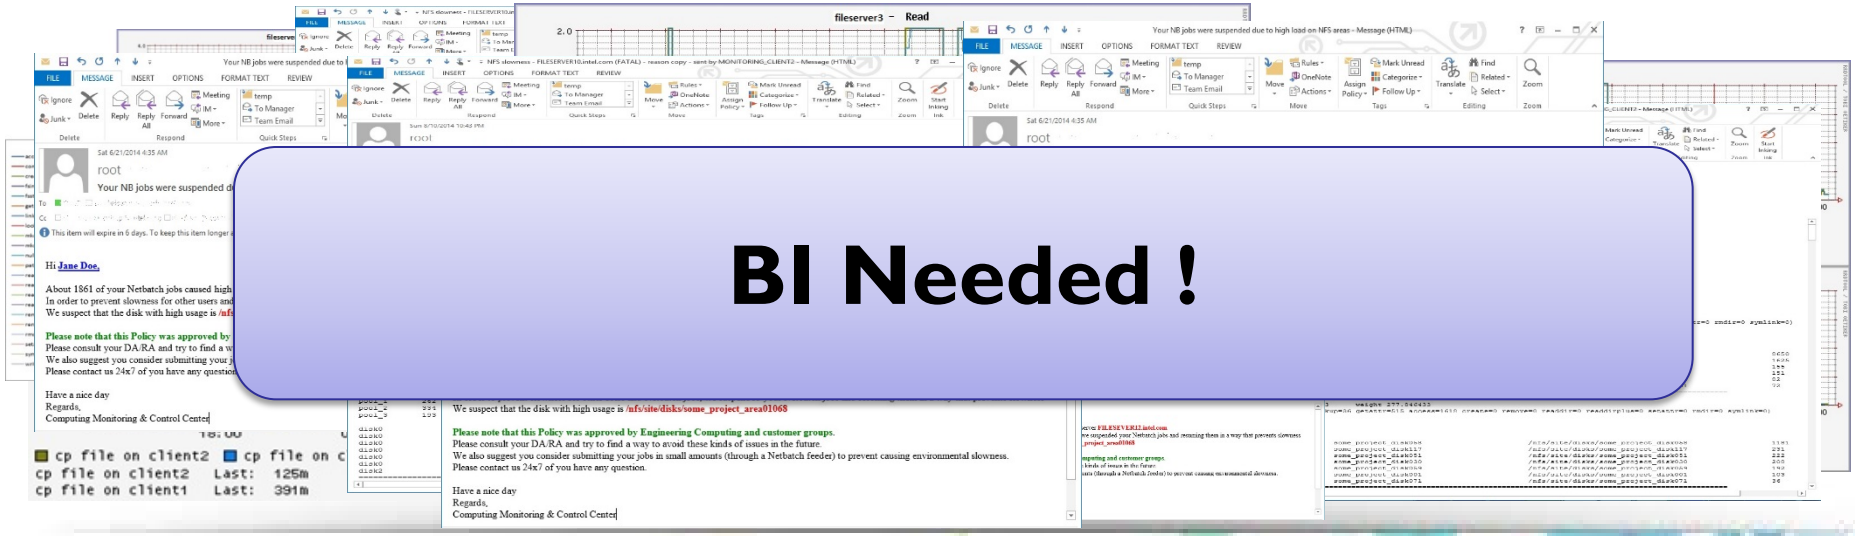

Trending

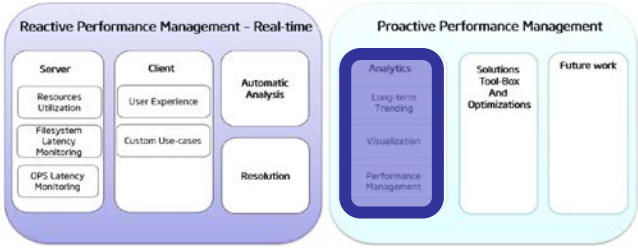

Top-Hitters

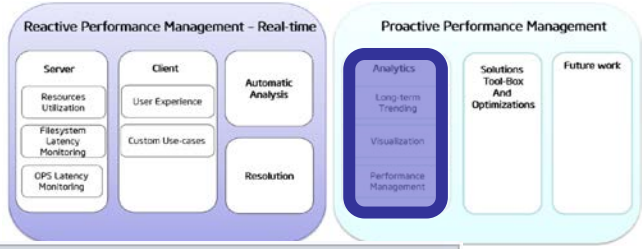

Currently applied filters

Work Week: [redacted]
severity: severity - CRITICAL, FATAL, MINOR;
cell: cell - [redacted]

Duration by Fileserver Model

	Fileserver	cell	tier	model	Business Groups	Duration
0	[redacted]	[redacted]	2	[redacted]	[redacted]	6525
1	[redacted]	[redacted]	3	[redacted]	[redacted]	4322
2	[redacted]	[redacted]	4	[redacted]	[redacted]	1773
3	[redacted]	[redacted]	1	[redacted]	[redacted]	1260
4	[redacted]	[redacted]	1		[redacted]	744
5	[redacted]	[redacted]	1	[redacted]	[redacted]	655
6	[redacted]	[redacted]	3	[redacted]	[redacted]	567
7	[redacted]	[redacted]	2	[redacted]	[redacted]	540
8	[redacted]	[redacted]	3	[redacted]	[redacted]	470
9	[redacted]	[redacted]	3	[redacted]	[redacted]	432
10	[redacted]	[redacted]	1	[redacted]	[redacted]	369
11	[redacted]	[redacted]	2	[redacted]	[redacted]	353
12	[redacted]	[redacted]	5	[redacted]	[redacted]	239
13	[redacted]	[redacted]	1	[redacted]	[redacted]	238
14	[redacted]	[redacted]	3	[redacted]	[redacted]	213
15	[redacted]	[redacted]	3	[redacted]	[redacted]	213
16	[redacted]	[redacted]	1	[redacted]	[redacted]	205
17	[redacted]	[redacted]		[redacted]	[redacted]	199
18	[redacted]					
19	[redacted]					

	User	cell	Total hours of suspend	Jobs count
0	[redacted]	[redacted]	104636	32652
1	[redacted]	[redacted]	48643	6625
2	[redacted]	[redacted]	48118	14470
3	[redacted]	[redacted]	39635	41259
4	[redacted]	[redacted]	39108	8682
5	[redacted]	[redacted]	35319	12703
6	[redacted]	[redacted]	27834	11950
7	[redacted]	[redacted]	23950	6618
8	[redacted]	[redacted]	18965	7157
9	[redacted]	[redacted]	17729	11981
10	[redacted]	[redacted]	17207	8606
11	[redacted]	[redacted]	10711	4639
12	[redacted]	[redacted]	9863	9891
13	[redacted]	[redacted]	7782	10842
14	[redacted]	[redacted]	7592	237219
15	[redacted]	[redacted]	7293	7344
16	[redacted]	[redacted]	5385	4945
17	[redacted]	[redacted]	4851	7889
18				1042
19				2632

Page 1 of 25

The fileservers with the most slowness and the users with the longest job suspends are what we want to take action on !

Drill-down for action

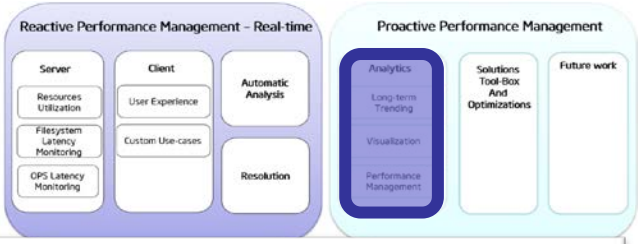

NFS Slowness – users & paths (by ops)

Currently applied filters

dsm_event_ww: 201413 - 201420;
severity: severity - CRITICAL, FATAL, MINOR;
cell: cell -

Filters

Apply All Filters

fileserver	path	cell	businessgroup	project	qslot	cl_fileserver	model	Total of Slowness	Times as top-hitter
	/nfs/ /proj/							392	5
	/nfs/ /proj/							383	3
	/nfs/ /proj/							145	17
	/nfs/ /disks/							96	14
	/nfs/ /proj/							48	1
	/nfs/ /proj/							45	3
	/nfs/ /proj/							40	1
	/nfs/ /disks/							32	3
	/nfs/ /proj/							31	1
	/nfs/ /disks/							31	5
	/nfs/ /disks/							31	1
	/nfs/ /disks/							30	1
	/nfs/ /disks/							28	2
root	/nfs/ /proj/							26	1
									2
									2
									1
									2
									1
									3

We can tell which problems are reoccurring by fileserver, specific export/path and activity

Our Solutions Tool-Box

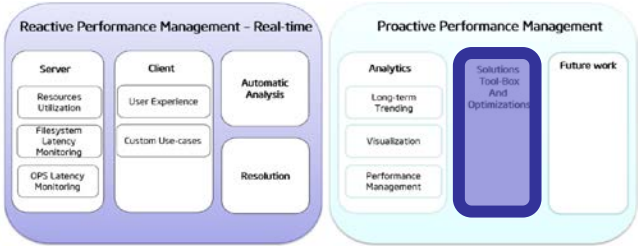

- Data migration – to appropriate tier or platform
- Batch dispatch configurations – slow ramp
- Local disk caching enablement
- Smarter Allocation – consider slowness history
- Customer Flow changes
- Refresh of old HW
- Backups tuning

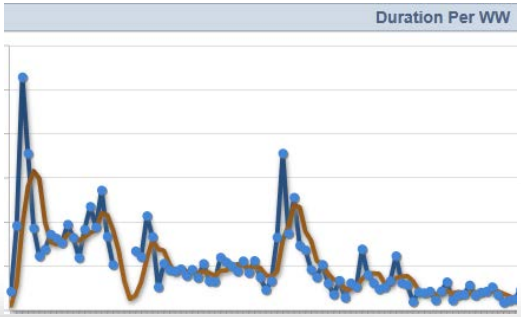

Future Plans

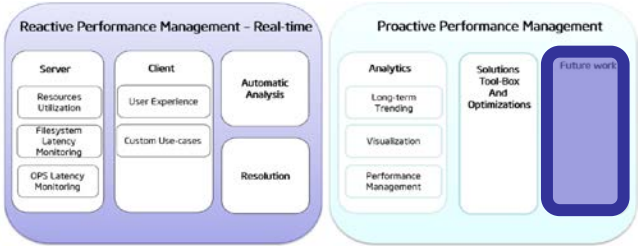

- Checking latencies usability
- Integration of new platforms
- Analyzing non-NFS traffic
- Monitoring SMB performance
- QoS implementation
- Automation for reoccurring top-hitters

Lessons Learnt

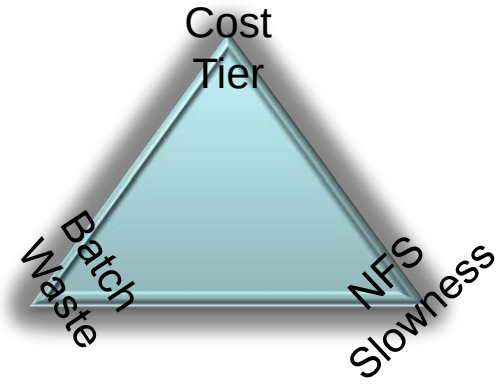

- User-experience monitoring
- Sometimes the simplest solution works best
- Automate everything, with caution

- NFS performance issues can be resolved
  - Sometimes a downtime may be required

Call For Action

- What we need from the NAS Vendors and/of the open source community:
  - Identify “Hot files/users/clients”
  - Alerts on counters with thresholds
  - **When** there is a performance problem ?
  - And **Why** would be ideal –

    CPU ? Cache ? Disks ? Network ?
- Standards and API ?

Q&A

- Any questions or feedback, please contact me at dany.felzenszwalbe@intel.com