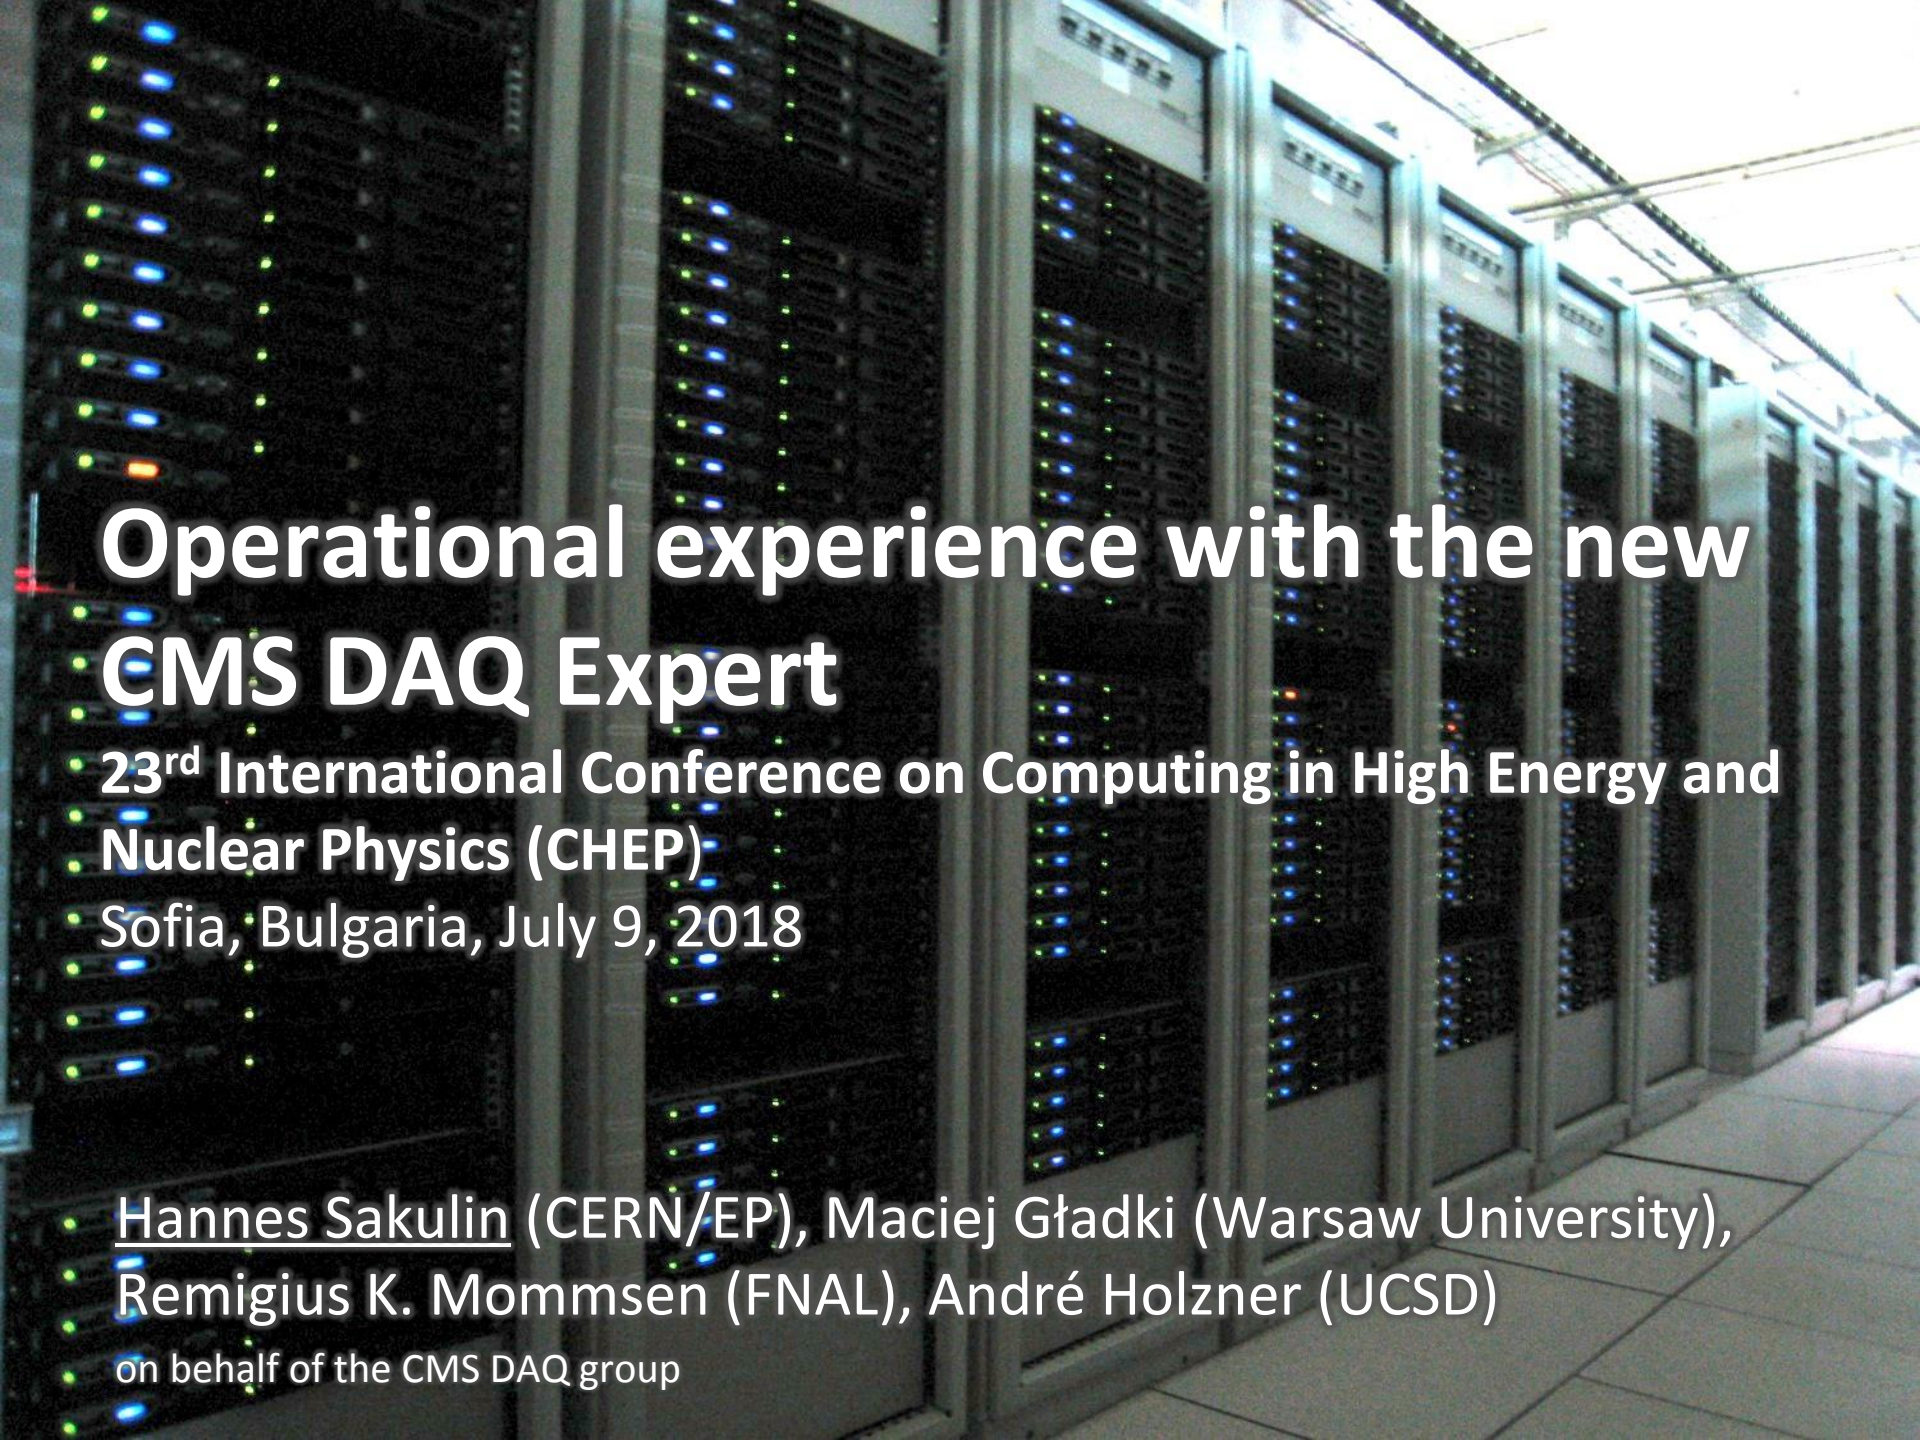

Operational experience with the new CMS DAQ Expert

23rd International Conference on Computing in High Energy and Nuclear Physics (CHEP)

Sofia, Bulgaria, July 9, 2018

Hannes Sakulin (CERN/EP), Maciej Gładki (Warsaw University), Remigius K. Mommsen (FNAL), André Holzner (UCSD)

on behalf of the CMS DAQ group

The CMS DAQ-2 system at a glance

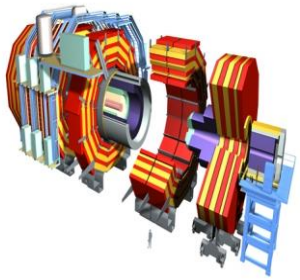

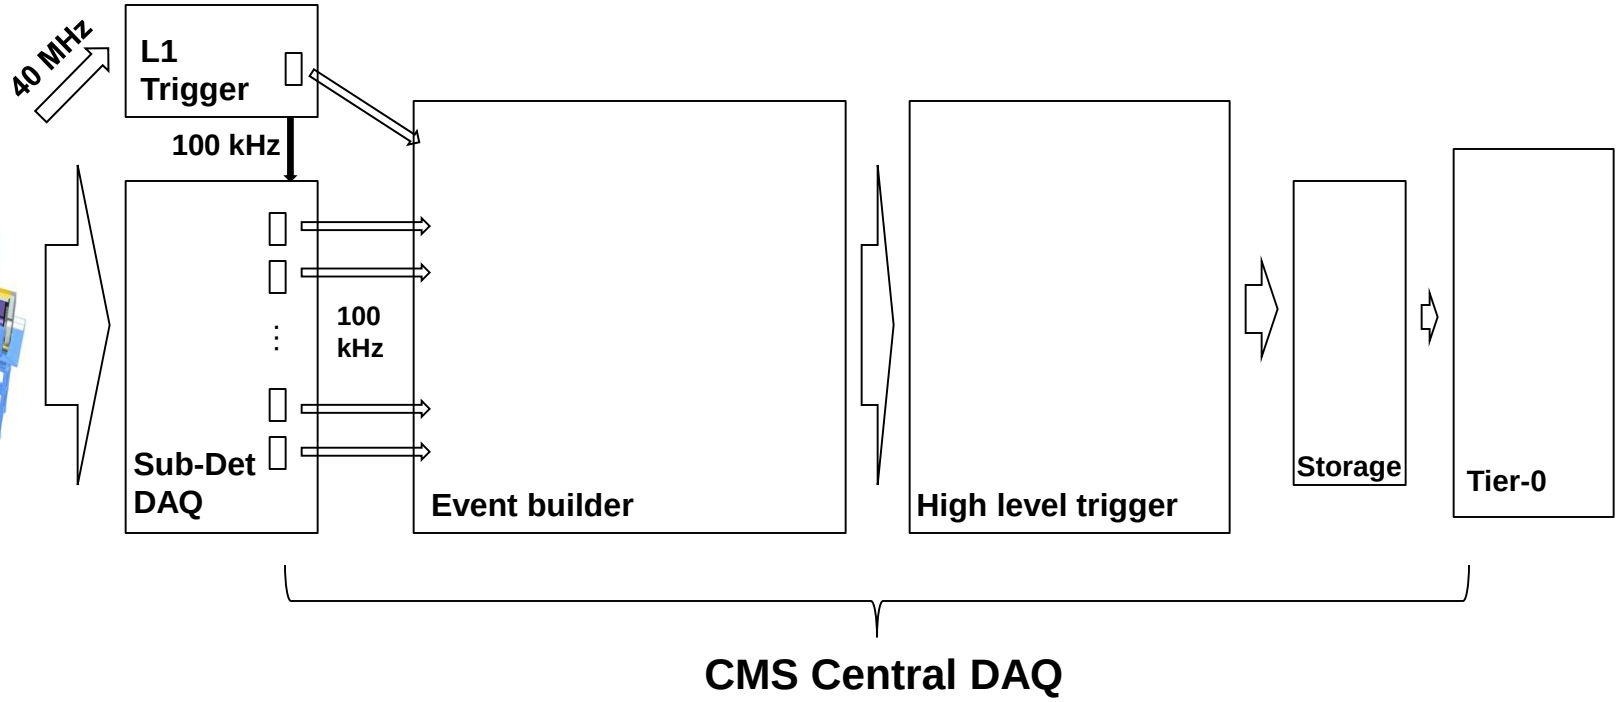

The CMS DAQ-2 system at a glance

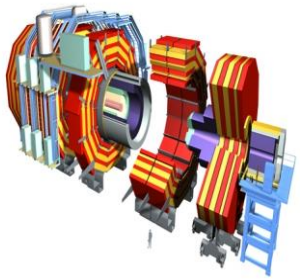

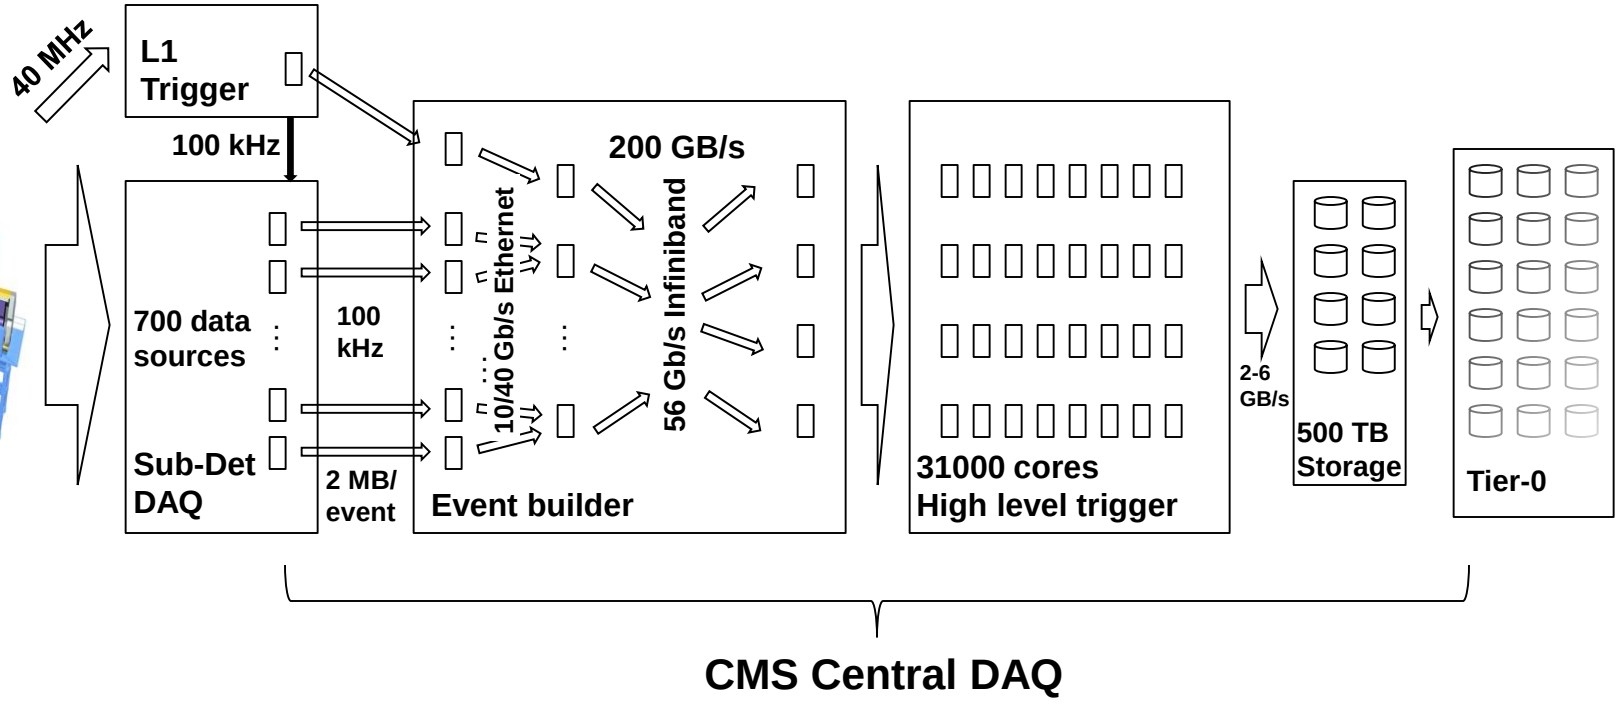

The CMS DAQ-2 system at a glance

In general it works reliably …

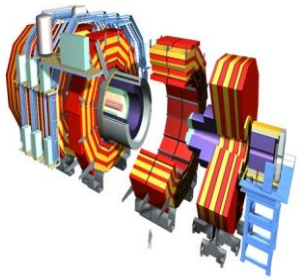

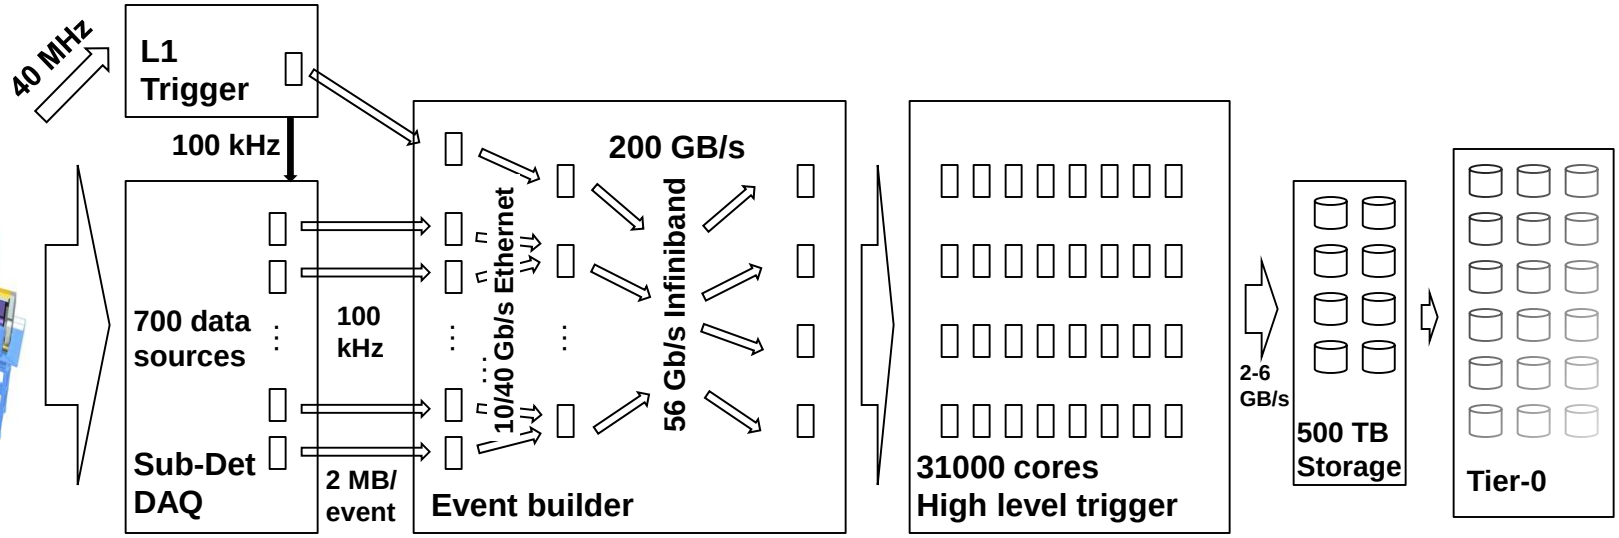

- **<0.2% of luminosity lost due to central DAQ,**
  2-4% percent of luminosity lost due to subdetector DAQ and trigger (LHC Run-2)

- **Built-in automation in run control(*)**
  - **reaction to LHC & HV state changes**
  - **recovery from expected problems (e.g. single event upsets)**

(*) CHEP 2013: H. Sakulin et. al. "Automating the CMS DAQ"

Sometimes unexpected problems may arise around or in the DAQ system …

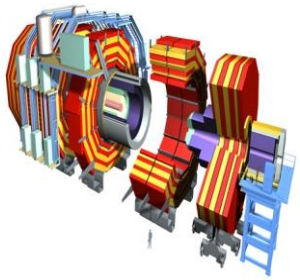

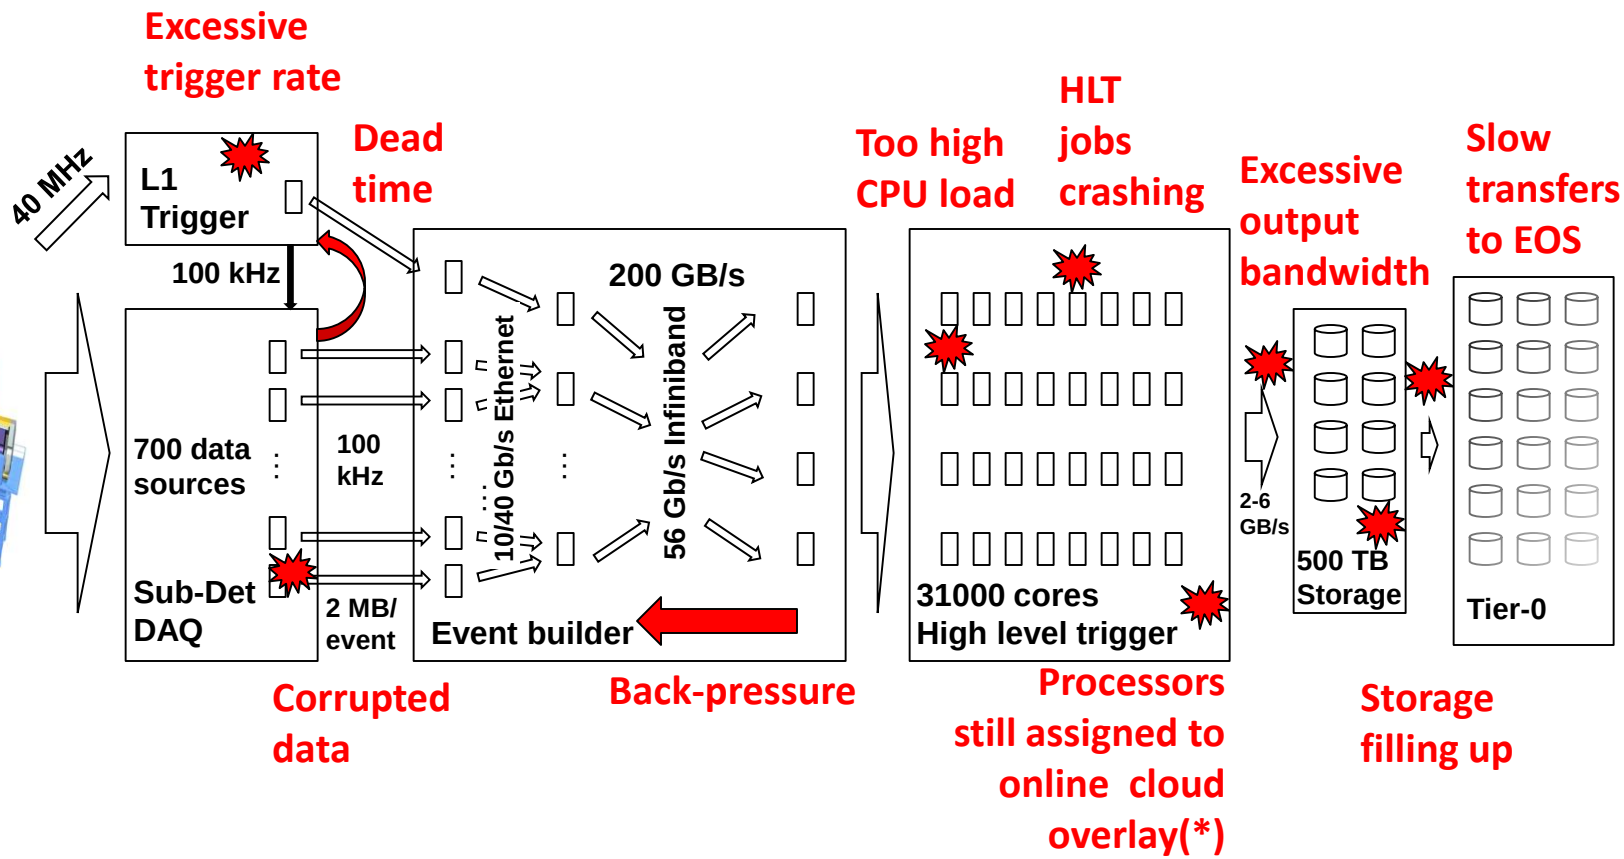

(*) tomorrow, 11.45, T7: “Experience with dynamic resource provisioning of the CMS online cluster using a cloud overlay”

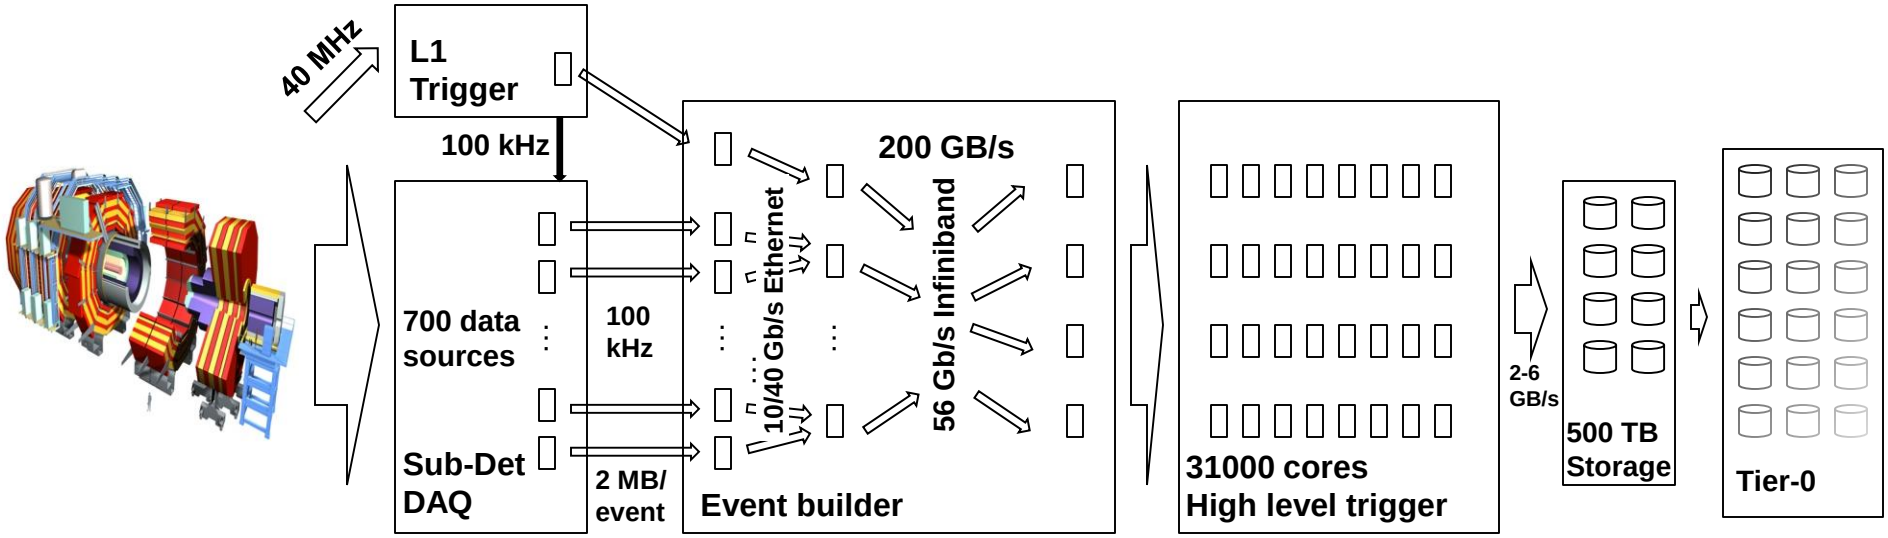
40 MHz
L1
Trigger
100 kHz
700 data
sources
Sub-Det
DAQ
100
kHz
2 MB/
event
200 GB/s
10/40 Gb/s Ethernet
56 Gb/s Infiniband
Event builder
31000 cores
High level trigger
2-6
GB/s
500 TB
Storage
Tier-0

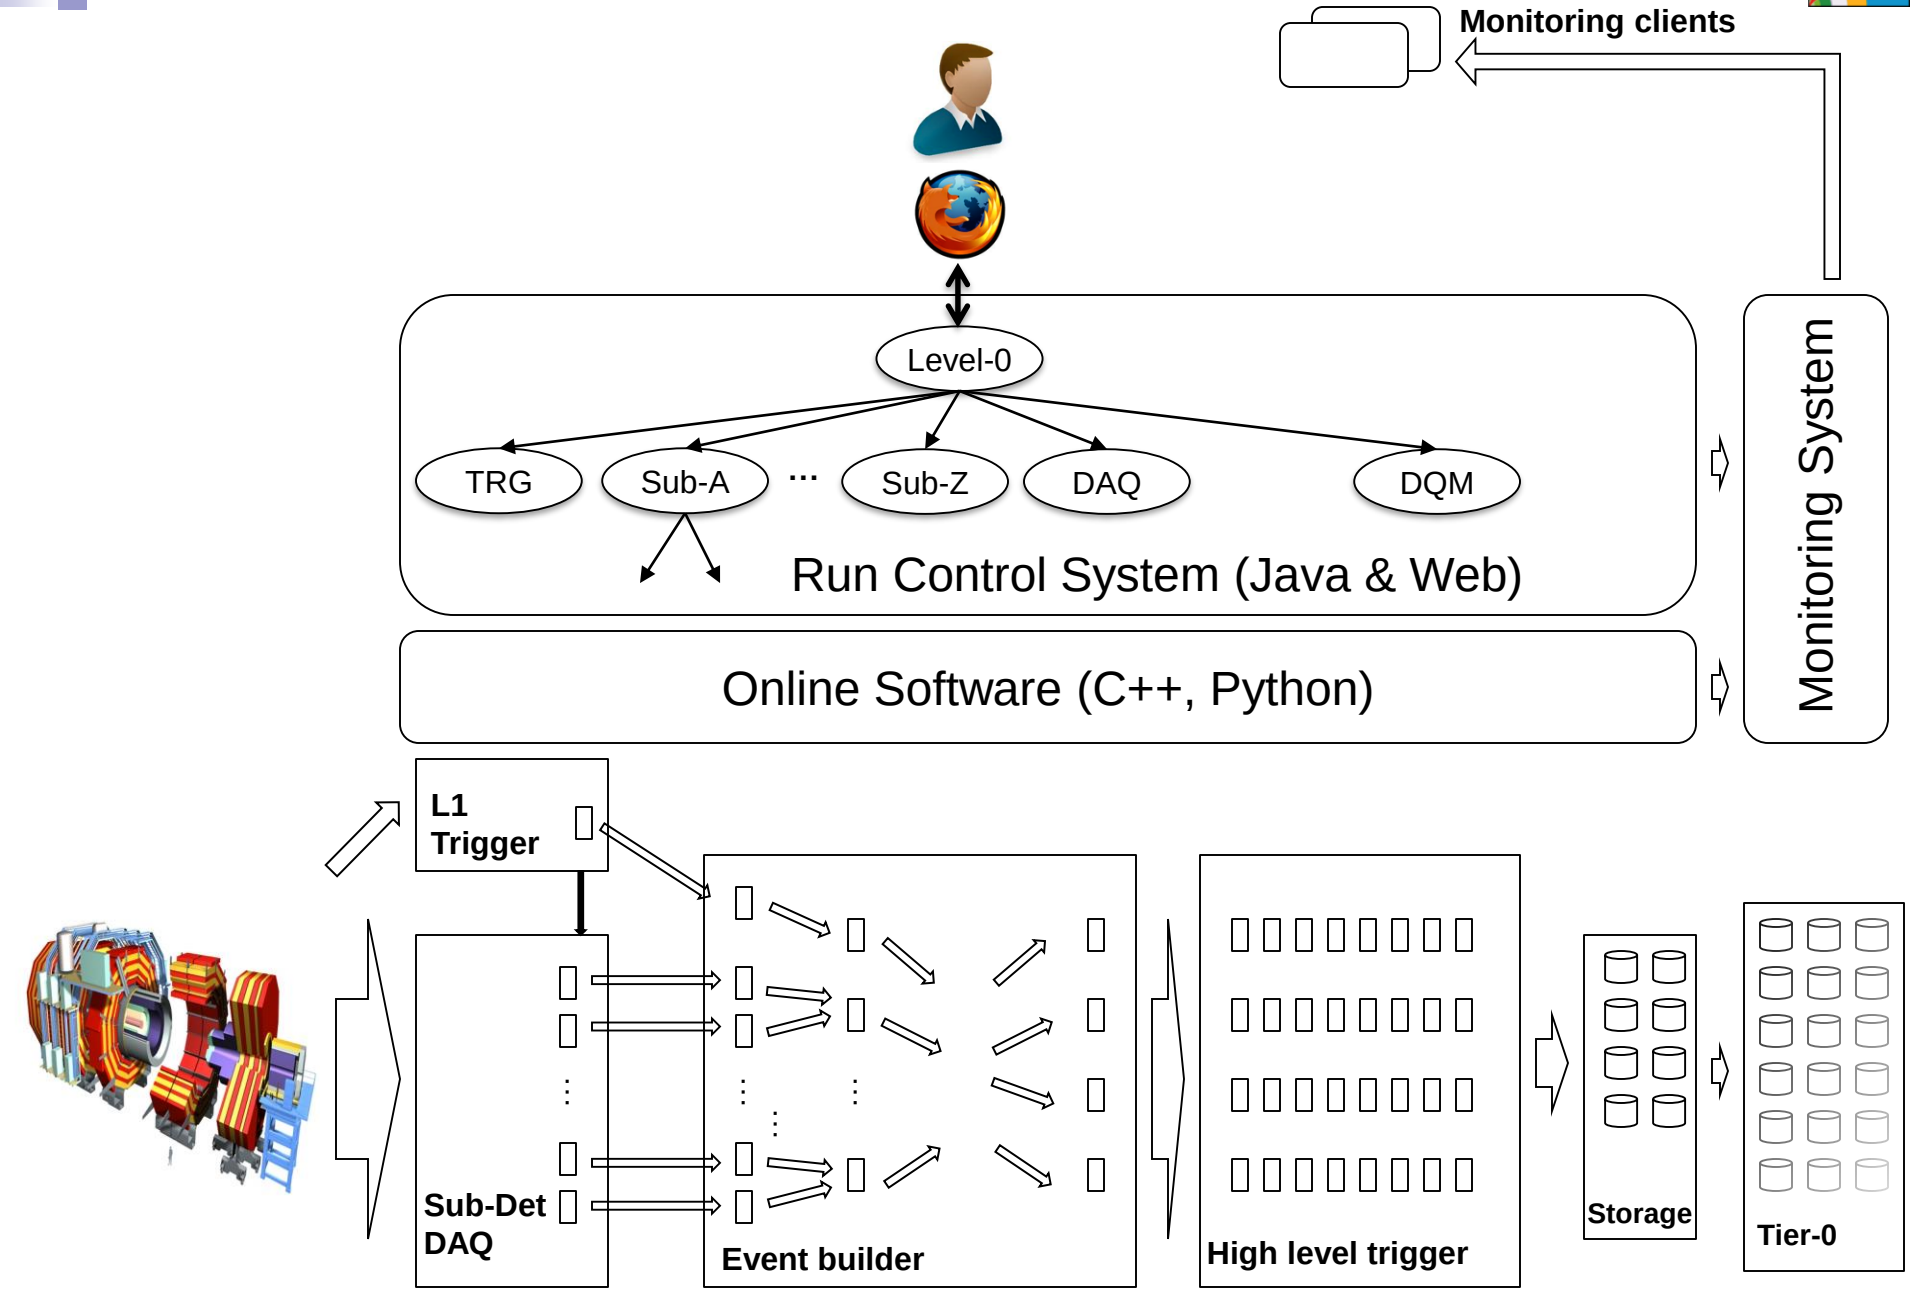
Monitoring clients
Level-0
TRG
Sub-A
Sub-Z
DAQ
DQM
Run Control System (Java & Web)
Online Software (C++, Python)
Monitoring System
L1 Trigger
Sub-Det DAQ
Event builder
High level trigger
Storage
Tier-0

A look at the monitoring clients reveals what went wrong …

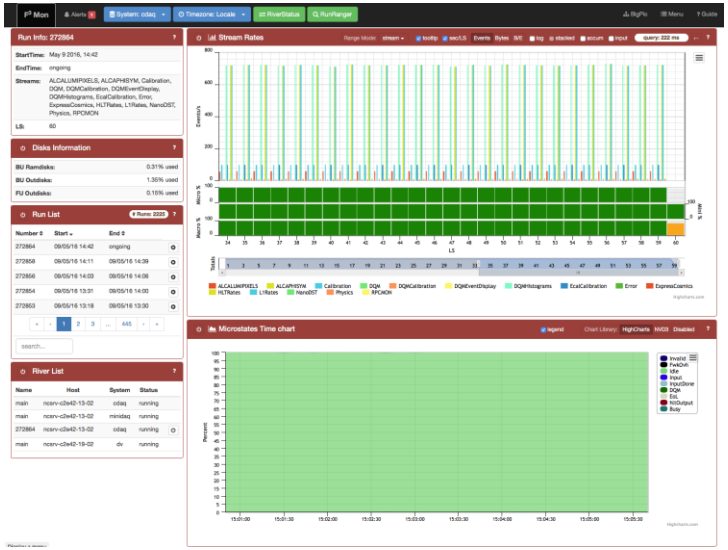

(Table Help)

	FEDBUILDER									EVB						
TTCP	T	%W	%B	frlpc	geoSlot:SrcId / TTSOnlyFEDSrcId	min Trg	max Trg	FB Name	RU	rate (kHz)	thru (MB/s)	size (kB)	#events	#frags in RU	#evts in RU	#evts requested
cpm-pri:59	-	-	-	s1d06-40-01	1:**1024**		9632	TCDS	e14-10-01	0.000	0.0	0.000±0.000	8064	512	0	0
CSC+:16	R	0.0	0.0	s1d06-34-01	1:**841**,**831**, 2:**842**,**832**, 3:**843**,**833**, 4:**844**,**834**, 5:**845**,**835**, 6:**846**,**836**, 7:**847**,837, 8:**848**,**838**, 18:**849**,**839**		9632	CSC+	e12-34-01	0.000	0.0	0.000±0.000	8064	512	0	0
CSC-:17	R	0.0	0.0	s1d06-34-01	9:**861**,**851**, 10:**862**,**852**, 11:**863**,**853**, 12:**864**,**854**, 13:**865**,**855**, 14:**866**,**856**, 15:**867**,**857**, 16:**868**,**858**, 17:**869**,**859**		9632	CSC-	e12-35-01	0.000	0.0	0.000±0.000	8064	512	0	0
EB+:0	-	?	?	s2d10-10-01	1:628, 2:629, 3:630, 4:634, 5:635, 6:636, 7:631, 8:632, 9:633 663			EB+1	e12-15-01	0.000	0.0	0.000±0.000	0	0	0	0
EB+:0	-	?	?	s2d10-18-01	1:643, 2:644, 3:645, 4:637, 5:638, 6:639, 7:640, 8:641, 9:642 663			EB+2	e12-26-01	0.000	0.0	0.000±0.000	0	0	0	0
EB-:1	R	0.0	0.0	s2d10-11-01	1:**610**, 2:**611**, 3:**612**, 4:**616**#FCRC=1, 5:**617**, 6:**618**, 7:**613**, 8:**614**, 9:**615 662**		9632	EB-1	e12-18-01	0.000	0.0	0.000±0.000	8064	512	0	0
EB-:1	R	0.0	0.0	s2d10-20-01	1:**625**, 2:**626**, 3:**627**, 4:**619**, 5:**620**#FCRC=1, 6:**621**, 7:**622**, 8:**623**, 9:**624 662**		9632	EB-2	e14-28-01	0.000	0.0	0.000±0.000	8064	512	0	0
EE+:2	-	?	?	s2d10-17-01	1:648, 2:649, 3:650, 4:654, 5:646, 6:647, 7:651, 8:652, 9:653 664			EE+	e12-19-01	0.000	0.0	0.000±0.000	0	0	0	0
EE-:3	W	100.0	0.0	s2d10-09-01	1:**606**, 2:**607**, 3:**608**, 4:**609**, 5:**601**, 6:W:100.0%**W602<80.0%**9605, 7:**603**, 8:**604**, 9:**605 661**	9605	9632	EE-	e12-10-01	0.000	0.0	0.000±0.000	8064	512	0	0
GCT:11	-	?	?	s1d06-29-01	1:745											
GT:33	R	0.0	0.0	s1d06-27-01	1:812, 2:**813 811**		9632	TRG	e15-22-01	0.000	0.0	0.000±0.000	8064	512	0	0
RETRI:4	-	?	?	-	810	-	-									
TIBTID:24	R	0.0	0.0	s1d06-04-01	1:**149**, 2:**150**, 3:**151**, 4:**152**, 5:**153**, 6:**154**, 7:**155**, 8:**156**, 9:**157**, 10:**158**, 11:**159**, 12:**160**, 13:**161**, 14:**162**, 15:**163**		9632	TIBTID4	e12-30-01	0.000	0.0	0.000±0.000	8064	512	0	0
TIBTID:24	R	0.0	0.0	s1d06-05-01	1:**134**, 2:**135**, 3:**136**, 4:**137**, 5:**138**, 6:**139**, 7:**140**, 8:**141**, 9:**142**, 10:**143**, 11:**144**, 12:**145**, 13:**146**, 14:**147**, 15:**148**		9632	TIBTID5	e13-10-01	0.000	0.0	0.000±0.000	8064	512	0	0
TIBTID:24	R	0.0	0.0	s1d06-07-01	1:**118**,**74**, 2:**119**,**75**, 3:**120**,**76**, 4:**121**,**77**, 5:**122**,**78**, 6:**123**,**79**, 7:**124**,**80**, 8:**125**,**81**		9632	TIBTID6a	e13-15-01	0.000	0.0	0.000±0.000	8064	512	0	0
TIBTID:24	R	0.0	0.0	s1d06-07-01	9:**126**,**82**, 10:**127**,**83**, 11:**128**,**84**, 12:**129**,**85**, 13:**130**, 14:**131**, 15:**132**, 16:**133**		9632	TIBTID6b	e15-13-01	0.000	0.0	0.000±0.000	8064	512	0	0
TIBTID:24	R	0.0	0.0	s1d06-01-01	1:**86**,**87**, 2:**88**,**89**, 3:**90**,**91**, 4:**92**,**93**, 5:**94**,**95**, 6:**96**,**97**, 7:**98**,**99**, 8:**100**,**101**		9632	TIBTID1a	e13-16-01	0.000	0.0	0.000±0.000	8064	512	0	0
TIBTID:24	R	0.0	0.0	s1d06-01-01	9:**102**, 10:**104**,**105**, 11:**106**,**107**, 12:**108**,**109**, 13:**110**,**111**, 14:**112**,**113**, 15:**114**,**115**, 16:**116**,**117**		9632	TIBTID1b	e13-17-01	0.000	0.0	0.000±0.000	8064	512	0	0
TIBTID:24	R	0.0	0.0	s1d06-02-01	1:**50**, 2:**51**, 3:**52**, 4:**53**, 5:**54**, 6:**55**, 7:**56**, 8:**57**, 9:**58**, 10:**590**, 11:**60**, 12:**61**	0	9632	TIBTID2	e13-22-01 [59]	0.000	0.0	0.000±0.000	0	512	0	8064
TIBTID:24	R	0.0	0.0	s1d06-03-01	1:**62**, 2:**63**, 3:**64**, 4:**65**, 5:**66**, 6:**67**, 7:**68**, 8:**69**, 9:**70**, 10:**71**, 11:**72**, 12:**73**		9632	TIBTID3	e13-25-01	0.000	0.0	0.000±0.000	8064	512	0	0
Summary				frlpc	geoSlot:SrcId	min Trg	max Trg	FB Name	RU	rate (kHz)	thru (MB/s)	size (kB)	#events	#frags in RU	#evts in RU	#evts requested
										0.000	Σ 0.0	Σ 0.000±0.000	Δ 8064	Σ 21504	Σ 0	Σ 8064

A look at the monitoring clients reveals what went wrong …

(Table Help) FEDBUILDER

TTCP	T	%W	%B	frlpc	geoSlot:SrcId / TTSOnlyFEDSrcId	min Trg	max Trg	FB Name	RU	rate (kHz) thru
cpm-pri:59	-	-	-	s1d06-40-01	1:1024		9632	TCDS	e14-10-01	0.000
CSC+:16	R	0.0	0.0	s1d06-34-01	1:841,831, 2:842,832, 3:843,833, 4:844,834, 5:845,835, 6:846,836, 7:847,837, 8:848,838, 18:849,839		9632	CSC+	e12-34-01	0.000
CSC-:17	R	0.0	0.0	s1d06-34-01	9:861,851, 10:862,852, 11:863,853, 12:864,854, 13:865,855, 14:866,856, 15:867,857, 16:868,858, 17:869,859		9632	CSC-	e12-35-01	0.000
EB+:0	-	?	?	s2d10-10-01	1:628, 2:629, 3:630, 4:634, 5:635, 6:636, 7:631, 8:632, 9:633 663			EB+1	e12-15-01	0.000
EB+:0	-	?	?	s2d10-18-01	1:643, 2:644, 3:645, 4:637, 5:638, 6:639, 7:640, 8:641, 9:642 663			EB+2	e12-26-01	0.000
EB-:1	R	0.0	0.0	s2d10-11-01	1:610, 2:611, 3:612, 4:616#FCRC=1, 5:617, 6:618, 7:613, 8:614, 9:615 662		9632	EB-1	e12-18-01	0.000
EB-:1	R	0.0	0.0	s2d10-20-01	1:625, 2:626, 3:627, 4:619, 5:620#FCRC=1, 6:621, 7:622, 8:623, 9:624 662		9632	EB-2	e14-28-01	0.000
EE+:2	-	?	?	s2d10-17-01	1:648, 2:649, 3:650, 4:654, 5:646, 6:647, 7:651, 8:652, 9:653 664			EE+	e12-19-01	0.000
EE-:3	W	100.0	0.0	s2d10-09-01	1:606, 2:607, 3:608, 4:609, 5:601, 6:W:100.0%W602<80.0%9605, 7:603, 8:604, 9:605 661	9605	9632	EE-	e12-10-01	0.000
GCT:11	-	?	?	s1d06-29-01	1:745					
GT:33	R	0.0	0.0	s1d06-27-01	1:812, 2:813 811		9632	TRG	e15-22-01	0.000
RETRI:4	-	?	?	-	810	-	-			
TIBTID:24	R	0.0	0.0	s1d06-04-01	1:149, 2:150, 3:151, 4:152, 5:153, 6:154, 7:155, 8:156, 9:157, 10:158, 11:159, 12:160, 13:161, 14:162, 15:163		9632	TIBTID4	e12-30-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-05-01	1:134, 2:135, 3:136, 4:137, 5:138, 6:139, 7:140, 8:141, 9:142, 10:143, 11:144, 12:145, 13:146, 14:147, 15:148		9632	TIBTID5	e13-10-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-07-01	1:118,74, 2:119,75, 3:120,76, 4:121,77, 5:122,78, 6:123,79, 7:124,80, 8:125,81		9632	TIBTID6a	e13-15-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-07-01	9:126,82, 10:127,83, 11:128,84, 12:129,85, 13:130, 14:131, 15:132, 16:133		9632	TIBTID6b	e15-13-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-01-01	1:86,87, 2:88,89, 3:90,91, 4:92,93, 5:94,95, 6:96,97, 7:98,99, 8:100,101		9632	TIBTID1a	e13-16-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-01-01	9:102, 10:104,105, 11:106,107, 12:108,109, 13:110,111, 14:112,113, 15:114,115, 16:116,117		9632	TIBTID1b	e13-17-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-02-01	1:50, 2:51, 3:52, 4:53, 5:54, 6:55, 7:56, 8:57, 9:58, 10:590, 11:60, 12:61	0	9632	TIBTID2	e13-22-01 [59]	0.000
TIBTID:24	R	0.0	0.0	s1d06-03-01	1:62, 2:63, 3:64, 4:65, 5:66, 6:67, 7:68, 8:69, 9:70, 10:71, 11:72, 12:73		9632	TIBTID3	e13-25-01	0.000
Summary				frlpc	geoSlot:SrcId	min Trg	max Trg	FB Name	RU	rate (kHz) thru
										0.000

… at least to the eyes of an expert

A look at the monitoring clients reveals what went wrong …

ECAL is 100% in Warning

FED 602 is in warning and last event is 9605

There's backpressure from DAQ

Rate is 0 kHz

(Table Help) FEDBUILDER

TTCP	T	%W	%B	frlpc	geoSlot:SrcId / TTSOnlyFEDSrcId	min Trg	max Trg	FB Name	RU	rate (kHz) thru
cpm-pri:59	-	-	-	s1d06-40-01	1:1024		9632	TCDS	e14-10-01	0.000
CSC+:16	R	0.0	0.0	s1d06-34-01	1:841,831, 2:842,832, 3:843,833, 4:844,834, 5:845,835, 6:846,836, 7:847,837, 8:848,838, 18:849,839		9632	CSC+	e12-34-01	0.000
CSC-:17	R	0.0	0.0	s1d06-34-01	9:861,851, 10:862,852, 11:863,853, 12:864,854, 13:865,855, 14:866,856, 15:867,857, 16:868,858, 17:869,859		9632	CSC-	e12-35-01	0.000
EB+:0	-	?	?	s2d10-10-01	1:628, 2:629, 3:630, 4:634, 5:635, 6:636, 7:631, 8:632, 9:633 663			EB+1	e12-15-01	0.000
EB+:0	-	?	?	s2d10-18-01	1:643, 2:644, 3:645, 4:637, 5:638, 6:639, 7:640, 8:641, 9:642 666			EB+2	e12-26-01	0.000
EB-:1	R	0.0	0.0	s2d10-11-01	1:610, 2:611, 3:612, 4:616#FCRC=1, 5:617, 6:618, 7:613, 8:614, 9:615 662		9632	EB-1	e12-18-01	0.000
EB-:1	R	0.0	0.0	s2d10-20-01	1:625, 2:626, 3:627, 4:619, 5:620#FCRC=1, 6:621, 7:622, 8:623, 9:624 662		9632	EB-2	e14-28-01	0.000
EE+:2	-	?	?	s2d10-17-01	1:648, 2:649, 3:650, 4:654, 5:646, 6:647, 7:651, 8:652, 9:653 664			EE+	e12-19-01	0.000
EE-:3	W	100.0	0.0	s2d10-09-01	1:606, 2:607, 3:608, 4:609, 5:601, 6:W:100.0%W602<80.0%9605, 7:603, 8:604, 9:605 661	9605	9632	EE-	e12-10-01	0.000
GCT:11	-	?	?	s1d06-29-01	1:745					
GT:33	R	0.0	0.0	s1d06-27-01	1:812, 2:813 811		9632	TRG	e15-22-01	0.000
RETRI:4	-	?	?	-	810	-	-			
TIBTID:24	R	0.0	0.0	s1d06-04-01	1:149, 2:150, 3:151, 4:152, 5:153, 6:154, 7:155, 8:156, 9:157, 10:158, 11:159, 12:160, 13:161, 14:162, 15:163		9632	TIBTID4	e12-30-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-05-01	1:134, 2:135, 3:136, 4:137, 5:138, 6:139, 7:140, 8:141, 9:142, 10:143, 11:144, 12:145, 13:146, 14:147, 15:148		9632	TIBTID5	e13-10-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-07-01	1:118,74, 2:119,75, 3:120,76, 4:121,77, 5:122,78, 6:123,79, 7:124,80, 8:125,81		9632	TIBTID6a	e13-15-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-07-01	9:126,82, 10:127,83, 11:128,84, 12:129,85, 13:130, 14:131, 15:132, 16:133		9632	TIBTID6b	e15-13-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-01-01	1:86,87, 2:88,89, 3:90,91, 4:92,93, 5:94,95, 6:96,97, 7:98,99, 8:100,101		9632	TIBTID1a	e13-16-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-01-01	9:102, 10:104,105, 11:106,107, 12:108,109, 13:110,111, 14:112,113, 15:114,115, 16:116,117		9632	TIBTID1b	e13-17-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-02-01	1:50, 2:51, 3:52, 4:53, 5:54, 6:55, 7:56, 8:57, 9:58, 10:590, 11:60, 12:61	0	9632	TIBTID2	e13-22-01 [59]	0.000
TIBTID:24	R	0.0	0.0	s1d06-03-01	1:62, 2:63, 3:64, 4:65, 5:66, 6:67, 7:68, 8:69, 9:70, 10:71, 11:72, 12:73		9632	TIBTID3	e13-25-01	0.000
Summary				frlpc	geoSlot:SrcId	min Trg	max Trg	FB Name	RU	rate (kHz) thru
										0.000

FED 59 has not sent any data

RU waits for data from FED 59

FED 59 is the culprit ➔ Talk to Tracker expert

But …

(Table Help) FEDBUILDER

TTCP	T	%W	%B	frlpc	geoSlot:SrcId / TTSOnlyFEDSrcId	min Trg	max Trg	FB Name	RU	rate (kHz) thru
cpm-pri:59	-	-	-	s1d06-40-01	1:1024		9632	TCDS	e14-10-01	0.000
CSC+:16	R	0.0	0.0	s1d06-34-01	1:841,831, 2:842,832, 3:843,833, 4:844,834, 5:845,835, 6:846,836, 7:847,837, 8:848,838, 18:849,839		9632	CSC+	e12-34-01	0.000
CSC-:17	R	0.0	0.0	s1d06-34-01	9:861,851, 10:862,852, 11:863,853, 12:864,854, 13:865,855, 14:866,856, 15:867,857, 16:868,858, 17:869,859		9632	CSC-	e12-35-01	0.000
EB+:0	-	?	?	s2d10-10-01	1:628, 2:629, 3:630, 4:634, 5:635, 6:636, 7:631, 8:632, 9:633 663			EB+1	e12-15-01	0.000
EB+:0	-	?	?	s2d10-18-01	1:643, 2:644, 3:645, 4:637, 5:638, 6:639, 7:640, 8:641, 9:642 663			EB+2	e12-26-01	0.000
EB-:1	R	0.0	0.0	s2d10-11-01	1:610, 2:611, 3:612, 4:616#FCRC=1, 5:617, 6:618, 7:613, 8:614, 9:615 662		9632	EB-1	e12-18-01	0.000
EB-:1	R	0.0	0.0	s2d10-20-01	1:625, 2:626, 3:627, 4:619, 5:620#FCRC=1, 6:621, 7:622, 8:623, 9:624 662		9632	EB-2	e14-28-01	0.000
EE+:2	-	?	?	s2d10-17-01	1:648, 2:649, 3:650, 4:654, 5:646, 6:647, 7:651, 8:652, 9:653 664			EE+	e12-19-01	0.000
EE-:3	W	100.0	0.0	s2d10-09-01	1:606, 2:607, 3:608, 4:609, 5:601, 6:W:100.0%W602<80.0%9605, 7:603, 8:604, 9:605 661	9605	9632	EE-	e12-10-01	0.000
GCT:11	-	?	?	s1d06-29-01	1:745					
GT:33	R	0.0	0.0	s1d06-27-01	1:812, 2:813 811		9632	TRG	e15-22-01	0.000
RETRI:4	-	?	?	-	810	-	-			
TIBTID:24	R	0.0	0.0	s1d06-04-01	1:149, 2:150, 3:151, 4:152, 5:153, 6:154, 7:155, 8:156, 9:157, 10:158, 11:159, 12:160, 13:161, 14:162, 15:163		9632	TIBTID4	e12-30-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-05-01	1:134, 2:135, 3:136, 4:137, 5:138, 6:139, 7:140, 8:141, 9:142, 10:143, 11:144, 12:145, 13:146, 14:147, 15:148		9632	TIBTID5	e13-10-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-07-01	1:118,74, 2:119,75, 3:120,76, 4:121,77, 5:122,78, 6:123,79, 7:124,80, 8:125,81		9632	TIBTID6a	e13-15-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-07-01	9:126,82, 10:127,83, 11:128,84, 12:129,85, 13:130, 14:131, 15:132, 16:133		9632	TIBTID6b	e15-13-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-01-01	1:86,87, 2:88,89, 3:90,91, 4:92,93, 5:94,95, 6:96,97, 7:98,99, 8:100,101		9632	TIBTID1a	e13-16-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-01-01	9:102, 10:104,105, 11:106,107, 12:108,109, 13:110,111, 14:112,113, 15:114,115, 16:116,117		9632	TIBTID1b	e13-17-01	0.000
TIBTID:24	R	0.0	0.0	s1d06-02-01	1:50, 2:51, 3:52, 4:53, 5:54, 6:55, 7:56, 8:57, 9:58, 10:590, 11:60, 12:61	0	9632	TIBTID2	e13-22-01 [59]	0.000
TIBTID:24	R	0.0	0.0	s1d06-03-01	1:62, 2:63, 3:64, 4:65, 5:66, 6:67, 7:68, 8:69, 9:70, 10:71, 11:72, 12:73		9632	TIBTID3	e13-25-01	0.000
Summary				frlpc	geoSlot:SrcId	min Trg	max Trg	FB Name	RU	rate (kHz) thru
										0.000

… DAQ shifters may not be experienced enough

… pinpointing the problem may take (considerable) time

… an on-call expert may need to be called …

… at any time of the day

New DAQ Expert

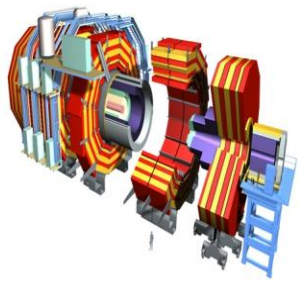

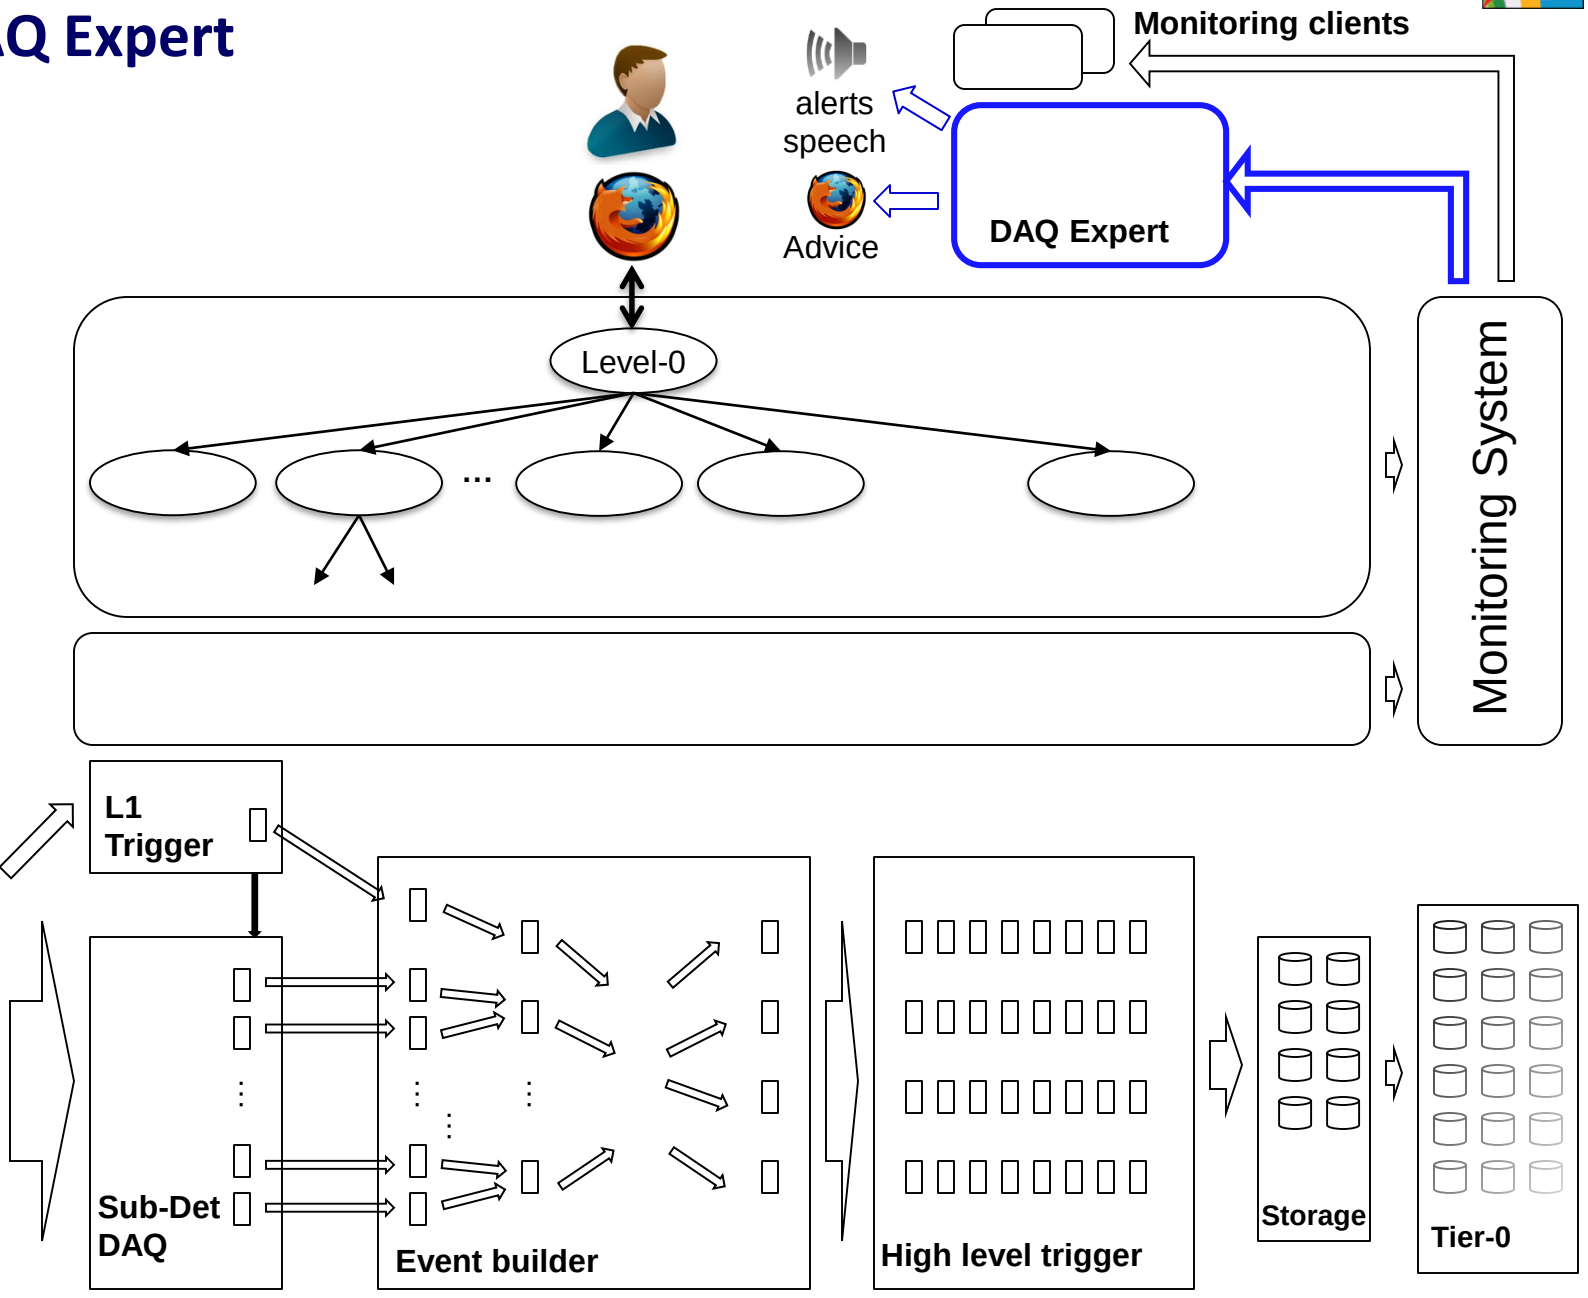

Solution

Aggregate monitoring data → Identify the problem & recovery→ Advise operators

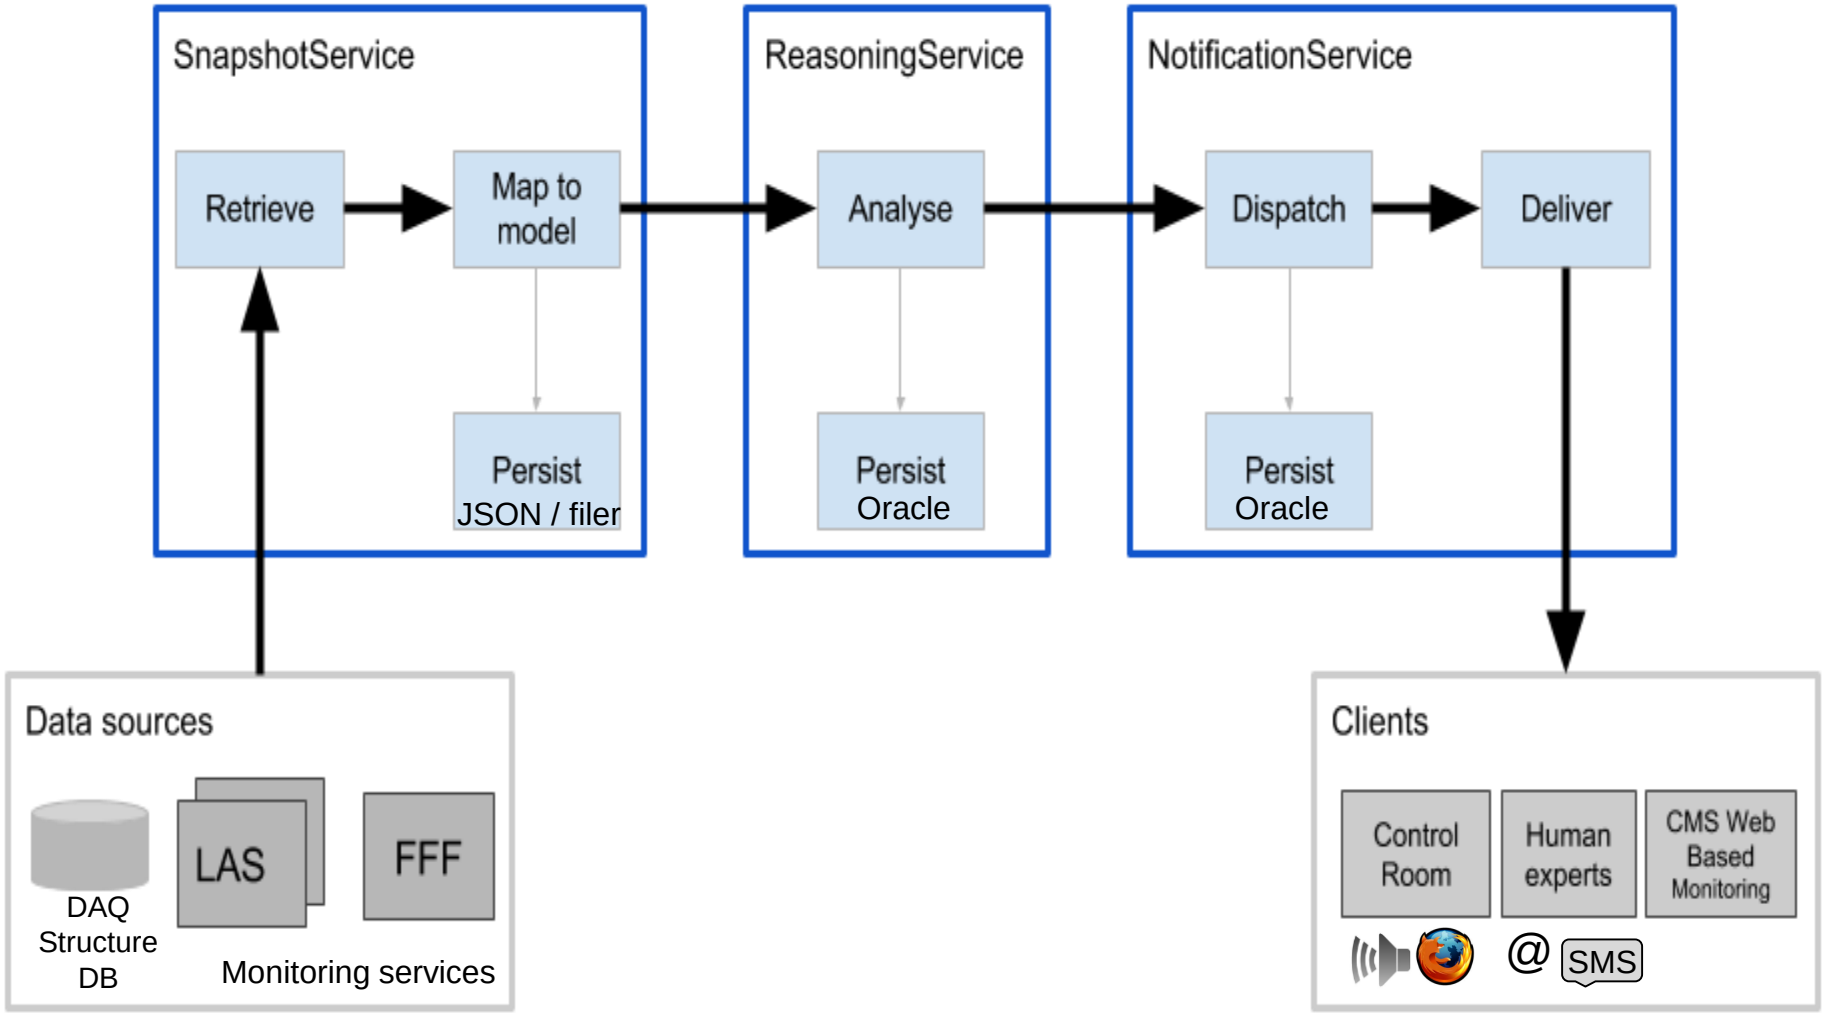

Reasoning

- Expert knowledge encapsulated in logic modules (LM)
- Each LM defines a condition
  - Satisfy method returning true or false
- Input data
  - Current snapshot of monitoring data
  - Output of other logic modules

Reasoning II

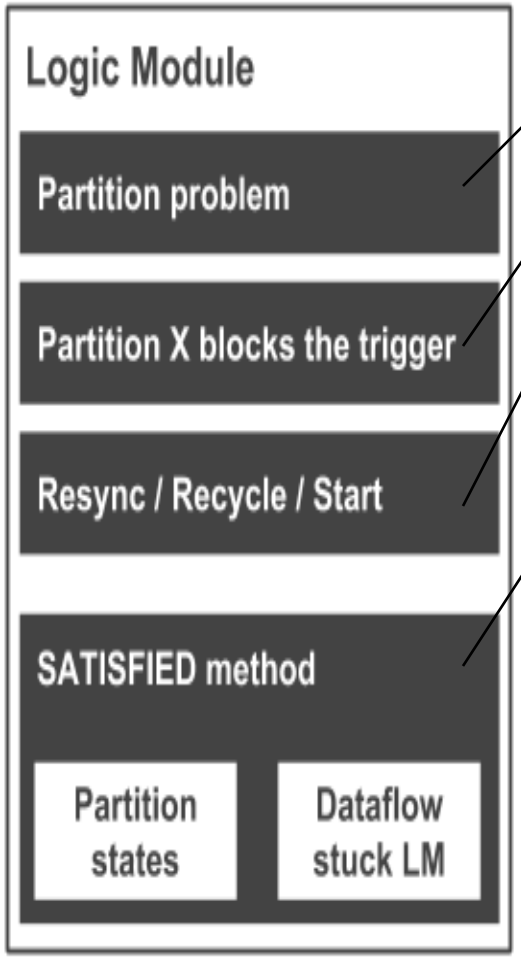


```java
public class PartitionProblem extends KnownFailure {

  public PartitionProblem() {
    this.name = "Partition problem";
    this.description = "Partition {{P}} in subsystem {{S}} is blocking triggers";
    this.action = new SimpleAction("Issue a TTCHardReset",
        "Problem not fixed: Stop current run & start a new run.",
        "Problem still not fixed: Call the DOC of {{S}}",
        "Problem fixed: Make an e-log entry.");
  }

  @Override
  public boolean satisfied(DAQ daq, Map<String, Boolean> results) {
    boolean result = false;
    if (!results.get(DataFlowStuck.class.getSimpleName()))
      return false;

    for (SubSystem subSystem : daq.getSubSystems()) {
      for (TTCPartition ttcp : subSystem.getTtcPartitions()) {
        TTSState currentState = getParitionState(ttcp);
        if (currentState == TTSState.OUT_OF_SYNC) {
          context.register("S", subSystem.getName());
          context.register("P", ttcp.getName());
          context.register("C", currentState.name());
          result = true;
        }
      }
    }
    return result;
  }
}
```

Order of LM execution defined by requirement graph

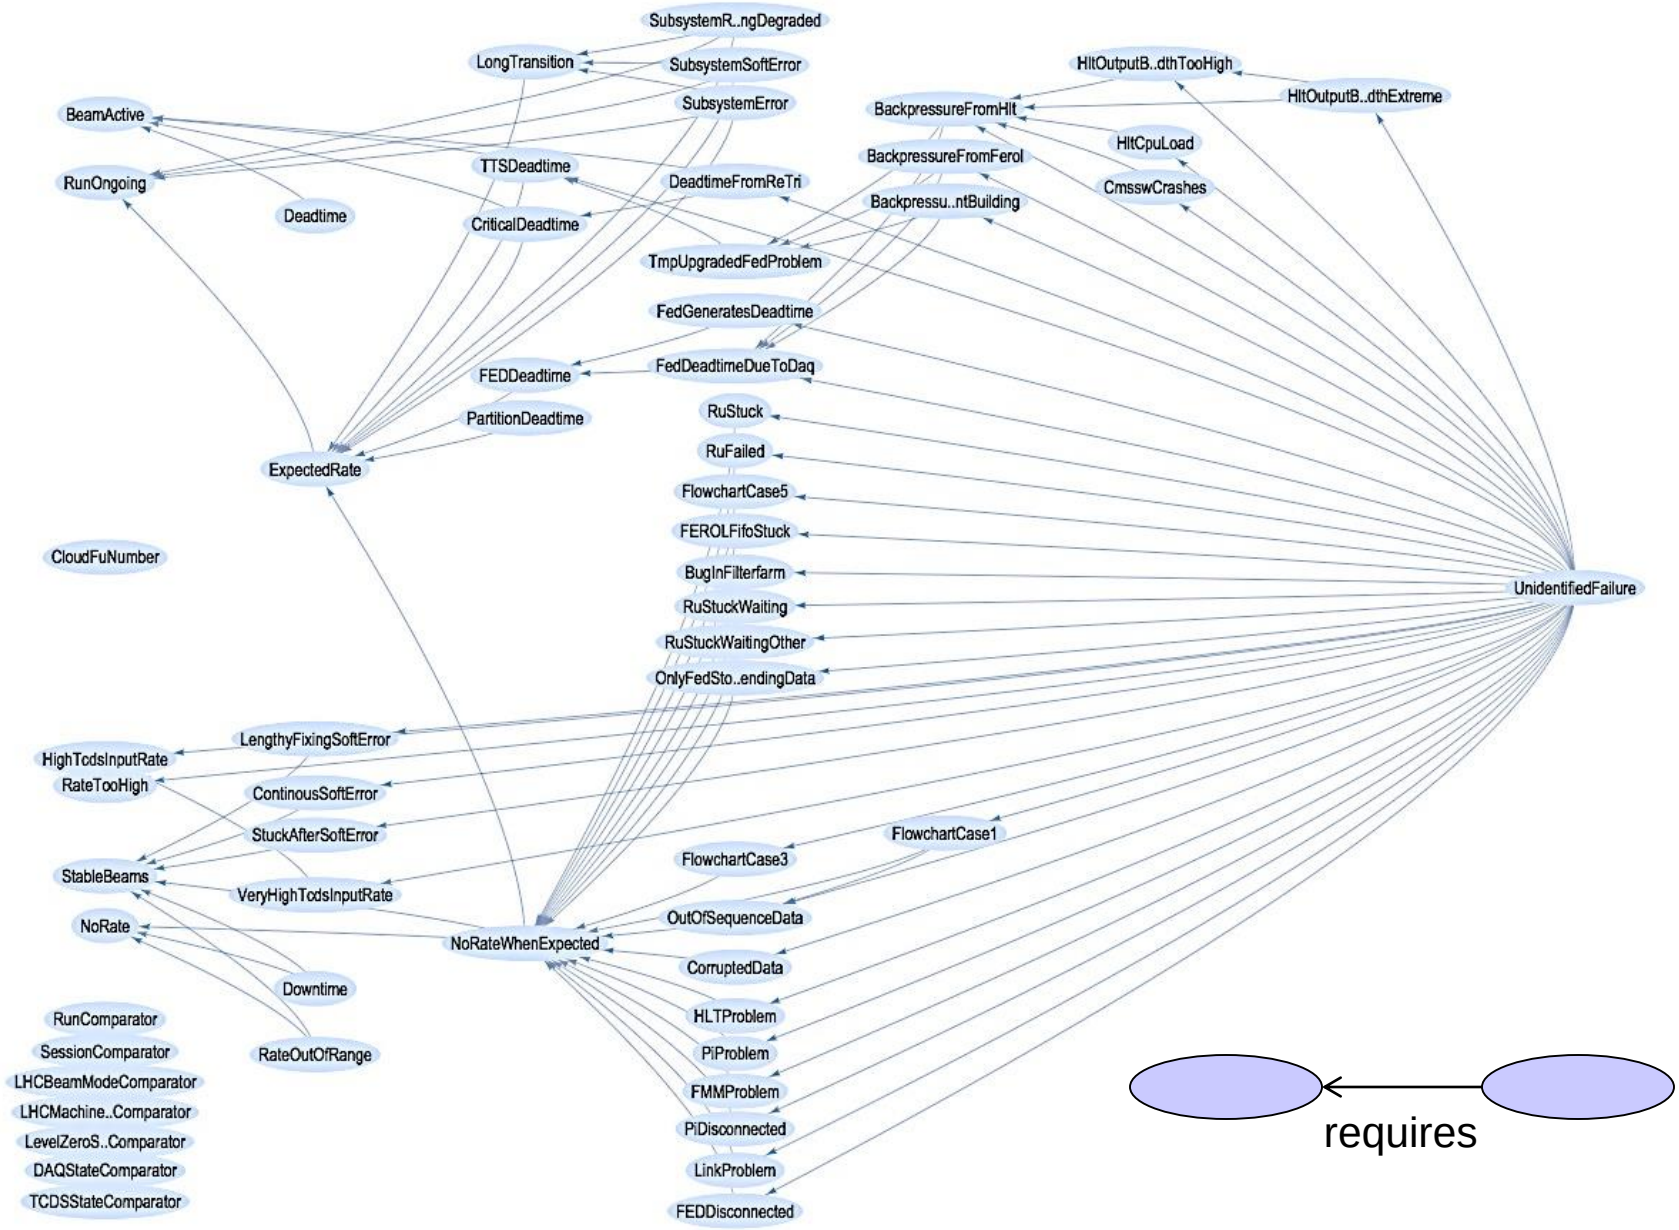

Dashboard: Current main problem(s) and history

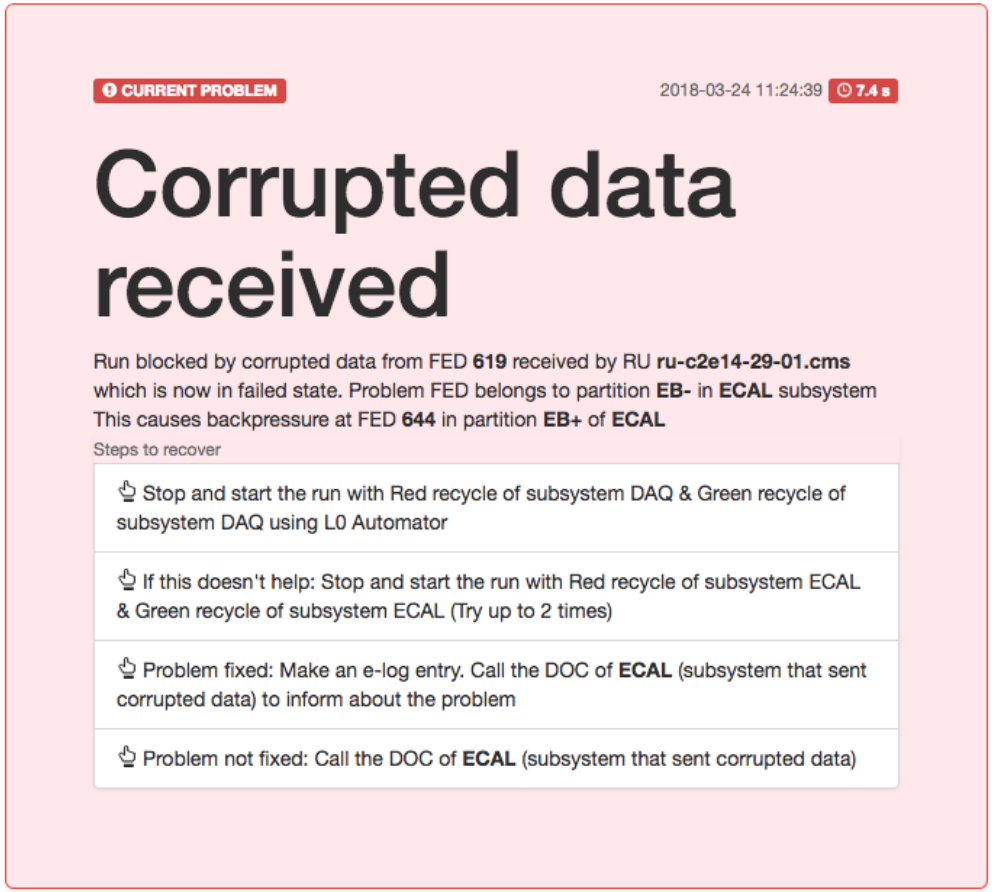

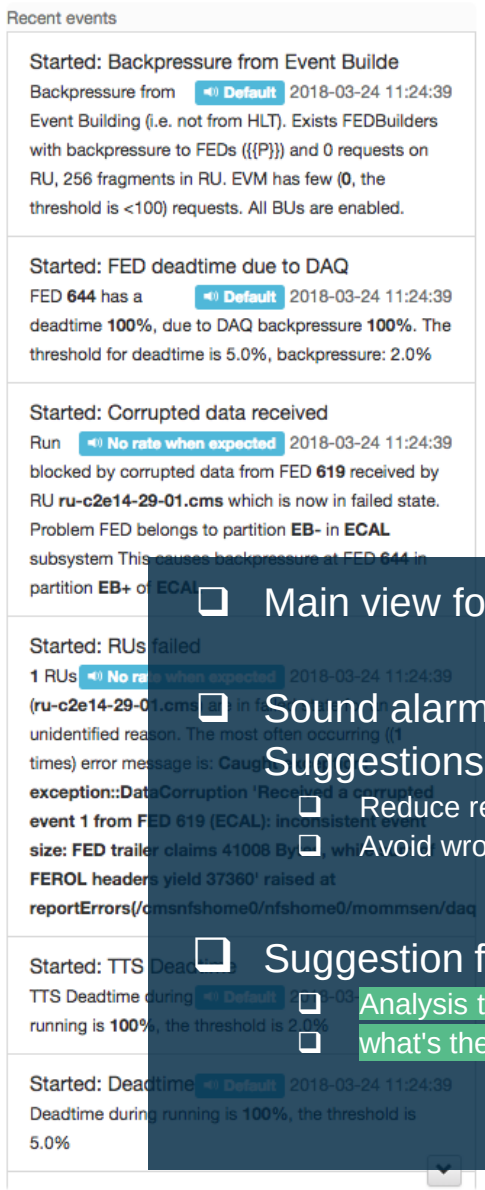

- Main view for control room
- Sound alarm and Suggestions to shifters:
  - Reduce reaction time
  - Avoid wrong decisions
- Suggestion format:
  - Analysis the problem
  - what's the best action to take

Browser

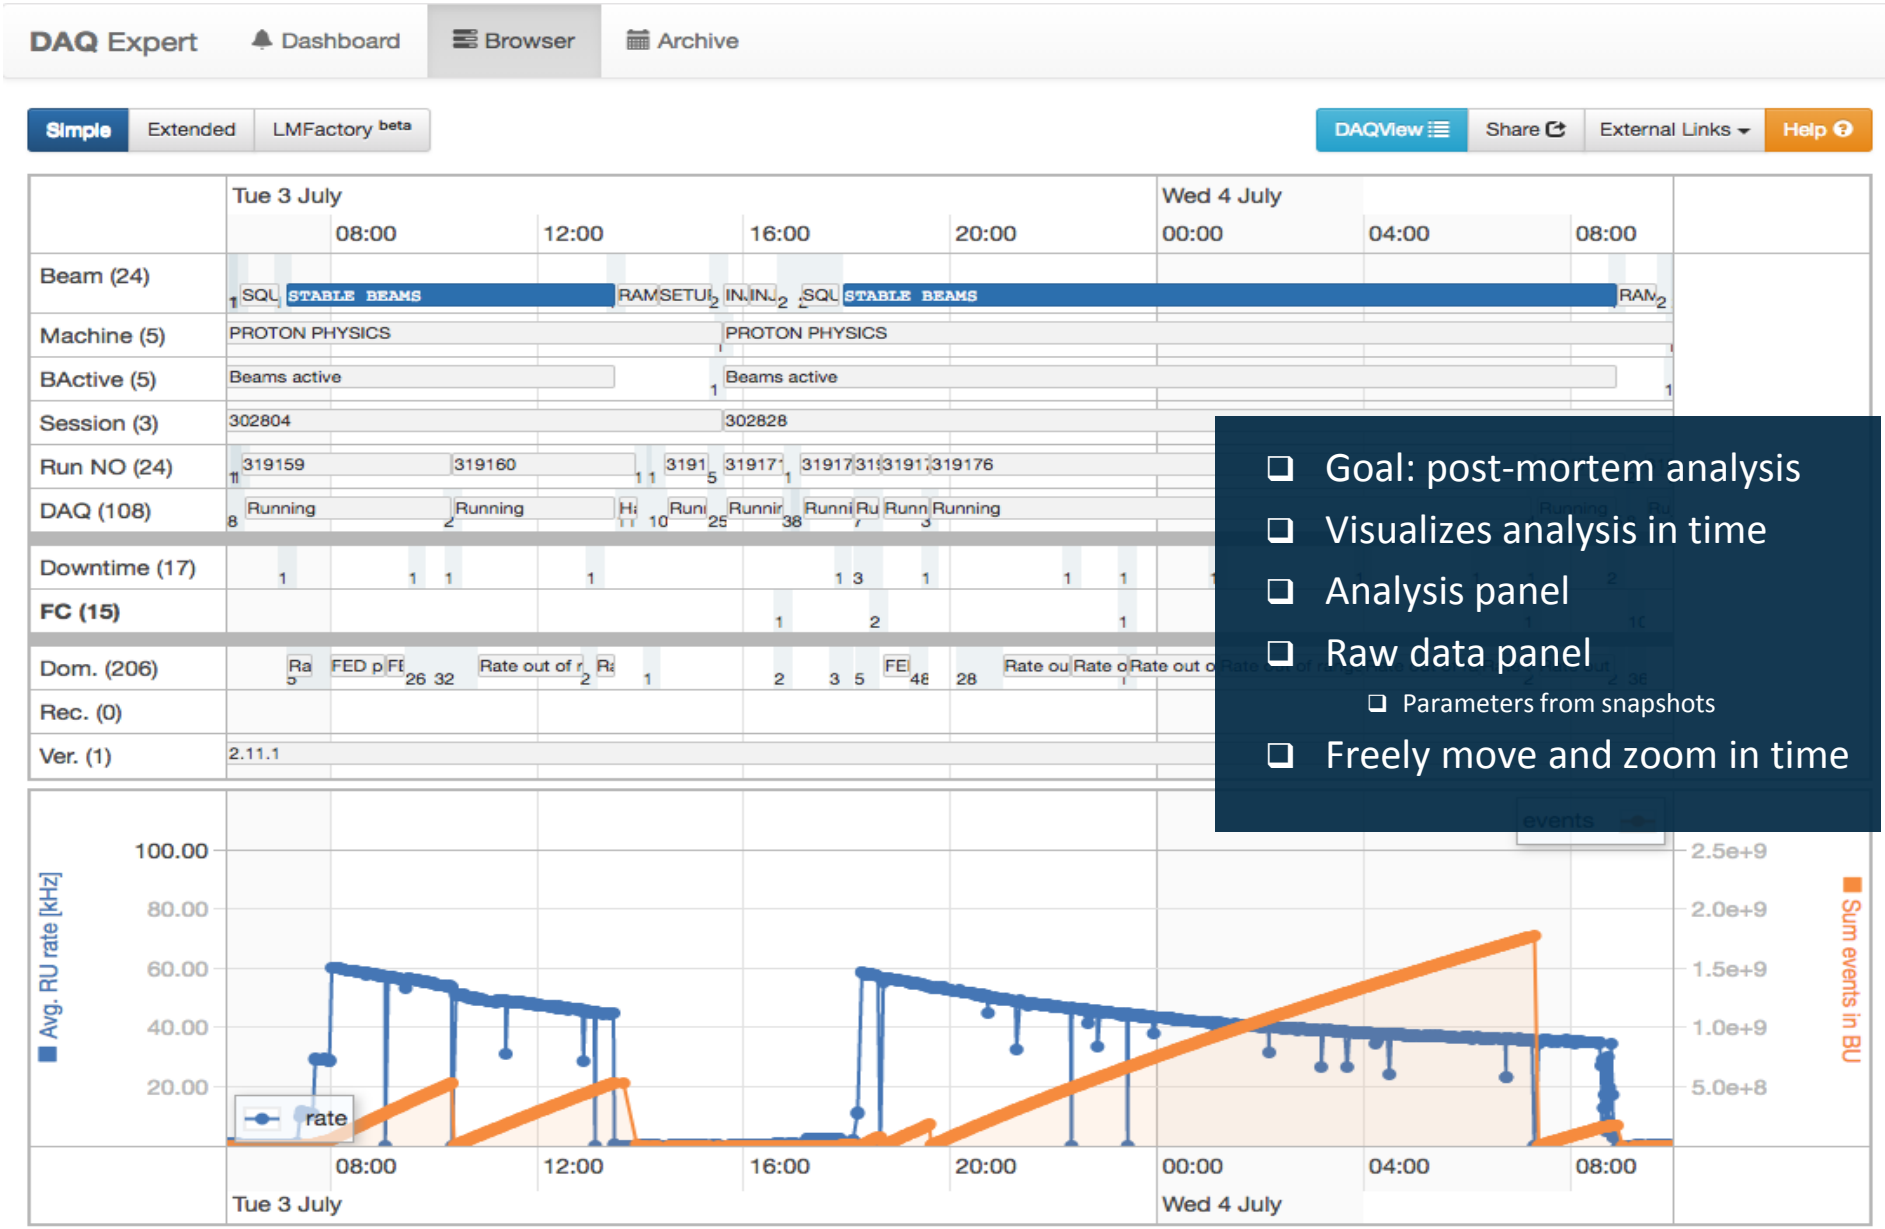
DAQ Expert
Dashboard
Browser
Archive
Simple
Extended
LMFactory beta
DAQView
Share
External Links
Help
Tue 3 July
Wed 4 July
08:00
12:00
16:00
20:00
00:00
04:00
08:00
Beam (24)
STABLE BEAMS
Machine (5)
PROTON PHYSICS
BActive (5)
Beams active
Session (3)
302804
302828
Run NO (24)
319159
319160
DAQ (108)
Running
Downtime (17)
FC (15)
Dom. (206)
Rec. (0)
Ver. (1)
2.11.1
Goal: post-mortem analysis
Visualizes analysis in time
Analysis panel
Raw data panel
Parameters from snapshots
Freely move and zoom in time
Avg. RU rate [kHz]
100.00
80.00
60.00
40.00
20.00
rate
events
Sum events in BU
2.5e+9
2.0e+9
1.5e+9
1.0e+9
5.0e+8

Archive

DAQ Expert Dashboard Browser Archive

source 5 selected event type Start date range 2018-07-03 00:00 - 2018-07-03 23:00 Help

Date	Event Type	Sender	Title	Displayed message	Source LM	Audible
2018-07-03 18:36:51	Start	Expert *2.11.1*	Started: FED stuck	TTCP TEC+ of TRACKER subsystem is blocking triggers, it's in WARNING TTS state, The problem is caused by FED 274 in WARNING	Fed stuck	No rate when expected
2018-07-03 18:35:54	Start	Expert *2.11.1*	Started: FED stuck	TTCP TEC+ of TRACKER subsystem is blocking triggers, it's in WARNING TTS state, The problem is caused by FED 274 in WARNING	Fed stuck	No rate when expected
2018-07-03 16:45:37	Start	Expert *2.11.1*	Started: Partition problem	Partition EB- in ECAL subsystem is in ERROR TTS state. It's blocking triggers.	Partition problem	No rate when expected
2018-07-03 05:54:50	Start	Expert *2.11.1*	Started: Partition problem	Partition EB- in ECAL subsystem is in ERROR TTS state. It's blocking triggers.	Partition problem	No rate when expected
2018-07-03 04:38:19	Start	Expert *2.11.1*	Started: Partition problem	Partition [EB+, EB-, EE+, EE- and 3 more] in [ECAL, TRACKER] subsystem is in [ERROR, OUT_OF_SYNC] TTS state. It's blocking triggers.	Partition problem	No rate when expected
2018-07-03 04:38:15	Start	Expert *2.11.1*	Started: FED stuck	TTCP [CALOL1, MUTF] of TRG subsystem is blocking triggers, it's in [BUSY, WARNING] TTS state, The problem is caused by FED [1354, 1356, 1358, 1380-1381] in [BUSY, WARNING]	Fed stuck	No rate when expected

first « 1 » last

6 entries (1 pages)

- ❑ Goal: browse all generated notifications
- ❑ Filter by key fields
- ❑ Inspect link

Technologies used

- Presentation layer
  - Web application
  - Javascript, VisJS
  - Bootstrap
  - React JS
- Backend
  - Micro service architecture
  - RESTFul services
  - Tomcat
  - Hibernate
  - Oracle
  - Jackson JSON serializer

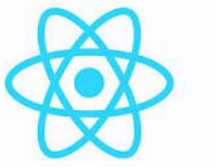

Determining the root cause

Determining the root cause

- Logic modules are not exclusive
- Sometimes a single problem satisfies O(10) conditions (logic modules)
- Task
  - Differentiate problem from symptom
  - But: depending on the situation, the same condition can be a problem or a symptom
- Solution: causality graph
  - Defines conditions that are possible causes for a condition
  - Relation defined inside the logic modules
    - `declareCause(LogicModuleRegistry.CorruptedData);`
    - `declareAffected(LogicModuleRegistry.TTSDeadtime);`

Causality graph

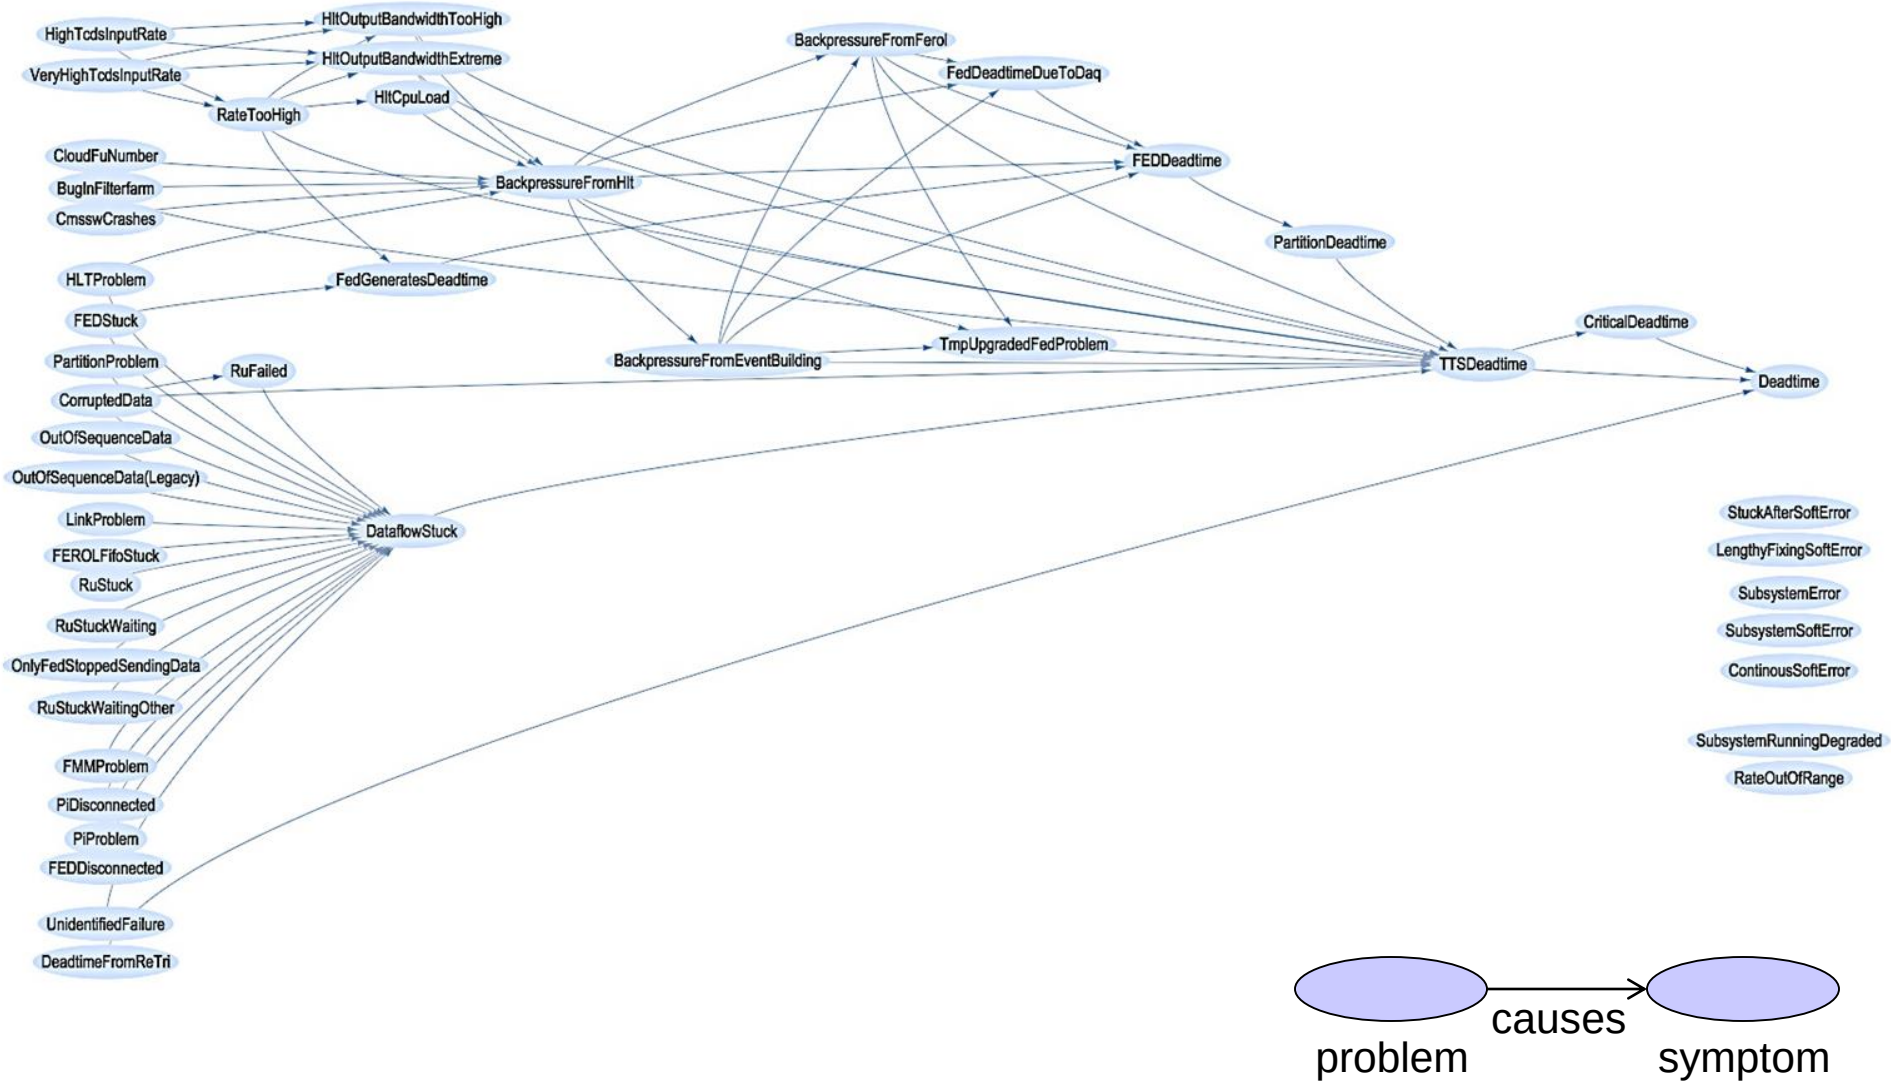

Causality graph: Simple example

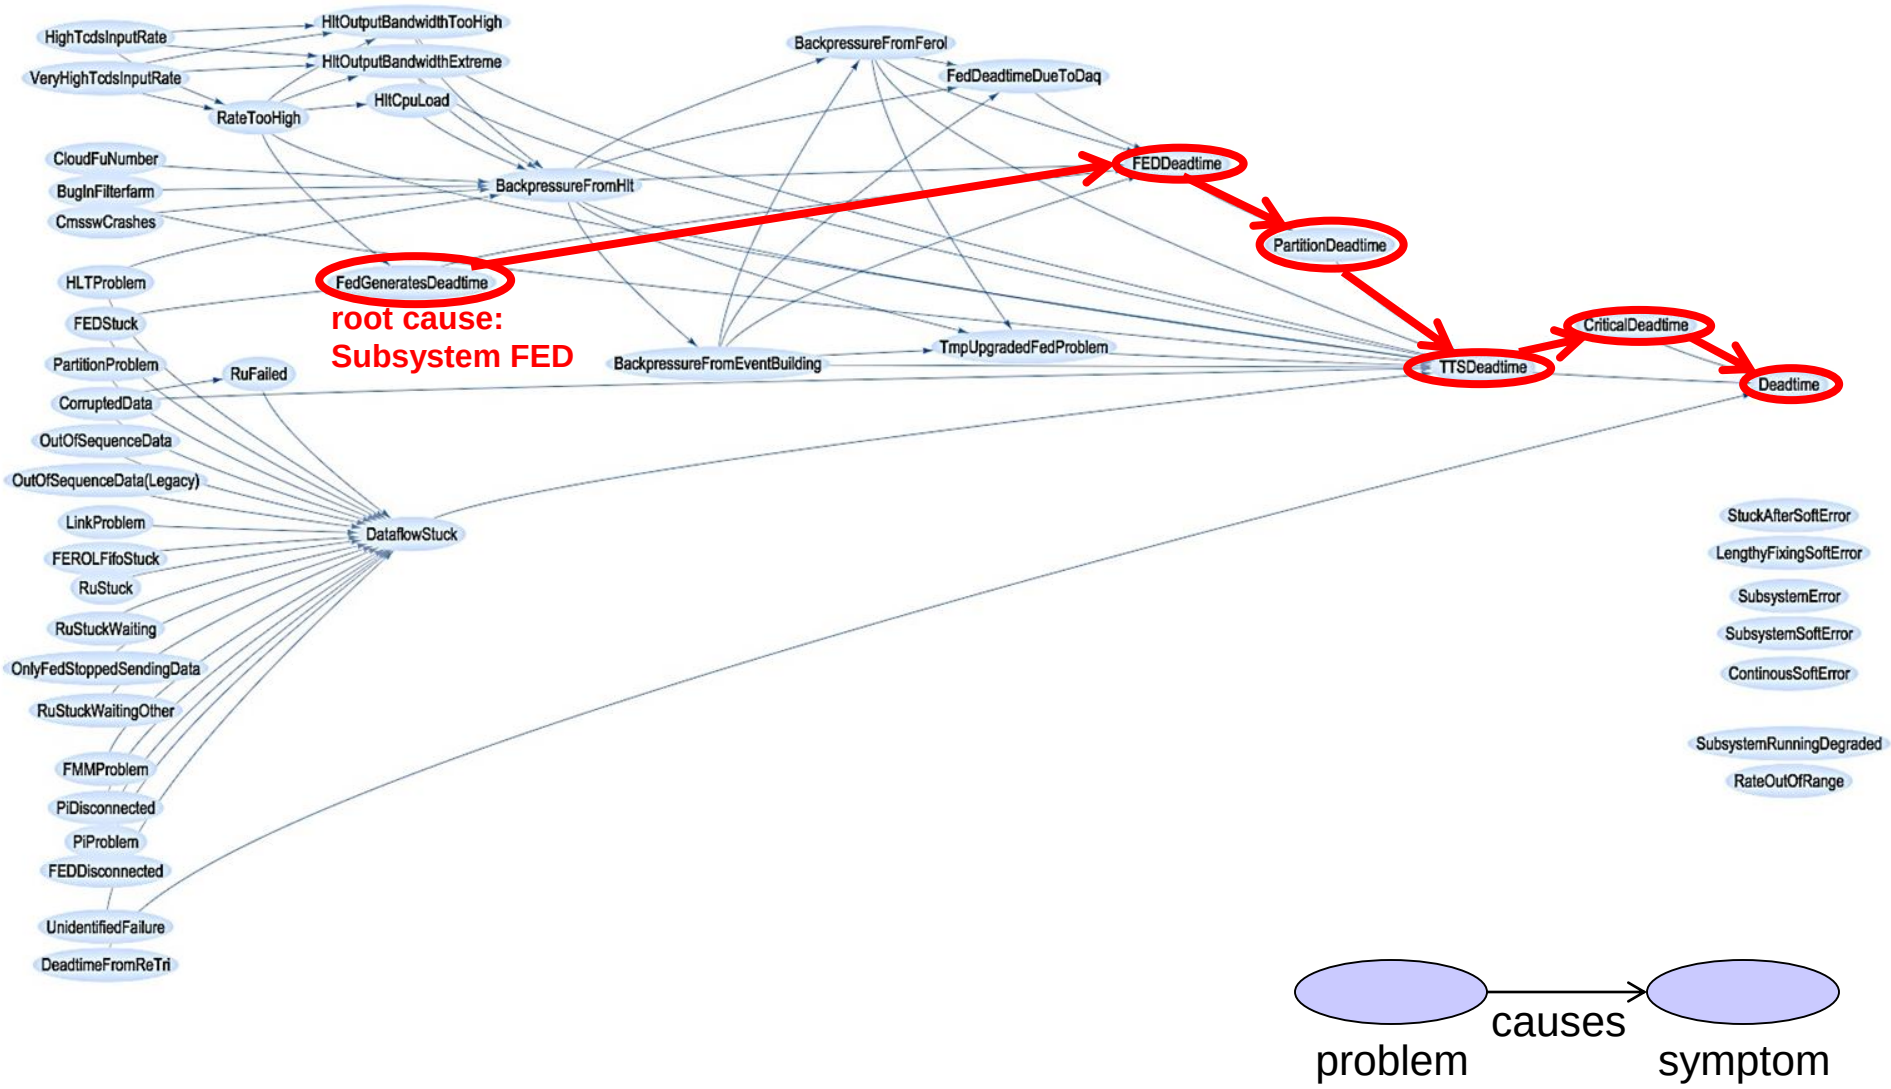

Causality graph: Simple example II

root cause:
trigger rate

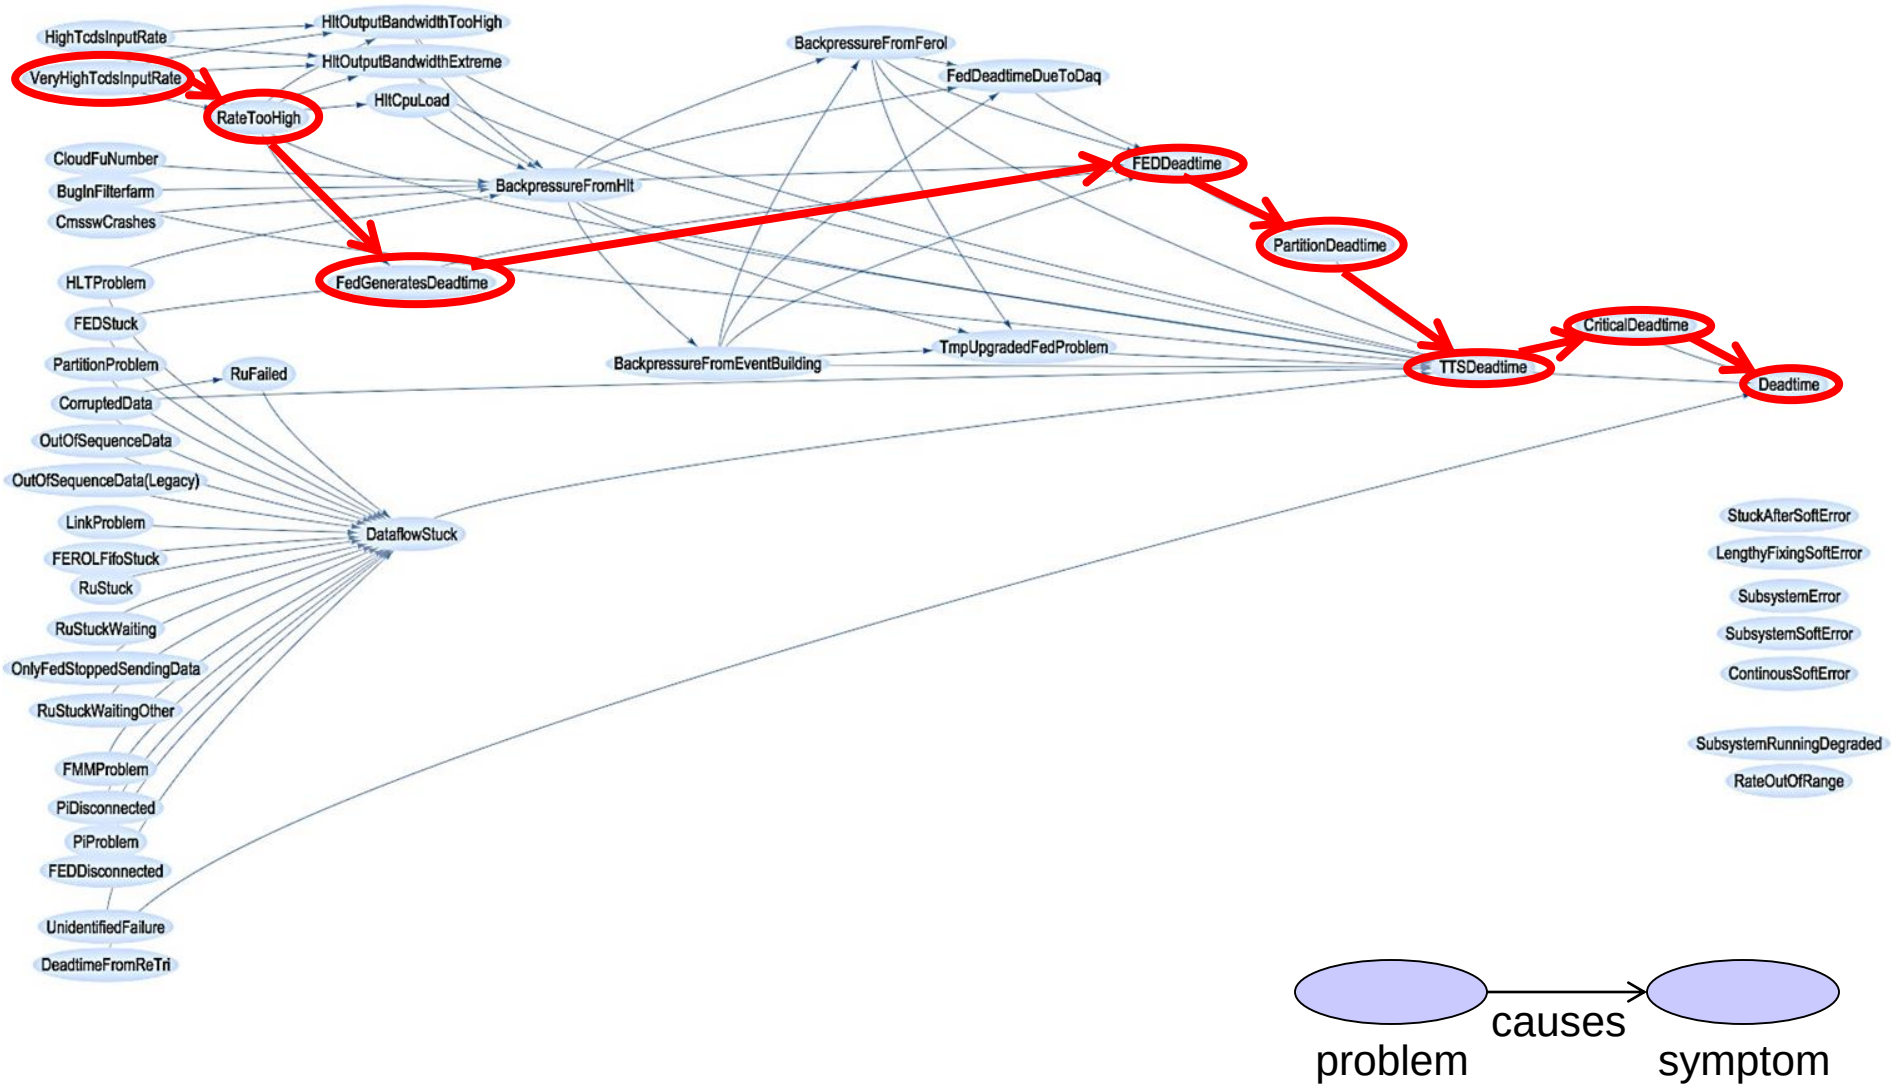

Causality graph: Simple example III

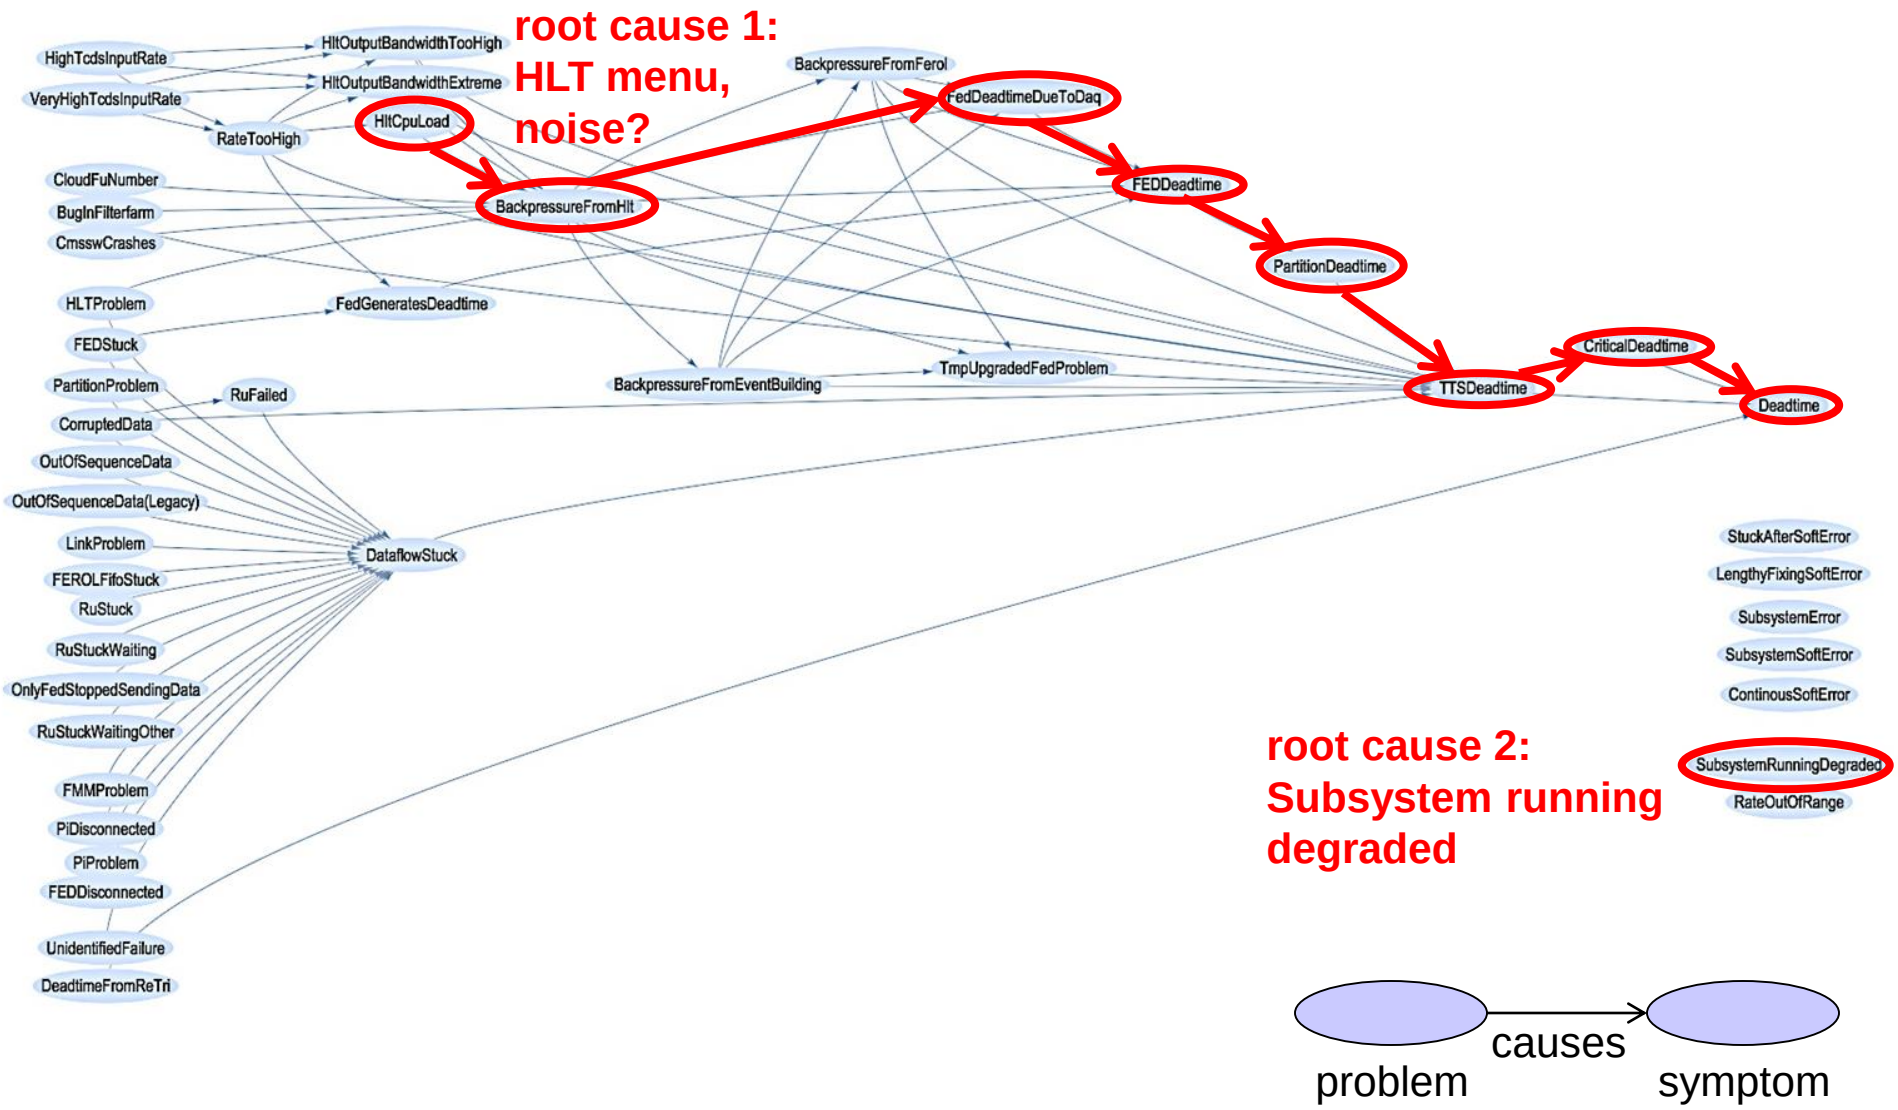

Causality graph: Example IV

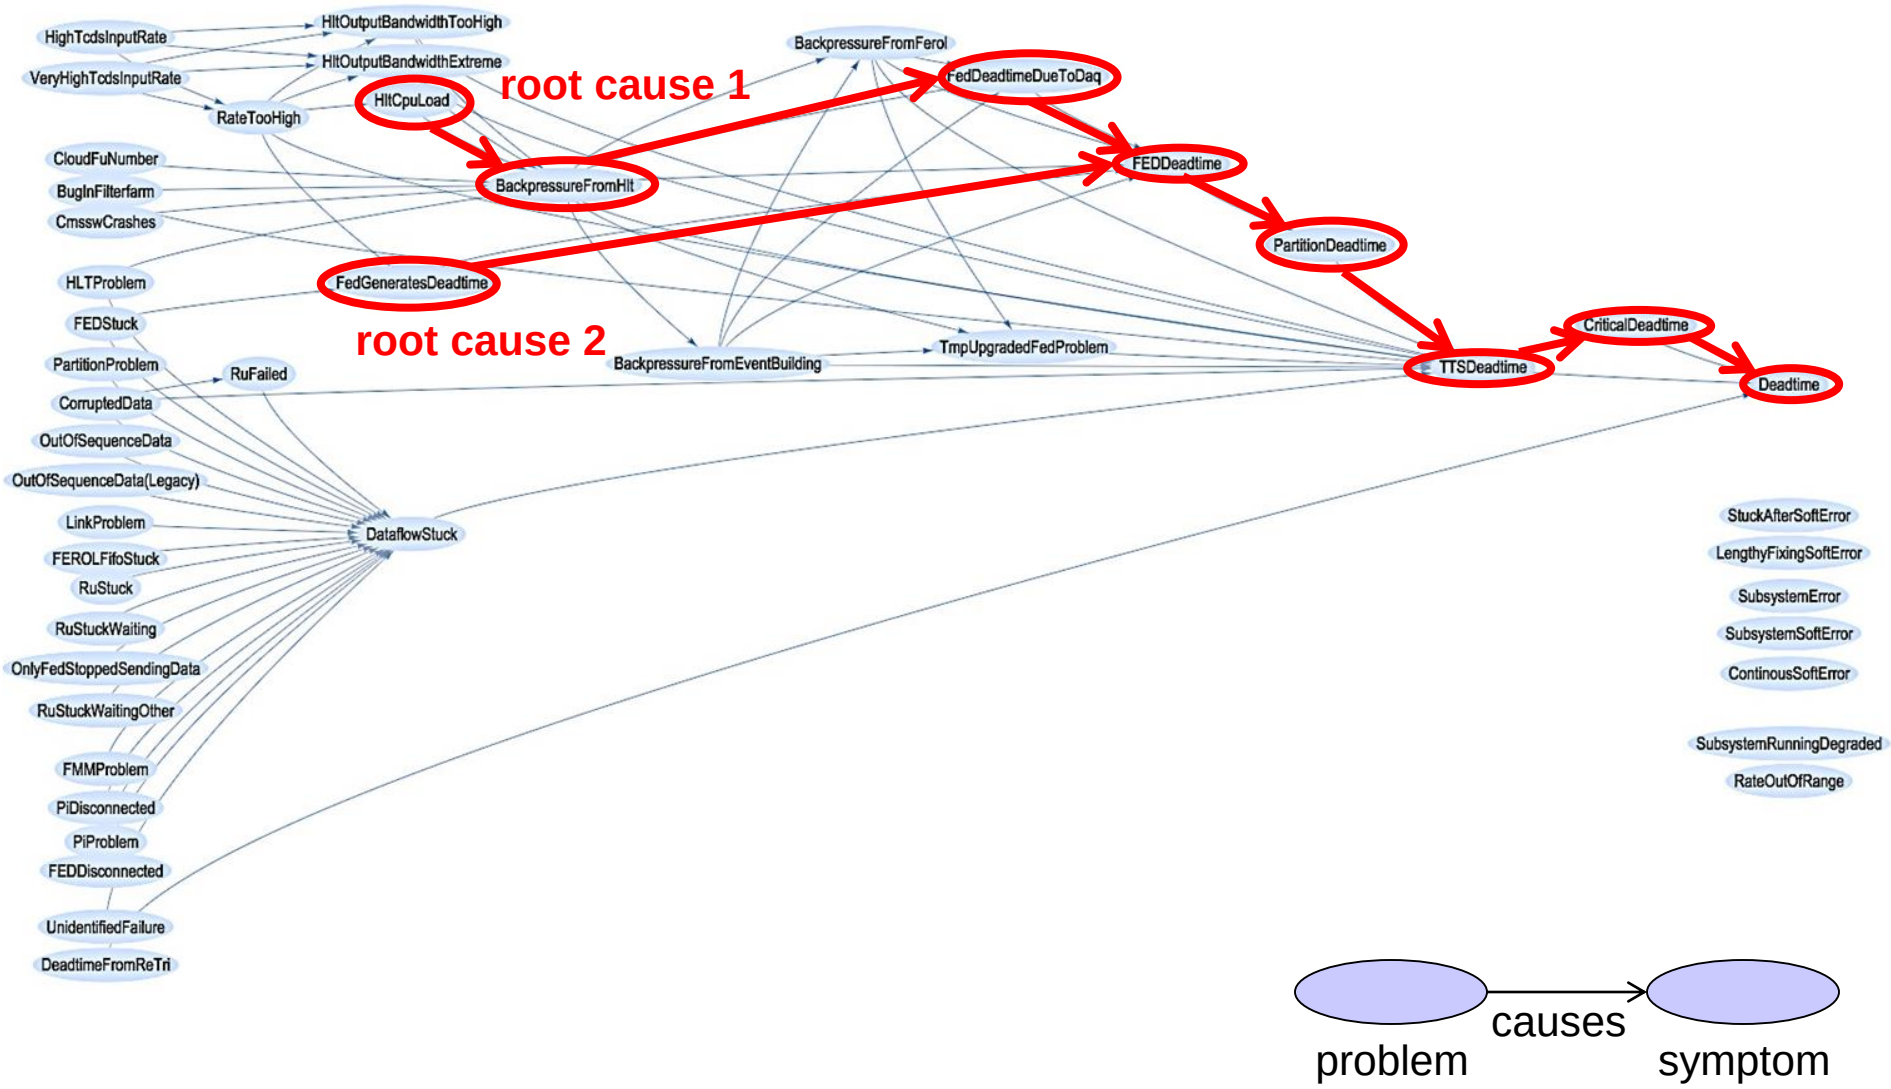

Causality graph: Example V

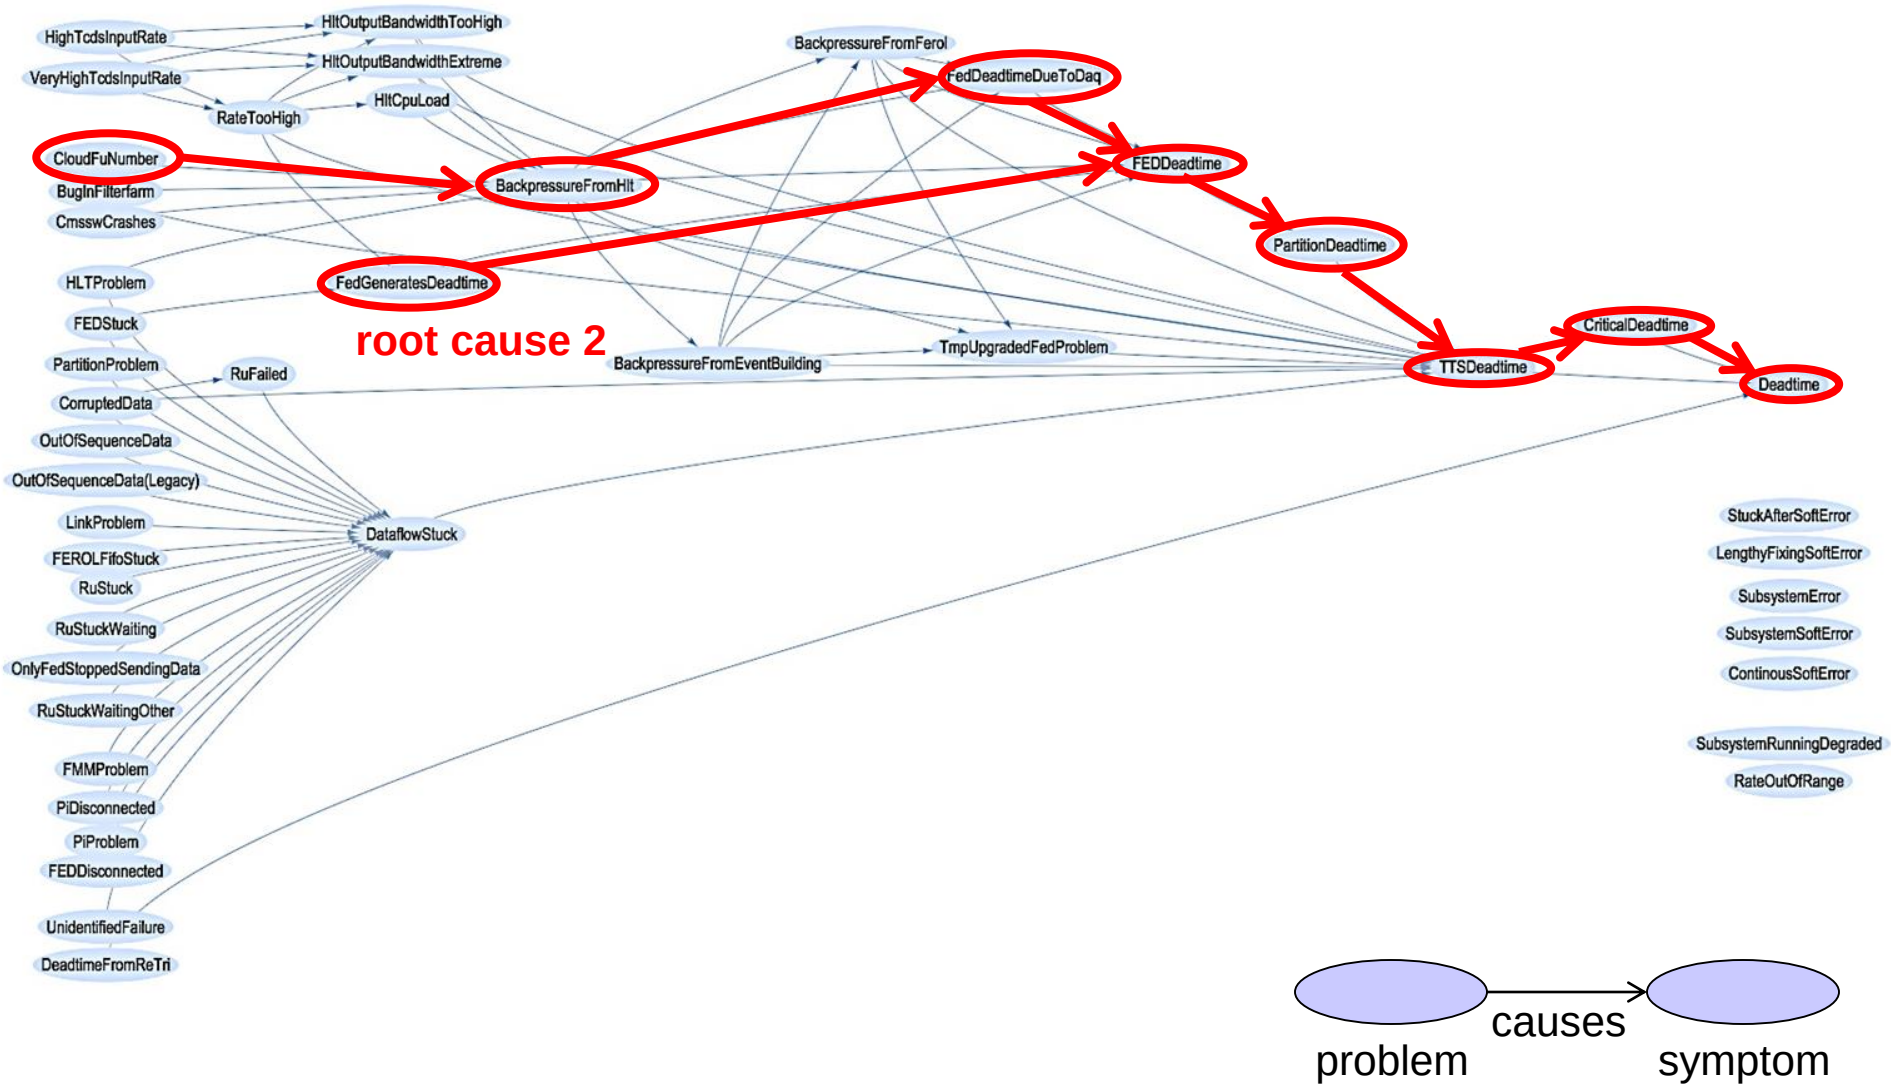

Root cause as found through causality graph

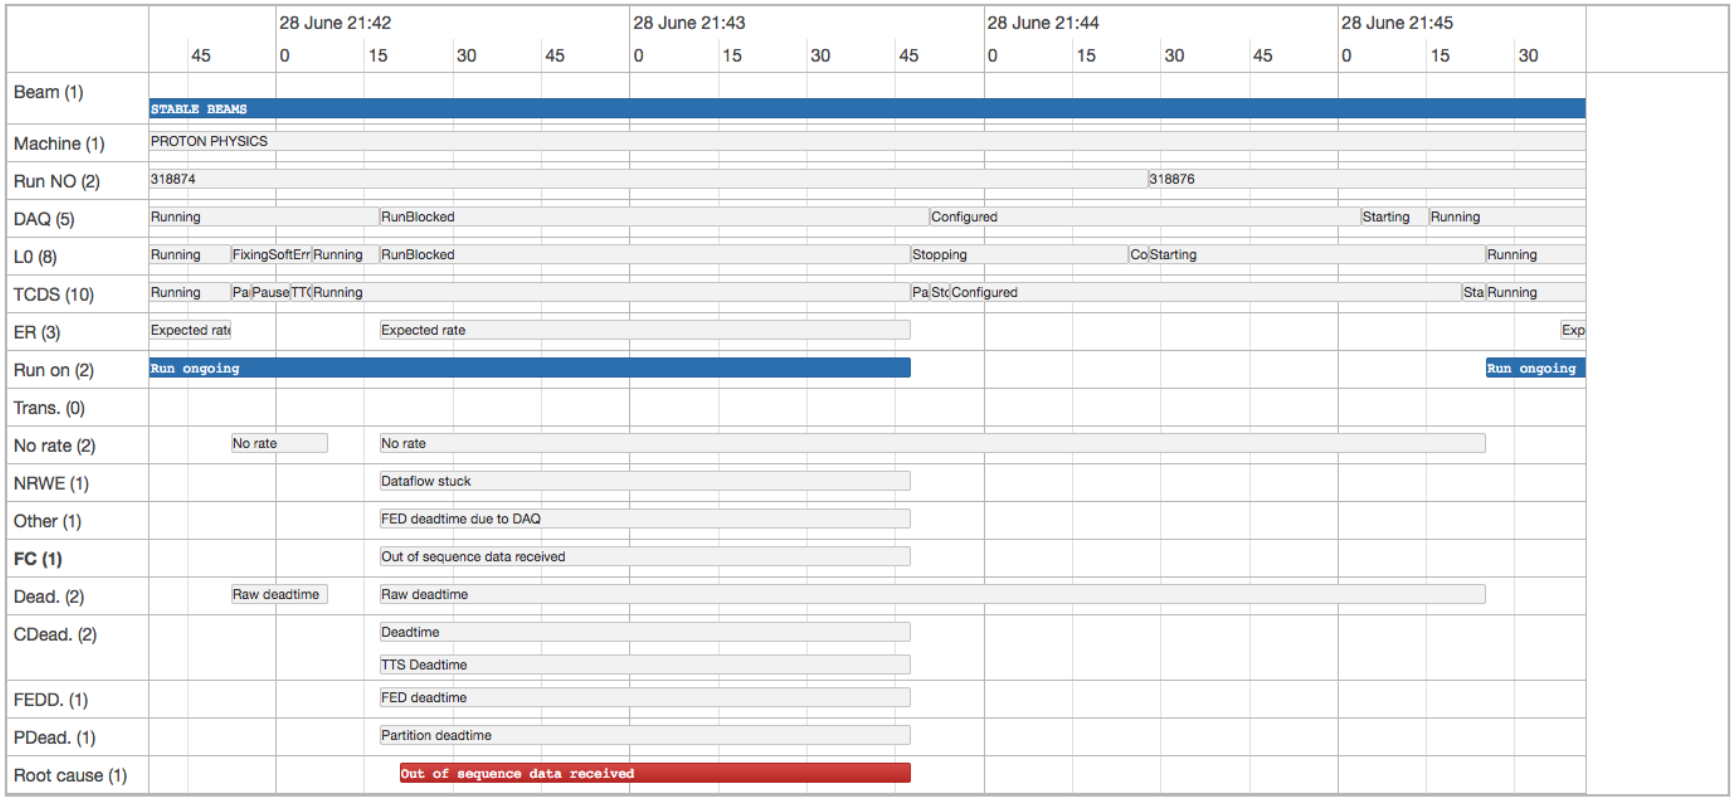

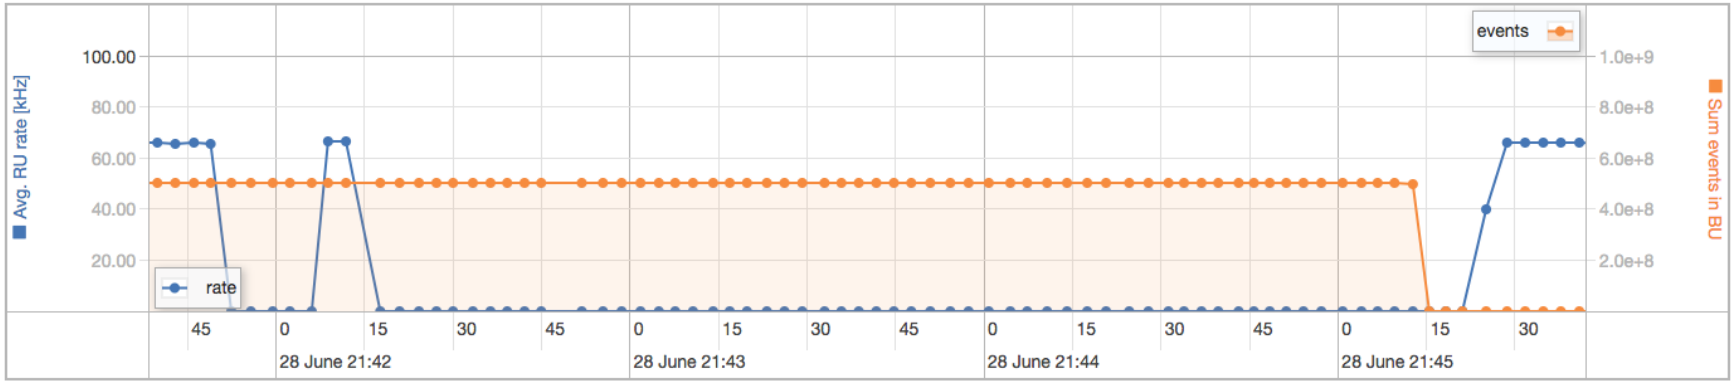

Recovery driven by DAQ Expert

Driving Recovery from the DAQ Expert

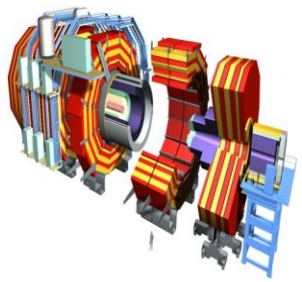

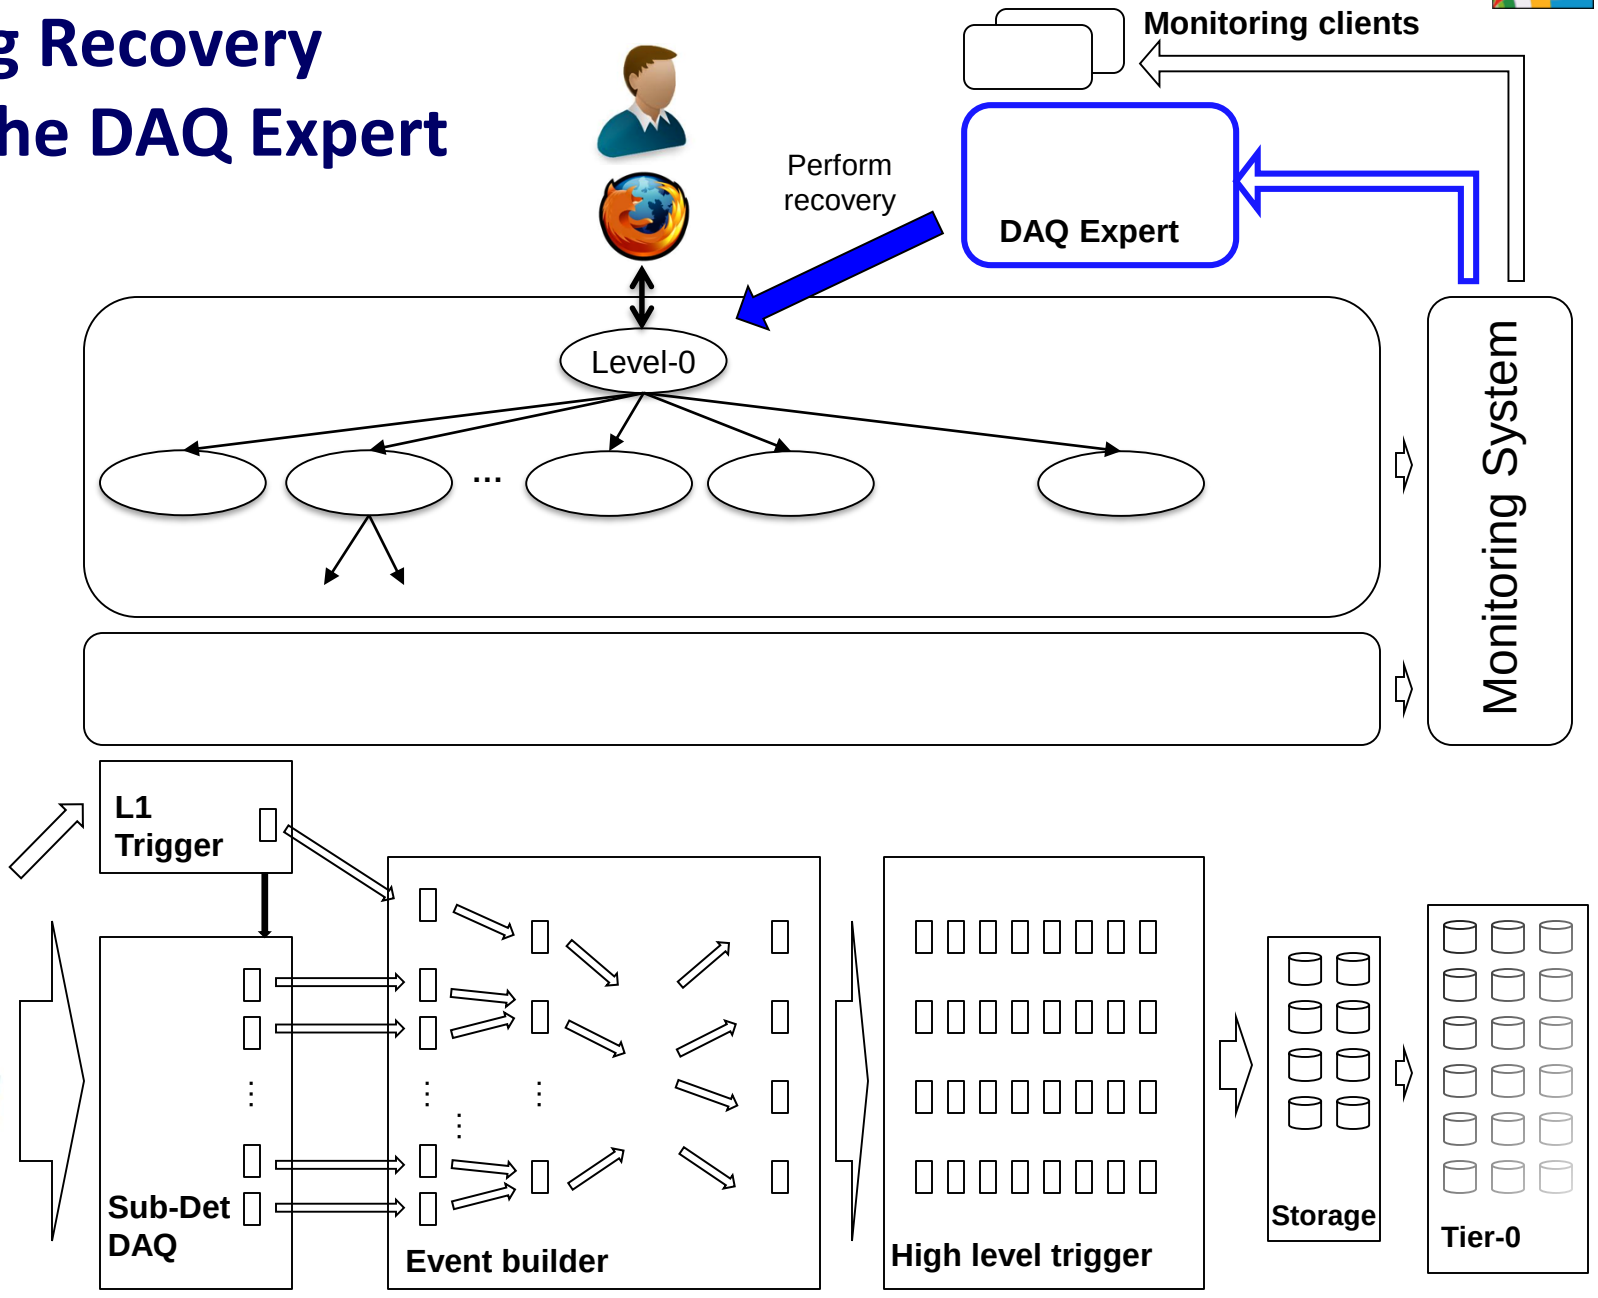

Driving Recovery from the DAQ Expert

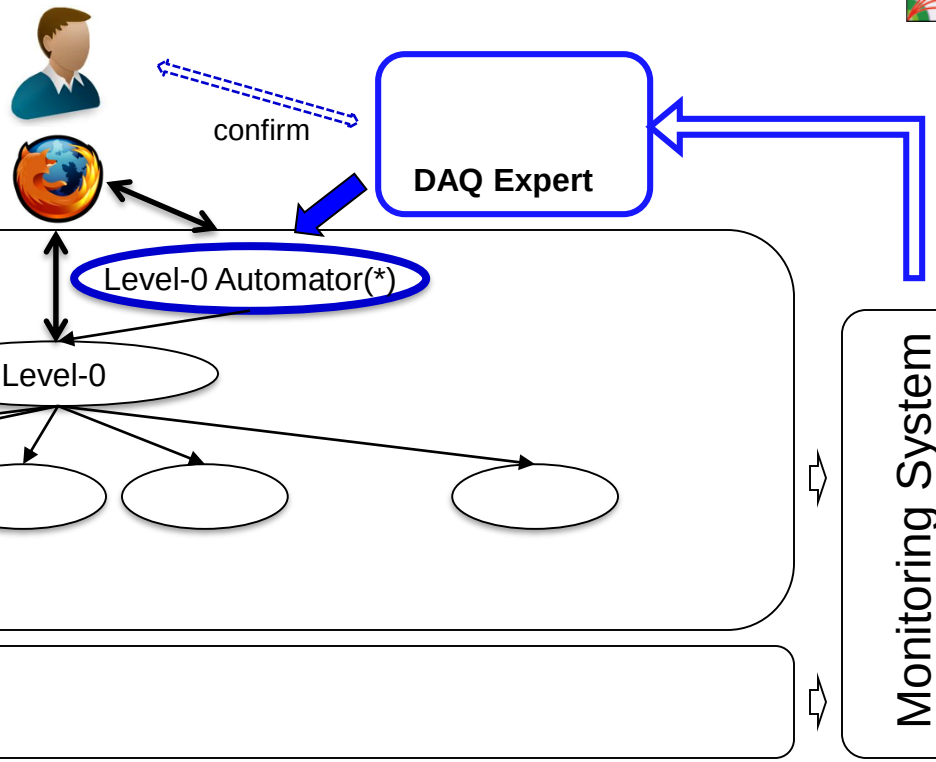

(*) CHEP 2016: H. Sakulin et. al., “New operator assistance features in the CMS Run Control System”

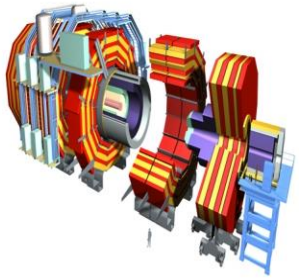

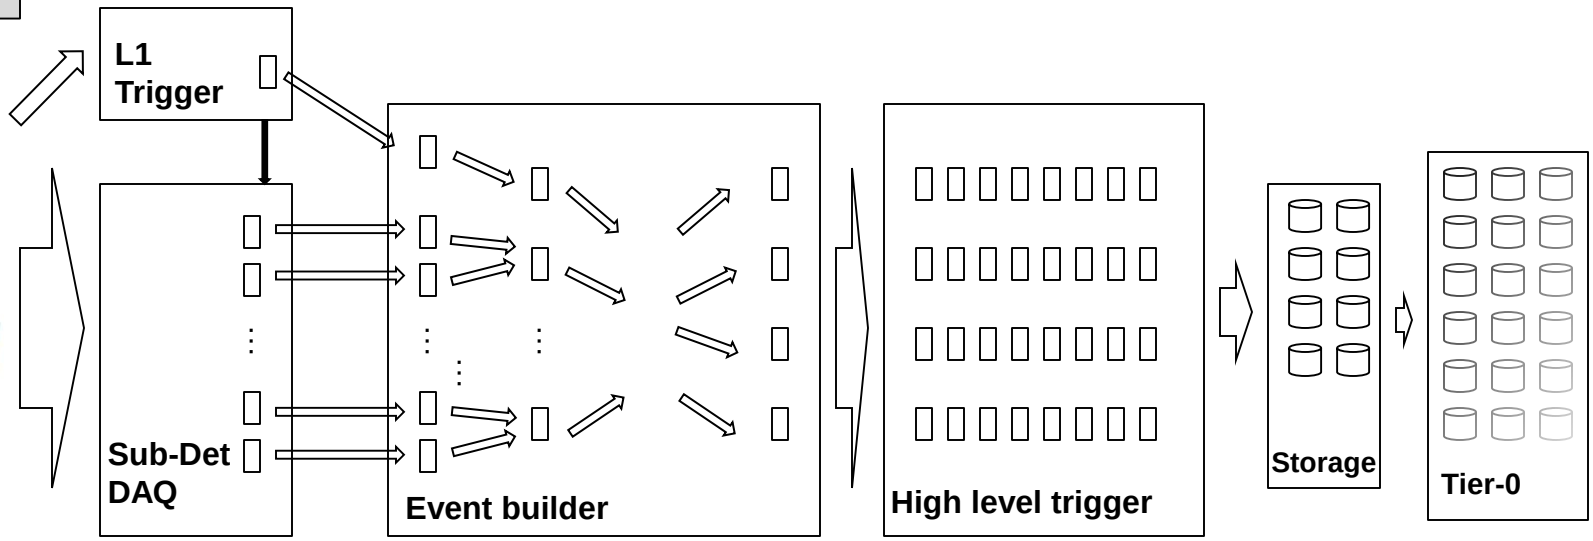

Recovery driven by DAQ Expert

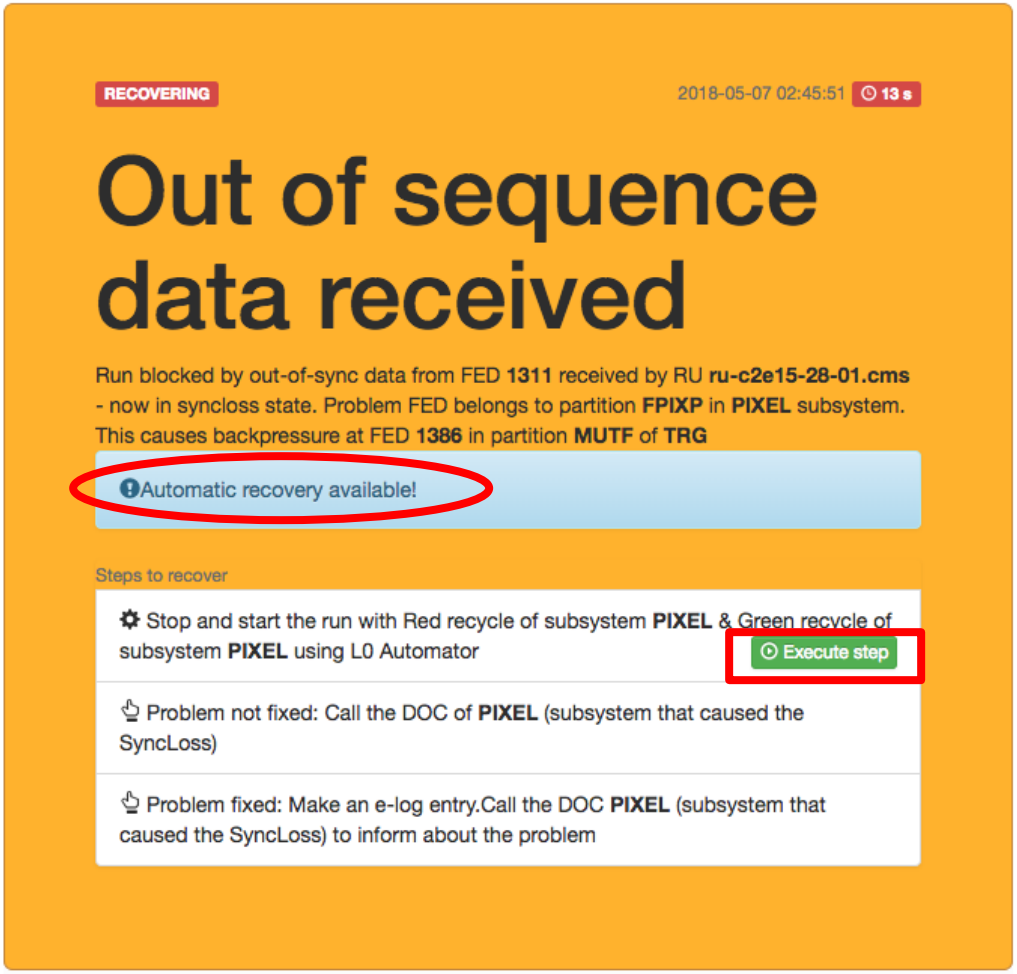

Recent problems

TTS Deadtime 2018-05-07 02:45:51 13 s
TTS Deadtime during running is **100%**, the threshold is 2.0% Show steps

Deadtime 2018-05-07 02:45:51 13 s
Deadtime during running is **100%**, the threshold is 5.0% Show steps

Recent events

Started: Upgraded FED problem (TMP)
High backpressure on fed(s) [1386, 1404] in partition(s) [MUTF, UGT] in subsystem(s) TRG is (last: 100%, avg: 79.3%, min: 39%, max: 100%) the threshold is 2.0%. This does not indicate a problem with these FEDs. This condition is only used as a basis for other backpressure analysis since upgraded FEDs have no deadtime monitoring. For legacy FEDs the deadtime is the basis for backpressure analysis. 2018-05-07 02:45:51

Started: Out of sequence data received
Run No rate when expected 2018-05-07 02:45:51 blocked by out-of-sync data from FED **1311** received by RU **ru-c2e15-28-01.cms** - now in syncloss state. Problem FED belongs to partition **FPIXP** in **PIXEL** subsystem. This causes backpressure at FED **1386** in partition **MUTF** of **TRG**

Started: TTS Deadtime
TTS Deadtime Default 2018-05-07 02:45:51 during running is **100%**, the threshold is 2.0%

Started: Deadtime
Deadtime during Default 2018-05-07 02:45:51 running is **100%**, the threshold is 5.0%

Started: Dataflow stuck
There No rate when expected 2018-05-07 02:45:51 is no rate when expected. The Data flow is stuck.

DAQ state: RunBlocked
New DAQ state DAQ state 2018-05-07 02:44:53 identified

Level Zero State: RunBlocked
New Level zero state identified 2018-05-07 02:44:53

Started: Run ongoing

Recovery driven by DAQ Expert

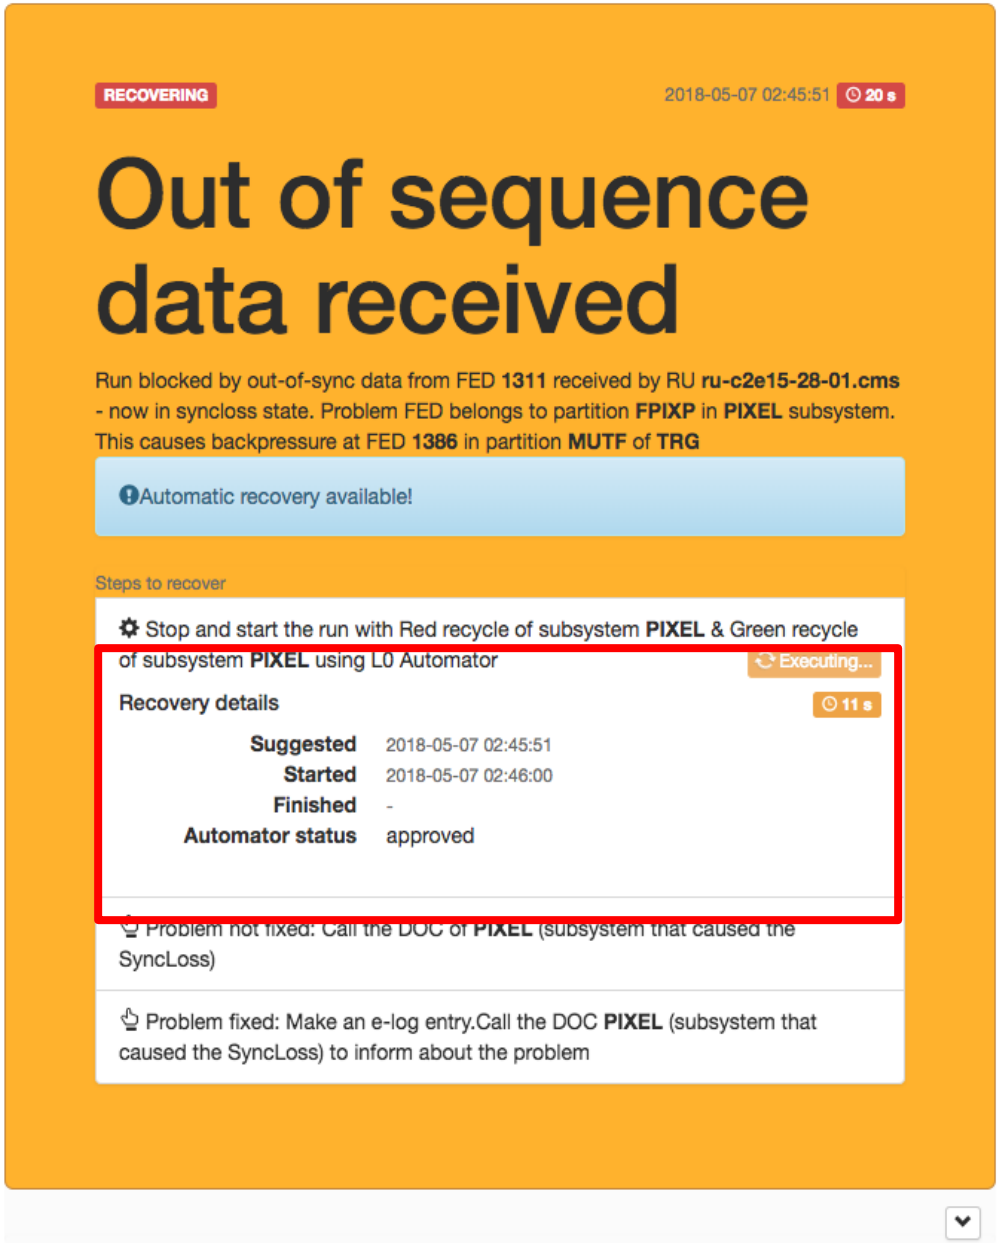

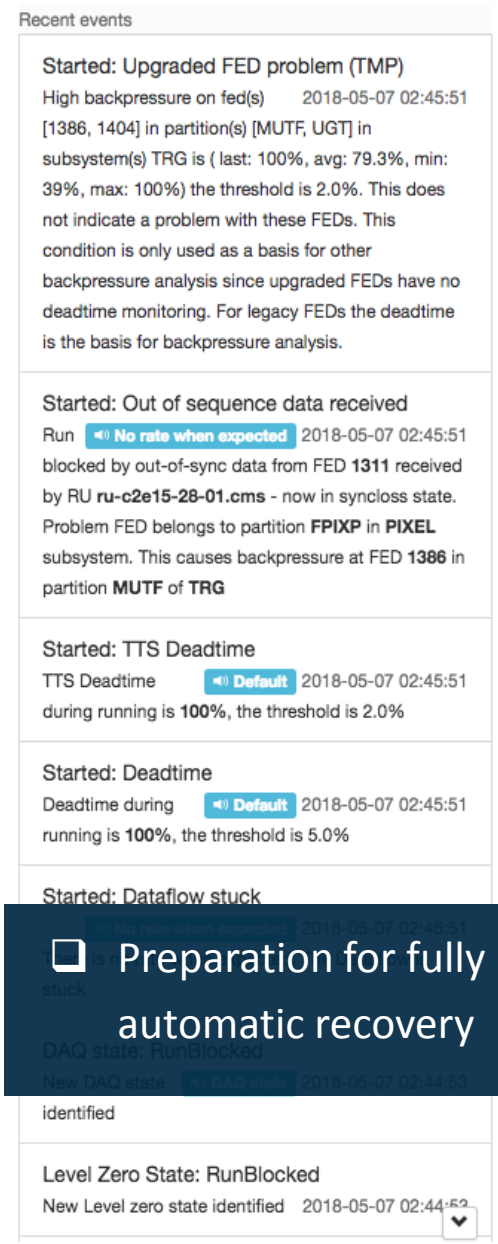

- Preparation for fully automatic recovery

Recovery driven by DAQ Expert

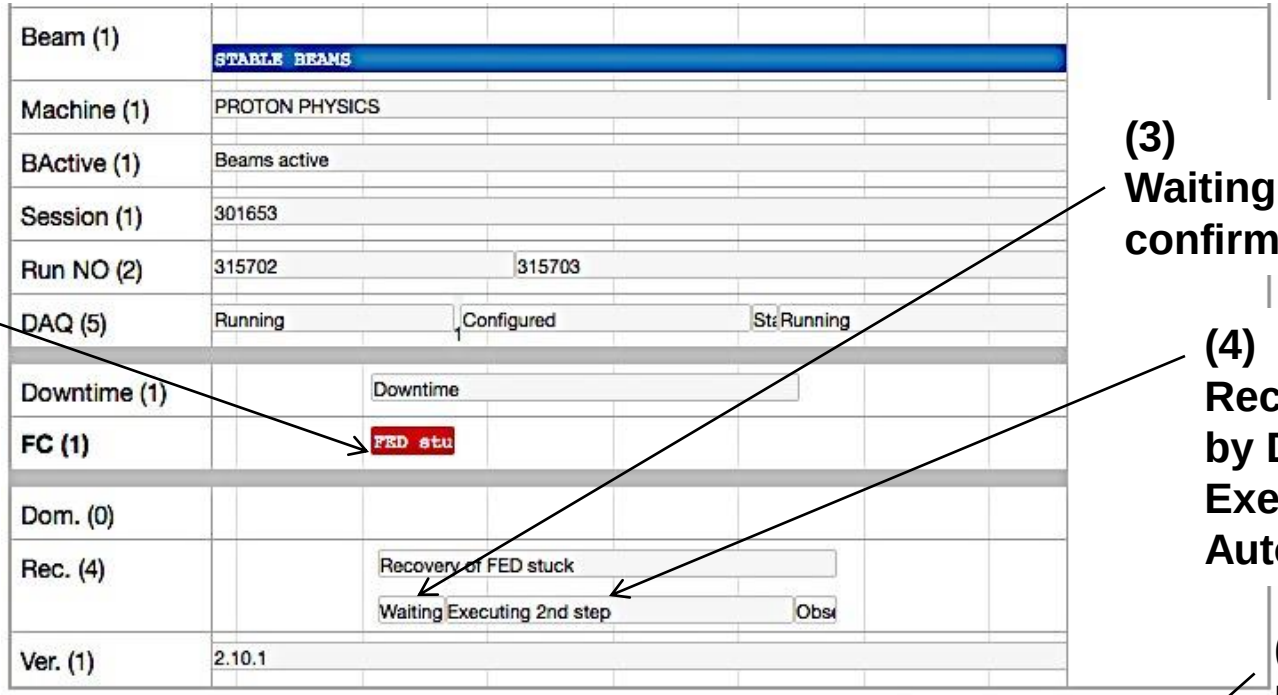

(1) Triggers stop

(2) Problem found

(3) Waiting for confirmation

(4) Recovery driven by DAQ Expert, Executed by L0 Automator FM

(5) Running again

Results

Results

- Difficult to provide hard metrics
  - ever-changing system
    - Continuous improvements
    - Upgrades
      - 2015: DAQ-2
      - 2016: L1 Trigger
      - 2017: Pixel
  - ever-changing shift crew
  - rare infrastructure problems
- Feed-back from shifters very positive
  - Finding problems has become less stressful

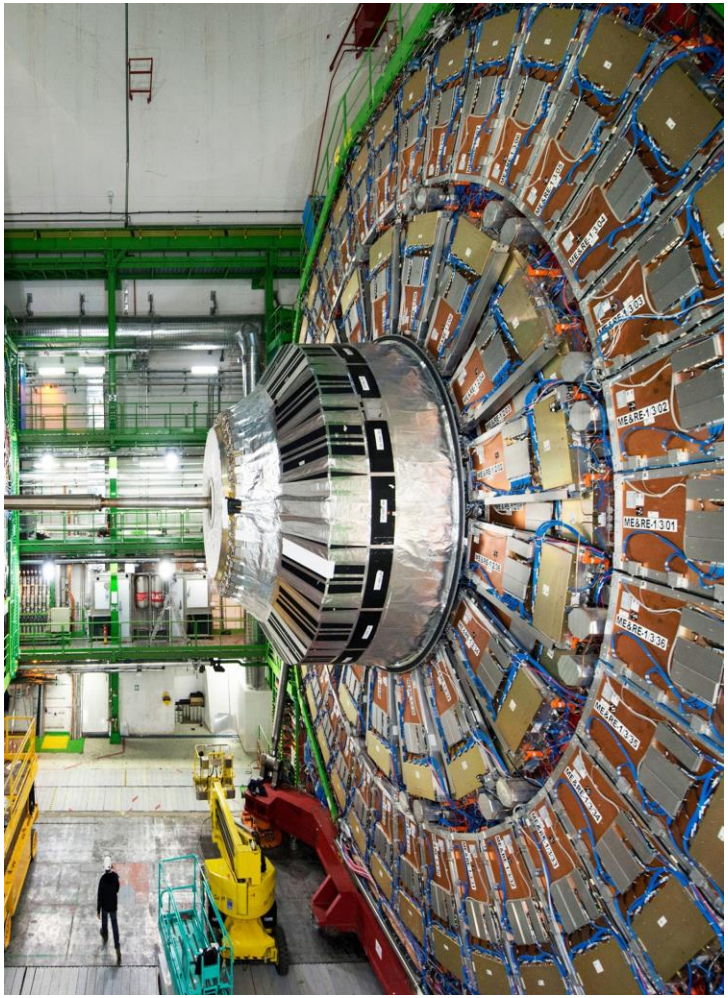

Results

- Observed higher correctness of recovery action
- Coverage: 95%
  - for problems stopping data taking during stable beams in 2018
- DAQ-Expert contributed to
  - Increasing data taking efficiency
  - Reducing number of calls to the on-call expert

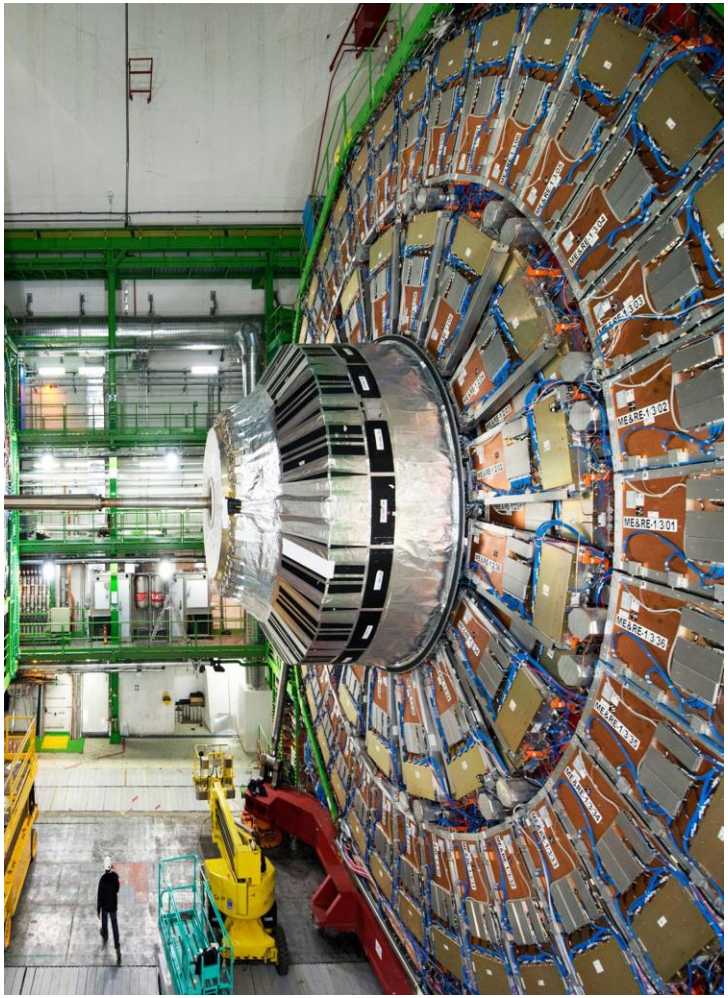

Results: Less luminosity lost

	Luminosity lost due to	
	Online systems: L1 Trigger + subsystem DAQ + Central DAQ	Central DAQ
2015 no Expert system	3.7%	0.20%
2016 DAQ Expert beta (from August)	2.6%	0.08%
2017 DAQ Expert	2.4%	0.02%
2018 with DAQ Expert (up to July 6)	1.8%	0.03%

- Proton fills with CMS magnet at 3.8 T, only
- Not showing luminosity lost due to infrastructure problems, tests, commissioning + dead time

Results: Number of night-time calls to the on-call

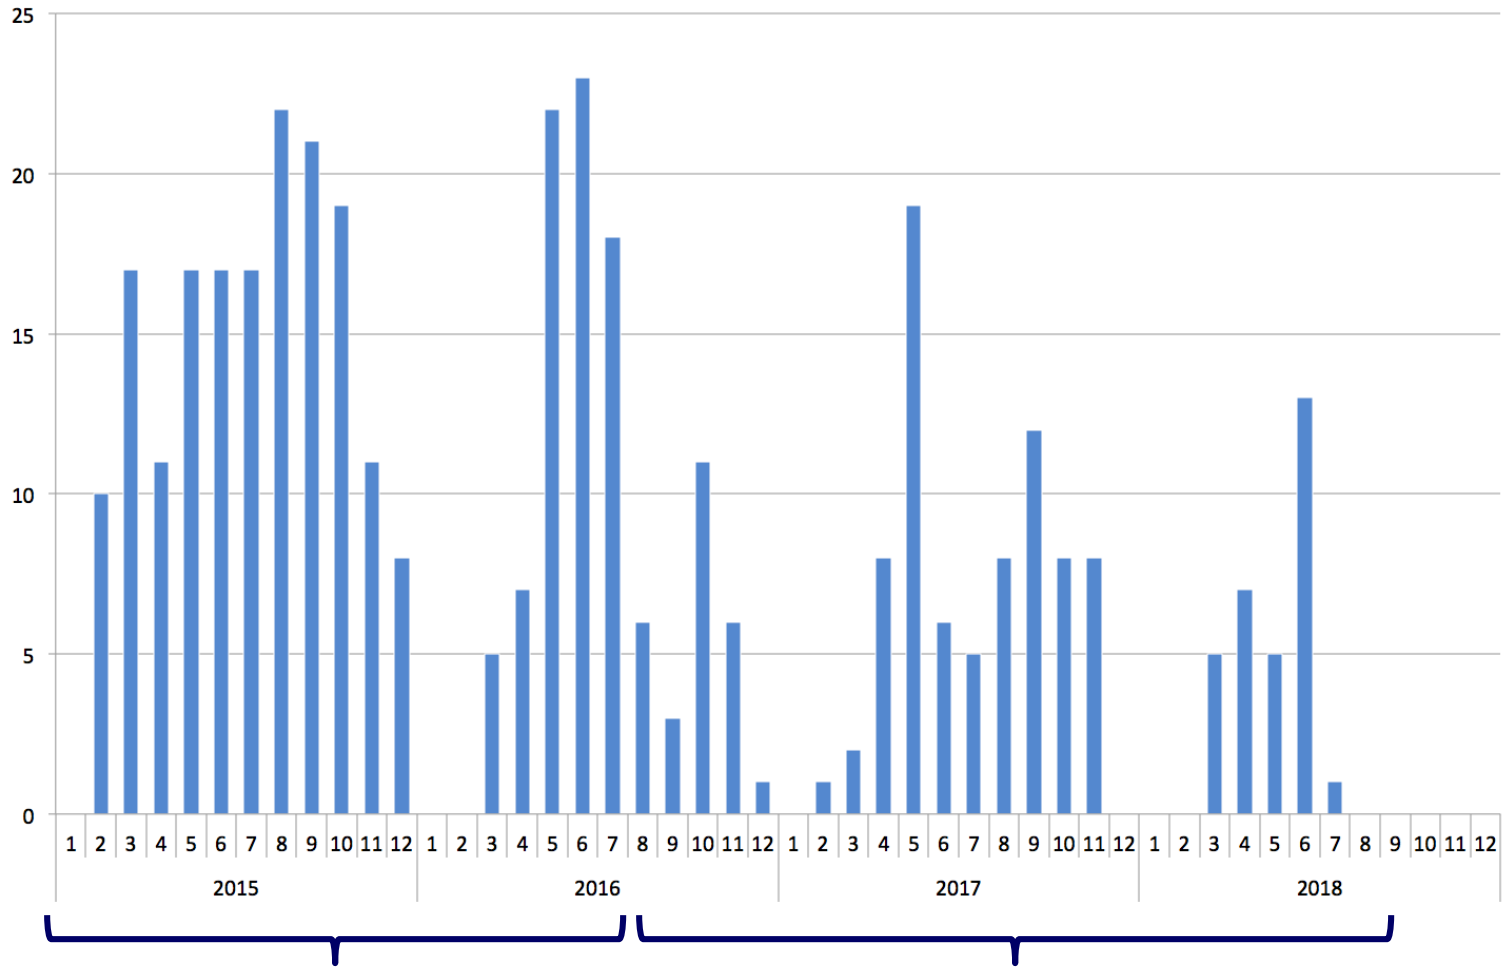

Based on data provided by IT / CS

Summary & Plans

1. **New expert tool for CMS operations**
   - Pure imperative language (Java) to implement reasoning
2. **Successful at CMS**
   - 95% coverage
   - Contribution to
     - increasing CMS efficiency
     - reducing need for external help
3. **Plan: automate recoveries completely**
   - for certain types of problems

Backup

Evolution of expert tools in DAQ

DAQDoctor

- DAQ- 1
- Run 1
- Perl
- Single person

2009

Esper-based tool research project

- Long Shutdown 1
- Declarative approach
  - Group not used to this paradigm

2013

DAQExpert

- DAQ-2
- Based purely on Java
  - Language widely known in the group
- modular

2016-2018

Results: Higher availability

	Availability during stable beams	
	Online systems: L1 Trigger + subsystem DAQ + Central DAQ	Central DAQ
2015 no Expert system	96.3%	99.80%
2016 DAQ Expert beta (from August)	97.4%	99.92%
2017 DAQ Expert	97.6%	99.98%
2018 with DAQ Expert (up to July 6)	98.2%	99.97%

- Percentage of delivered integrated luminosity for which systems were available, proton fills with CMS magnet at 3.8 T, only
- Data from CMS Web Based Monitoring