sPHENIX Annual MIE Review
1.6 DAQ/Trigger

Martin Purschke

July 14-15, 2021

BNL

DAQ/Trigger (WBS 1.6)

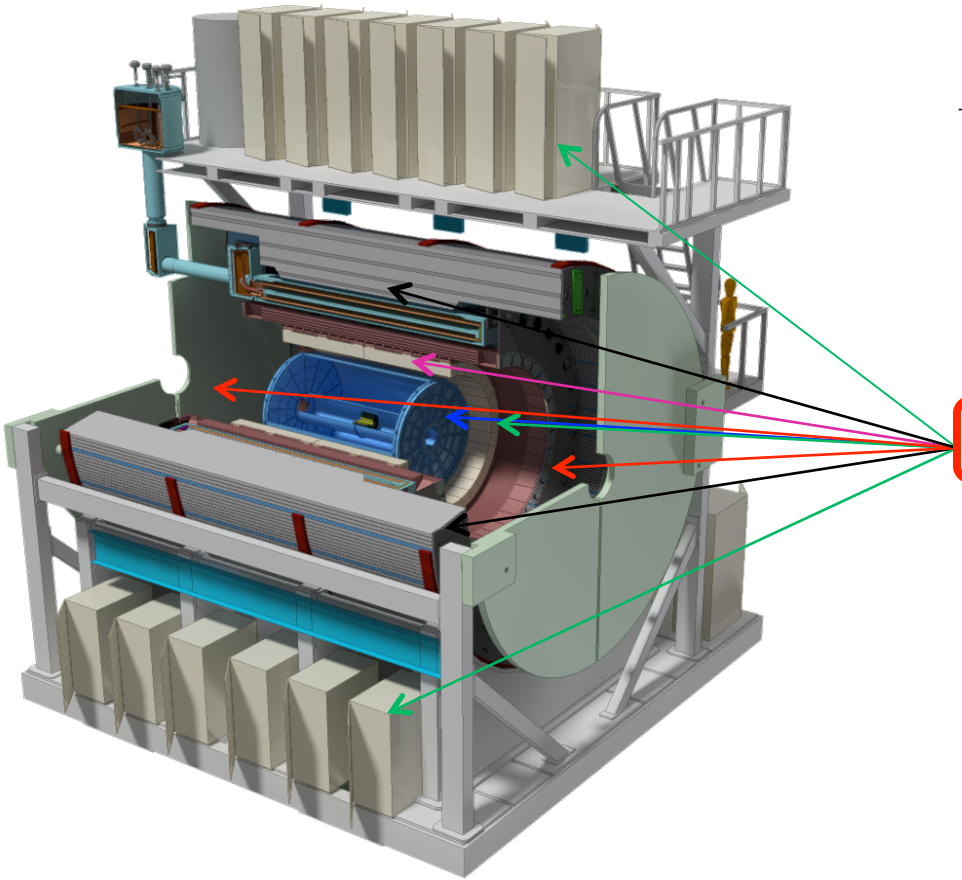

WBS	sPHENIX MIE Project Elements
1.1	Project Management
1.2	Time Projection Chamber
1.3	Electromagnetic Calorimeter
1.4	Hadron Calorimeter
1.5	Calorimeter Electronics
1.6	DAQ-Trigger
1.7	Minimum Bias Trigger Detector

The Subsystem Technical Overview

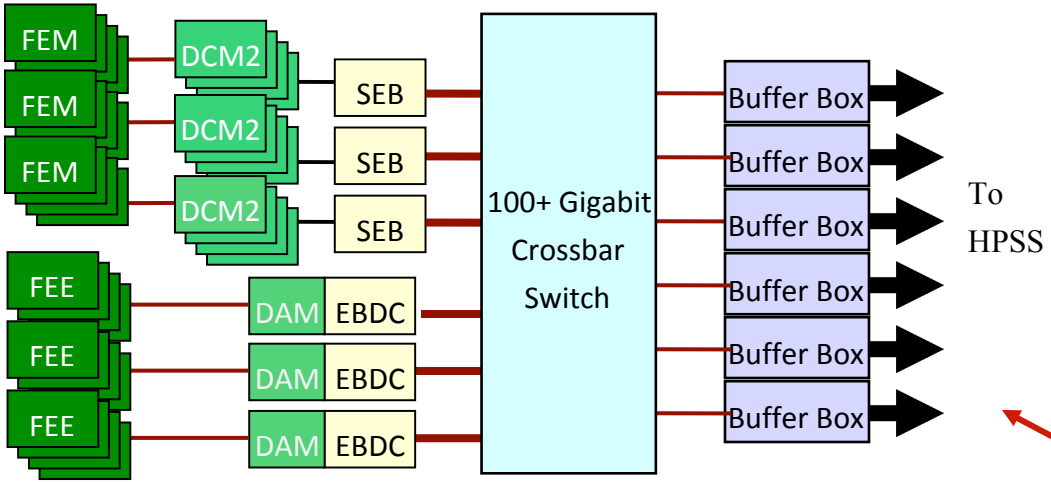

4 Components:

- Core DAQ system (1.6.1)
- Local Level 1 systems (1.6.2)
- Global Level 1 system (1.6.3)
- Timing system (1.6.4)

The 1.6.1 overview here shows how things hang together, not all of it is "ours"…

Local Level 1

Global Level 1

Slow control system

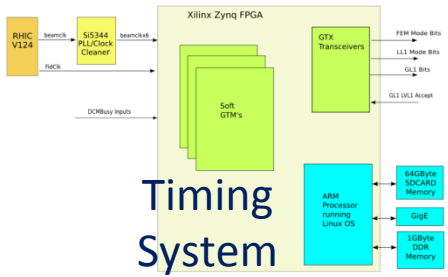

Glossary of acronyms in the backup slides

The Subsystem Technical Overview

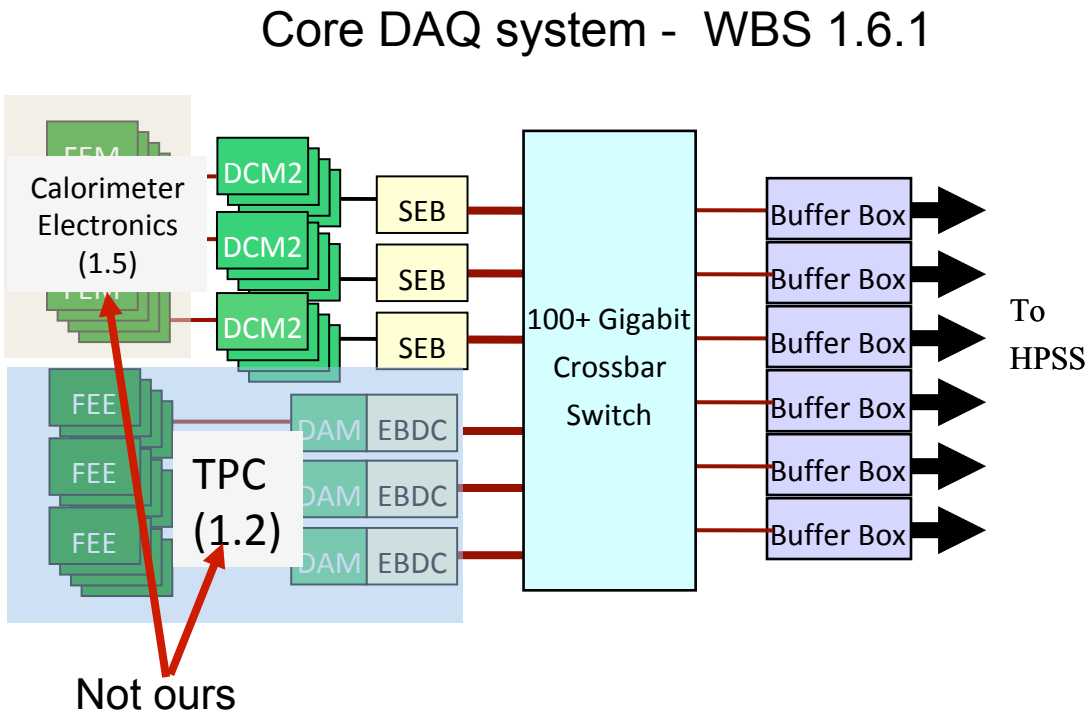

The full 1.6.1 system on the last slide was shown for context

Here is what the 1.6.1 Core DAQ system is responsible for:

- DCM2s (in hand, re-used)
- Subevent Buffers (SEBs), (commercial PCs)
- Network switch
- Buffer Boxes (commercial file servers/PCs)

Scope and Deliverables

Take Data with a ~90% live time at 15KHz event rate.

Buffer and send Data for permanent storage to the SDCC-hosted HPSS system

Provide the Local Level 1 hardware and firmware for triggering

Provide the Global Level-1 hardware and firmware for the experiment

Provide the Timing System to distribute the clocks and trigger info to the front-end electronics

Configuration of all components / overall configuration management

Online/near-line monitoring technology of data / detectors

Subsystem Collaborators

- 1.6 Data Acquisition and Trigger (M. Purschke, BNL, L2)
  - **1.6.1 The core DAQ system (M. Purschke)**
  - **1.6.2 Local-Level 1 Triggers (Jamie Nagle, U Colorado, L3)**
  - **1.6.3 Global Level 1 Trigger**
  - **1.6.4 Timing system**

} **(M. Purschke)**

Accounting-wise, these remain separate, but have been merged into one for all intents and purposes (same hardware platform)

Additional contributors:

John Haggerty, Eric Mannel, Cheng-Yi Chi, Dennis Peripelitsa, Jin Huang, Chris Pinkenburg, John Kuczewski, Jan Bernauer, Joe Mead, Danny Padrazzo, Wei Wang, Steve Boose

Schedule Performance and Schedule To Go

DAQ		50	399		30-Mar-20 A	28-Oct-21	20-Apr-20	05-May-21
DAQ Production		50	399		30-Mar-20 A	28-Oct-21	20-Apr-20	05-May-21
S260100	Procure DAQ Production Main Switches - Contract Award(s)		0	100%	30-Mar-20 A		16-Jun-20	
S260300	Procure DAQ Production Main Switches - Delivery Acceptance		3	100%	21-Apr-20 A	24-Apr-20 A	09-Dec-20	15-Dec-20
S257400	Procure DAQ Production SEBs - Contract Award(s)		0	100%	30-Sep-20 A		16-Jun-20	
S258200	Procure DAQ Production ATPs - Contract Award(s)		0	100%	30-Sep-20 A		16-Jun-20	
S258400	Procure DAQ Production ATPs - Delivery Acceptance		25	100%	22-Oct-20 A	30-Nov-20 A	12-Apr-21	16-Apr-21
S259300	Procure DAQ Production Buffer Boxes - Contract Award(s)		0	100%	28-Oct-20 A		16-Jun-20	
S257600	Procure DAQ Production SEBs - Delivery Acceptance		20	100%	28-Oct-20 A	30-Nov-20 A	08-Mar-21	12-Mar-21
S259500	Procure DAQ Production Buffer Boxes - Delivery Acceptance	154	130	16%	23-Nov-20 A	01-Jun-21	29-Apr-21	05-May-21
S255700	Procure DAQ Production Boards - Contract Award(s)	5	0	0%	01-Jun-21		18-May-20	
S256500	Procure DAQ Production Crates - Contract Award(s)	44	0	0%	01-Jun-21		20-Apr-20	
S258700	Procure DAQ Production jSEB Slow control computers - Delivery Acceptance	120	5	0%	01-Jul-21	08-Jul-21	23-Jun-20	29-Jun-20
S255900	Procure DAQ Production Boards - Delivery Acceptance	5	5	0%	23-Sep-21	29-Sep-21	08-Mar-21	12-Mar-21
S256700	Procure DAQ Production Crates - Delivery Acceptance	44	5	0%	22-Oct-21	28-Oct-21	08-Mar-21	12-Mar-21

Ready for Operation:

- Timing System 1/2022
- DAQ 1/2022
- Global Level 1 1/2022
- Local Level 1 12/2021

Cost Performance

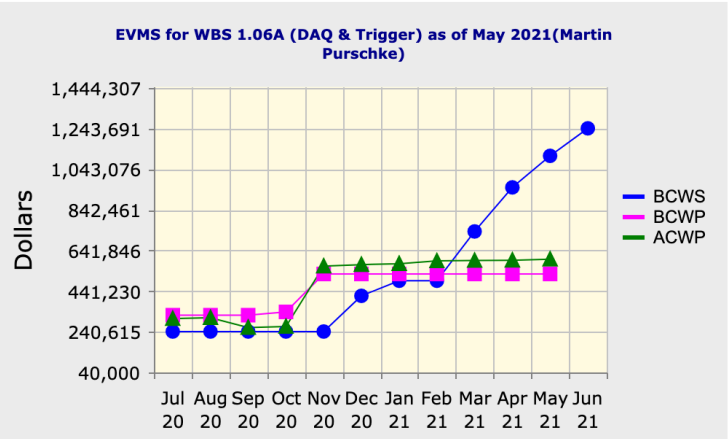

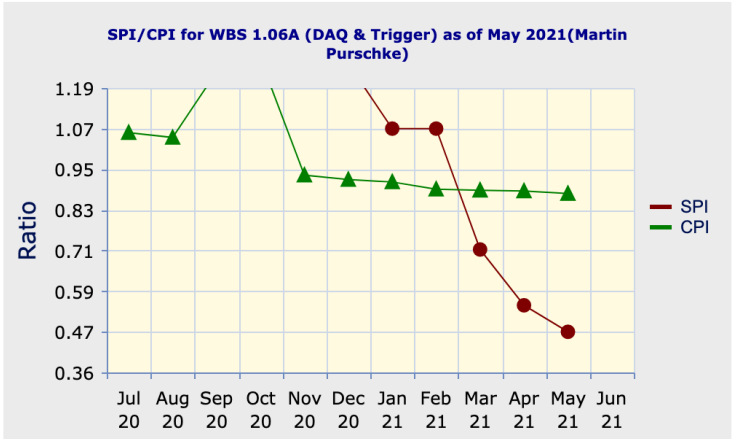

	BCWS	BCWP	ACWP	SV in $	SV in %	CV in $	CV %	SPI	CPI
Current:	155,676	0	5,129	-155,676	-100%	-5,129	0%	0.00	0.00
Cumulative:	1,109,290	525,778	599,107	-583,512	-53%	-73,329	-14%	0.47	0.88
	BAC								
At Complete:	1,245,090								

- "Bufferbox" order (5 addt'l storage systems and two racks) delayed by 15 weeks due to parts shortage
- Some mitigation: Vendor delivered the (empty) racks on schedule so the rack installation could proceed

Status of Reviews

- Local-Level 1 3/2021
  - **4 recommendations , next slide**

- Software and computing review 3/2021
  - **No recommendations for DAQ**

LL1 Review Recommendations

1. Additional tests of the present LL1 board should be planned and carried out before June, 2021.
2. Changes in the LL1 board should be finalized as soon as possible, so that the next round of prototype boards can be ready for testing in summer, 2021. The next round of testing should be aimed at a test that can verify the operation of the MBD LL1 by taking advantage of the existing laser system and the actual MBD detector which can be operated on the bench in 1008.
3. The full layout of the LL1 system should be made, showing fibers, crates, boards, interconnects, power, and interfaces to calorimeter electronics and GL1/GTM electronics.
4. The protocol and data sent to the GL1/GTM should be specified and documented before September, 2021.

Status and Highlights - 1

Data rate and volume corrected downwards

Numbers taken from the beam-use proposal

Significantly lower data volumes due to the introduction of a beam crossing angle, fewer useless off-vertex collisions (TPC)

Numbers include 30% uncertainty to the high side

Run 1: Au+Au: 13 weeks @ 60% RHIC uptime x 60% sPHENIX uptime ⟶ 43 billion events 72 PB 73Gbit/s

Run 2: p+p, p+A: 21 weeks @ 60% RHIC uptime x 60% sPHENIX uptime ⟶ 69 billion events 78 PB 49Gbit/s

Run 3: Au+Au: 24.5 weeks @ 60% RHIC uptime x 80% sPHENIX uptime ⟶ 107 billion events 180 PB 97Gbit/s

These are conservative uptime figures establishing *minimum* sampled luminosity goals.
We can still write up to 135Gbit/s if needed, so significantly higher uptimes are ok

Status and Highlights -2

Installed in Bldg 1008 right now

102 14TB disks
~ 1.1 PB usable

1 Bufferbox

1 Rack configured as SEBs

1 Rack configured as EBDCs

Multiple System Readout Milestone

- We are reading time-aligned data from the GL1 (that actually defines “time”) together with data taken with the calorimeter digitizer system
- The calorimeter digitizer system does not provide the original absolute beam clock value in the data stream
- We enlist 8 digitizer channels to digitize 8 bits of that counter, brought out on a GL1 header, connected to digitizer inputs

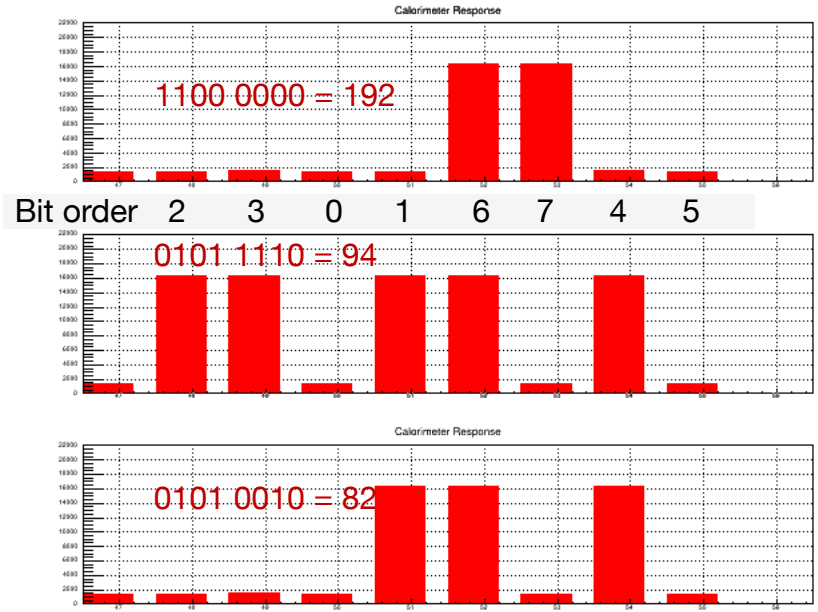

With 8 bits we get 0…255 to compare with the full counter.
You will see the 192 - 94 - 82 sequence on the next slide

Status and Highlights - 5

Multiple System Readout – Proof of event alignment

Beam Clock from GL1

evt nr	Bclk
2	192
3	94
4	82
5	242
6	66
7	94
8	181
9	41
10	71
11	12
12	57
13	119
14	88
15	84
16	191
17	158
18	237
19	235
20	175
. . .	

Beam Clock from data

evt nr	Bclk
2	192
3	94
4	82
5	242
6	66
7	94
8	181
9	41
10	71
11	12
12	57
13	119
14	88
15	84
16	191
17	158
18	237
19	235
20	175
. . .	

We are reading the 2 systems' data together and compare the numbers (within the 8 bit limit)

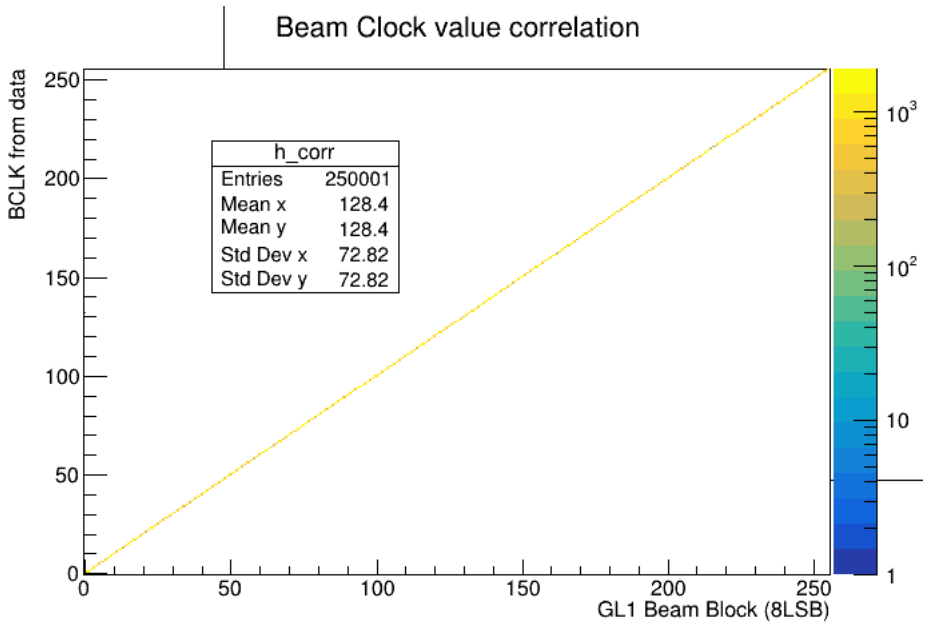

500 Million Event Challenge

Test of the data logging system and offline event building

Use 15 SEBs/EBDCs = 15 streams

Run 4 “volleys” one after another to get 60 files, 240 packets combined, about the number of packets and event size we expect

This tests the entire data logging chain with realistic rates (with 1/6 capacity installed)

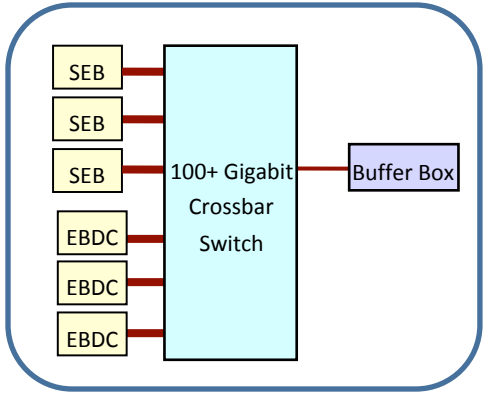

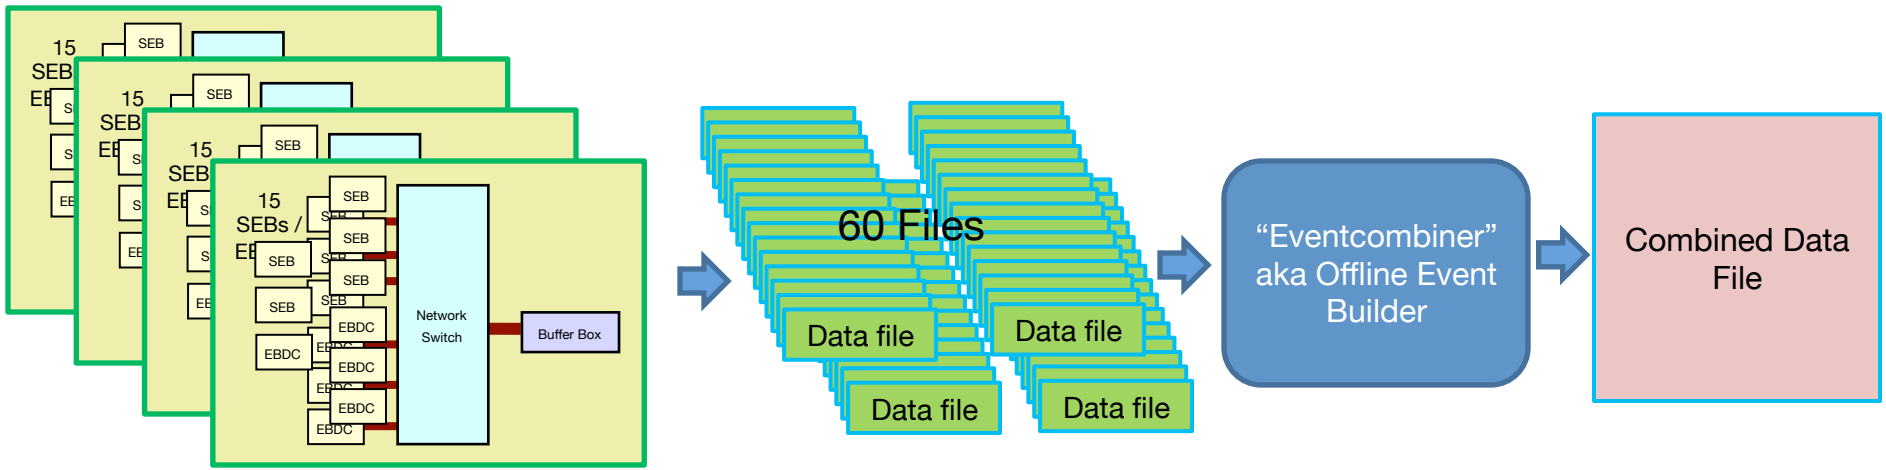

ES&H

Our system does not have gas, mechanical, etc hazards

The principal identified hazard is electrical (voltages > 50V)

Majority of components are commercial off the shelf and UL or better listed

Custom-developed components are subject to the relevant safety reviews and sign-offs by the ESRC

All such components are fused, have breakers, etc (all in racks with “buckets”)

Standard smoke / water detection protection

ES&H (and QA) Documentation and approval for one of the few custom components

BROOKHAVEN NATIONAL LABORATORY
BROOKHAVEN SCIENCE ASSOCIATES
UPTON, N.Y. 11973
Exploring Life's Mysteries, Protecting its Future

NSLS II NATIONAL SYNCHROTRON LIGHT SOURCE II

DRAWN BY	J. Mead	9/9/2016
CHECKED BY	A. Dellapenna	9/1/2016
VACUUM APPROVAL	NA	
ENGINEER APPROVAL	J. Mead	9/9/2016
SUPERVISOR APPROVAL	O. Singh	9/1/2016
ES&H APPROVAL	B. Lee	--/--/----
QA APPROVAL	E. Cheswick	--/--/----

SCHEMATIC DIAGRAM

RF ZBPM

ZDFE

B — SHEET SIZE

LT-EL-SR-EU-BPM-0601-03 — DRAWING/PART NUMBER

WBS# ----

ESH&Q RISK LEVEL A-2

SHEET 1 C

Issues and Concerns

- Hardware projects impacted by Covid (and red tape)
  - **LL1 SOW on hold at BNL for a long time**
  - **Routine procurements (e.g. board assembly at standard company) held up for weeks**
  - **Production GL1/Timing hardware under schedule pressure now**
  - **Still waiting for one more SOW to be finalized**

- Parts shortage with vendors
  - **Most COVID-related**
  - **Some due to other market developments (disk storage)**

Remaining Risks

Risk Identification			Risk Handling Plan (Mitigations)	Residual Risk (Po			
Risk ID Number	Risk Title	IF/THEN	Risk Handling Plan (Mitigations)	Residual Risk Likelihoc	Low Cost Impac	Likely Cost Impac	High Cost Impac
sPH_DAQ&Tr_001	DAQ Prototype does not meet specifications	If tests with the various prototype stages reveal problems, then DAQ prototype throughput and	Acquire more or more expensive PCs	5%	10	35	74
sPH_DAQ&Tr_005	TPC produces higher data rate than we can store	If the TPC or other subsystem cannot meet the envisioned data reduction specifications, then data	Invest in more local storage, change compression algorithms	30%	50	100	150
sPH_DAQ&Tr_007	DAQ DCM boards are late	If Nevis contract is delayed, then board production will be late	Work with BNL procurement, maintain Critical Procurement list for POB	20%	0	0	0

- The residual risk DAQ&Tr_001 is low
  - The overall throughput numbers were corrected downwards
  - Tests with 1/6 of the eventual system showed the proper throughput at the uncorrected levels
  - The residual risk is that a bottleneck will manifest itself with the full system
- Risk DAQ&Tr_005 can only manifest itself once the TPC is installed
  - the TPC is the main system with a downwards-corrected data rate that further mitigates that risk
- Risk DAQ&Tr_007 is being worked on
  - No technical issues addressed here, only administrative “red tape”

Summary

- We are on a good track with the DAQ
- “Slice test” (test of calorimeter, TPC, GL1) imminent with full DAQ support
- The DAQ (the real thing) in routine use for detector tests, test beams, …
- Float lost due to COVID
- Event Building, entire data handling system in place
- Solid commissioning plan with support for other subsystems (backup slides)

Back Up

Short Glossary

- FEE: Front End Electronics (TPC readout)
- FEM: Front End Module (calorimeter readout)
- DAM: Data Aggregator Module (for TPC)
- EBDC: Event-Buffering Data Collection computer (for TPC)
- SEB: Sub-Event Buffer
- DCM-2: Data Collection Module, version 2

Site Connectivity

We have plenty of unused fibers, more than we ever need for sPHENIX (72 fibers total, 36 pairs)

We plan to instrument 400Gbit/s (4 pairs) as uplink

(97 Gbs needed in Year 3 - 4x headroom)

Each bufferbox has 100Gbit/s connectivity

Hardware (transceiver optics) is in hand (good – COVID-related shortage)

The high-speed connection will be established in August/September 2021.

We are waiting for the final rack arrangement and then instrument the “last 50 yards”

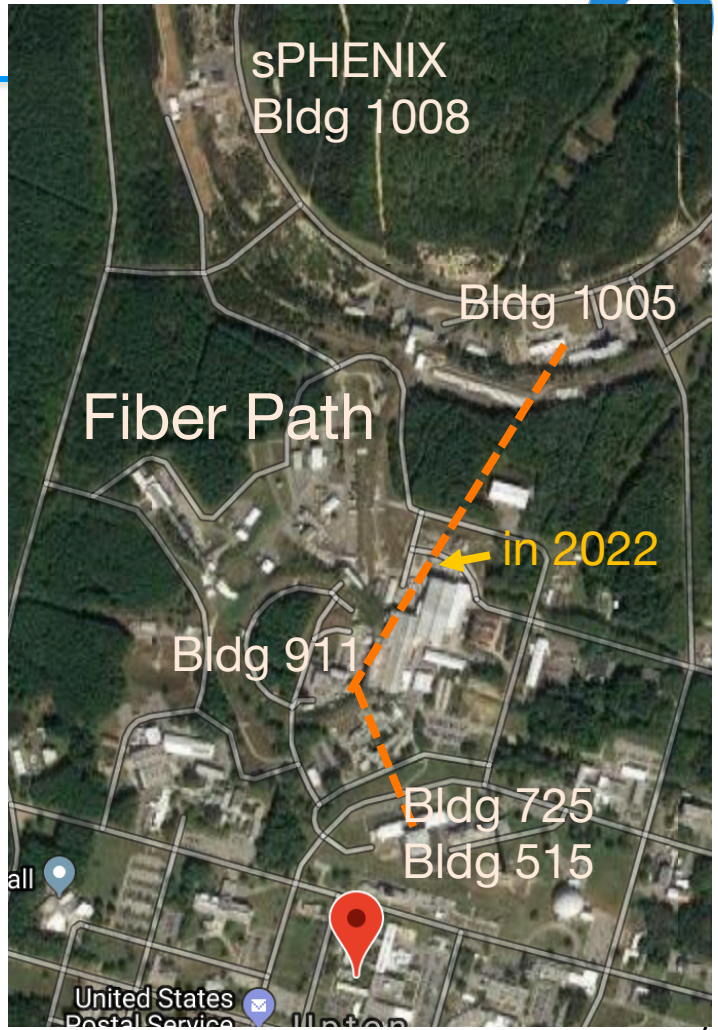

Commissioning Plan (1)

We have a large number of the zDFE ”small” GL1GTM units in hand. They can support a handful of detector systems, but usually support one

This is a picture/collage from the 2018/19 Fermilab test beams

INTT, MVTX, TPC, EmCAL, HCAL

Ok, that didn't actually happen then (today it would)

You see, we run several setups with the close-to-final trigger/timing hardware with ***completely independent*** DAQ systems

Every system can do its own thing – test, start, stop, debug, time in, what have you

Commissioning Plan (2)

With enough GL1/GTM units in hand, each subsystem can debug/test etc on their own

Can even independently run, say, TPC South and TPC North, ditto for EmCal or MBD

That achieves the PHENIX-style “partionining” in a way that PHENIX never did (you still had one trigger system, and partitions were NOT independent)

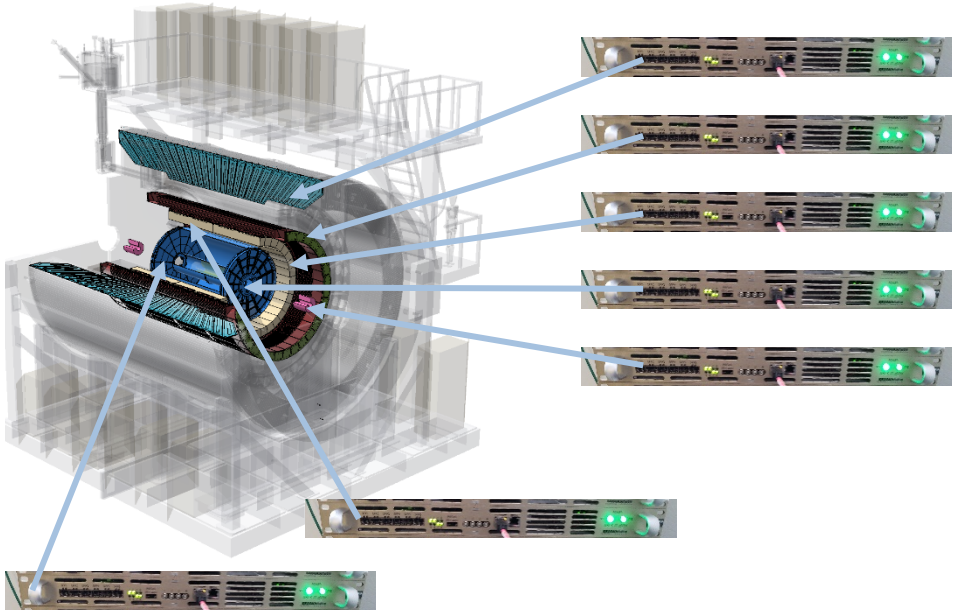

And that is really close to the final setup

Same packet Ids, etc

Same setup/config scripts

Etc

For some periods, we will try to run them all together, too

Commissioning Plan (3)

Then you connect your system's timing fibers to the "Big Iron" unit, and all runs together as one experiment

I expect that to greatly speed up the commissioning

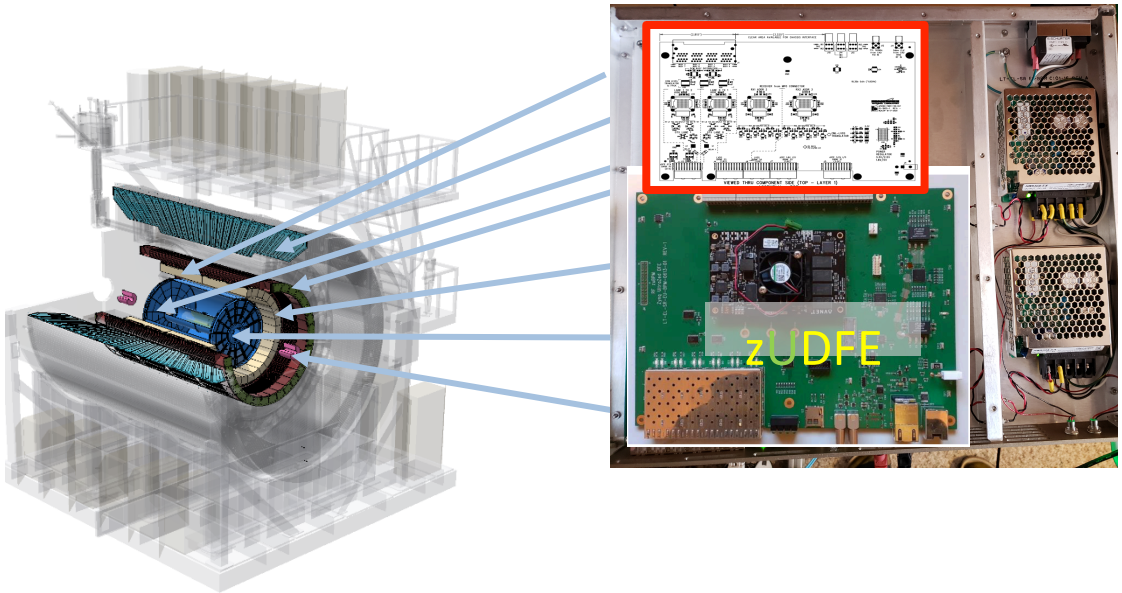

Status and Highlights – 5

500 Million Events Challenge – Offline event building check

We then read back the combined data file and the 60 originals into our analysis framework

Word by word we check that the packets in the combined file are the same as in the originals, and that they are all there and accounted for

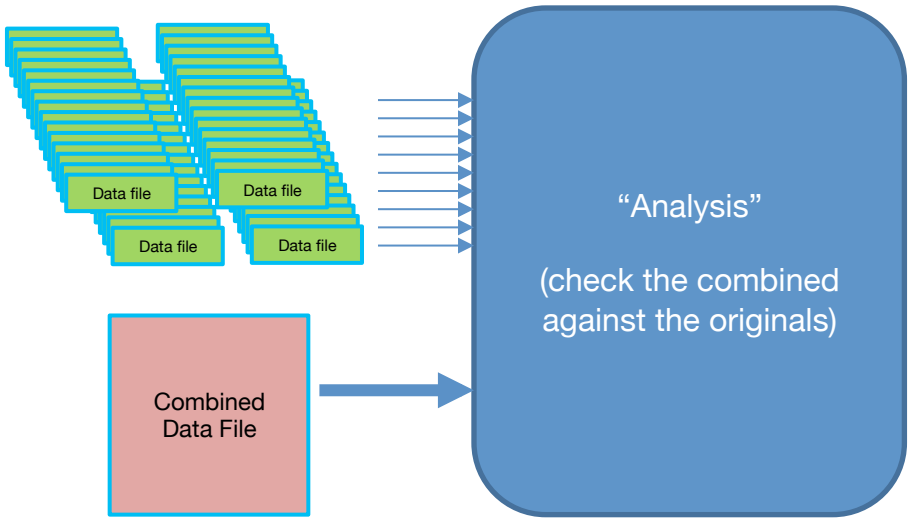

100% success rate!

It shows that our framework is able to handle the combined stream, as well as 60 on-the-fly streams

Keep in mind that those files are generated by the real DAQ processes