DUNE DAQ-CF interface

slides by Giles Barr

for Technical Board

Sept 30th 2015

Overview of Single-Phase DAQ

For CD1:

- ADCs in cold, - pipe data out: 80 1Gbit/s serial links/APA (=12000 links/10kt), ~1TB/s/10kt
- Zero suppress in FPGAs, do sensible trigger get all physics without moving most of the data.

Other ideas:

- Don't do zero suppression in cold, read all 4TB/s/40kt into computers and process there. – More flexibility, but needs 'lots of computers'.
- Difficulty is right now we want to pin down power/space requirements - not easy with such nascent designs.

Risks

Before this interface document signoff, appropriate risks that it may change must be established for both LBNF and DUNE.

The funding for DAQ has not been secured yet. It is likely that the ultimate provider of the DAQ in one or many of the caverns chooses to implement it differently from the assumptions here, and this has the risk of incurring additional costs in both projects.

Additionally, the current designs have large uncertainties on power consumption. Experts have refused to give me estimates, they say it is impossible to pin down. Some chip manufacturers deliberately do not specify a way to estimate power consumption of their devices, because it is hard with FPGAs to do this with any accuracy.

For the 'lots of computers' option, the calculation goes like this….

- One computer pulls about 300W can sink about 80Gbits/s = 10 GB/s (they say, not tried by me).
- We tested computers can just about read data from memory to CPU for processing at 1GB/s/core; on a multicore about 10GB/s

- This is DAQ, we must never lose the data (1TB/s/10kt), which never stops coming, so we need at least a factor two, better a factor three in capacity for overhead. With a factor three, 40kt = 1200 machines.
- The processor takes ~200W, memory ~100W, some DUNE DAQ designs, receive and sort data on FPGA cards with custom firmware. Hard to estimate power well but have added 3 x 100W. Then add 20% for PSU inefficiency and 20% kW->kVA and you get 940VA/machine.
- At 3kW/rack, this is an enormous number of racks, could push it to 10kW/rack, but that is O(100 racks), so a sizeable computer center.

Reminder: This is for the 'lots of computers' option. The FPGA DAQ design satisfies our requirements fine.

Cross check: Data-center design tool at amazon web services.

- Ask tool to quote for 1200 machines, each with 16 cores and 1000TB of storage (good for 25secs of raw data). Ask for max network speed tool allows, which is 10GB/s
- It estimates 88 racks and 671kW. So our estimate is OK.

Side remark: AmazonWS points at good articles on data center layout at www.energystar.gov, e.g. picture on right showing hot-cold-hot aisle layout. Worth a read (me too).

This will probably give correct method of going from #racks to floor area.

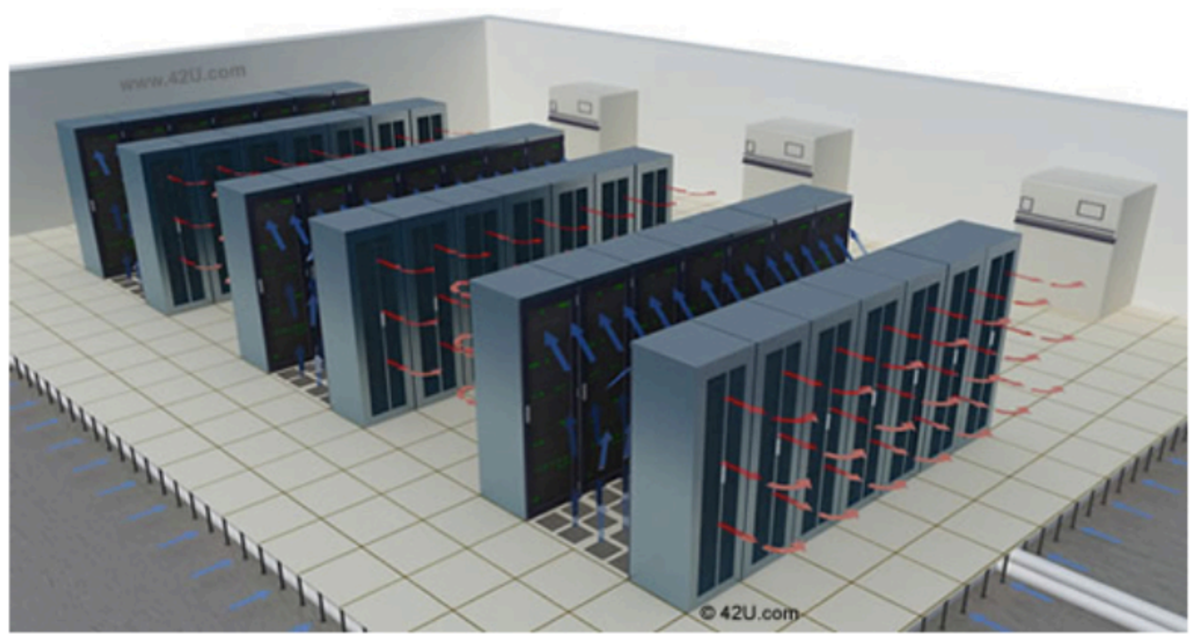

Figure 3: Hot aisle/cold aisle server row orientation. (Photo courtesy of 42U.com)

Component Power estimates

I have attempted to estimate power for each type of item in racks, including non-DAQ ones from Terri.

Dark blue = less certainty.

Crucial ones:

ATCA shelf– New values today from Matt, based on measurements at 35t and scaled up.

Felix receiver computers – see two slides back

Trigger nodes – similar calculation

Max W/rack CUA&surfac	3000
Power factor	0.8

Element	Power per 10kt	Multiplier (racks)	Power per rack	Space in U (per rack)	Multiplier (per rack)	element in U	Power in VA per element	Power estimate (W)
Ethernet switch	18750	75	250	2	2	1	125	100
SSP PD readout +calib	36750	75	490	7	7	1	70	56
DUNE timing	4688	75	63	4	1	4	63	50
MPOD LoV for preamps	70313	75	938	9	1	9	938	750
Rack protection	0	75	0	1	1	1	0	0
Totals for Det	130500	75	1740	23				
ATCA shelf	65625	75	875	5	1	5	875	700
Boardreader CPU SSP	18750	75	250	2	1	2	250	200
Boardreader CPU RCE	18750	75	250	2	1	2	250	200
Trigger node	9375	1	9375	20	10	2	938	750
Server node	4500	1	4500	16	8	2	563	450
Central 1G switch	500	1	500	5	5	1	100	80
Central 10G switch	500	1	500	1	1	1	500	400
Data storage node	1563	1	1563	10	5	2	313	250
Totals for RCE DAQ	119563		1125	52				
SubTotal move things	16438							
Felix receiver comp	281250	75	3750	16	4	4	938	750
Copper2Optic at flange	56250	75	750	8	4	2	188	150
Less RCE stuff	-84375	75	-1125	-7				
Delta for Felix	253125		3375	17				
Part-PCIe receiver comp	98438	75	1312.5	8	2	4	656	525
Copper2Ethnet at flange	56250	75	750	8	4	2	188	150
Local switch	7500	75	100	1	1	1	100	80
Less RCE stuff	-84375	75	-1125	-7				
Delta for Part-PCIe	77813		1038	10				
DualPhase uTCA	180000	20	9000	72	12	6	750	600
DualPhase FE	6000	20	300	0	12	0	25	20
DualPhase BE	25000	20	1250	0	1	0	1250	1000
DualPhase Online	125000	1	125000	0	1	0	125000	100000
Total for DualPhase	336000		10550	72				
Delta for Dual Phase	216438							
Copper2Ethnet at flange	56250	75	750	8	4	2	188	150

System power estimates

Considered 4 design philosophies: ('RCE' = reference design from CDR but with optical coupling removed, 'LBNO'=alternate deign, a.k.a. scaled up WA105, Felix = CERN suggestion to read all data into computers, part-PCIe =some mix of the other philosophies).

Considered combinations of 'on cryostat (flange), in central utility or surface.

Power in VA (not including A/C)			At Flange						At CUA			At Surface	
		Total power/10kt	Total power per rack at SP flange	Non-DAQ power/10kt at flange	Non-DAQ U/flange at flange	DAQ power/10kt at flange	DAQ U/flange at flange for SP	Data transfer/10kt cavern to CUA in GB/s	DAQ power/10kt at CUA	#racks at CUA	Data transfer /10kt CUA-surface in GB/s	DAQ power/10kt at surface	#racks/10kt at surface
A. Logic and LC at flange TC at CUA	1. RCE design	262,063	3,115	130,500	23	103,125	9	0.06	16,438	6	0.06	12,000	4
	2. LBNO design	478,500	DP	130,500	23	211,000	DP	1000	125,000	42	0.06	12,000	4
	3. Part-PCIe design	321,125	3,903	130,500	23	162,188	16	0.06	16,438	6	0.06	12,000	4
	4. Felix design	496,438	6,240	130,500	23	337,500	24	0.06	16,438	6	0.06	12,000	4
B. Logic and LC at flange TC at Surface	1. RCE design	262,063	3,115	130,500	23	103,125	9	0.06	12,000	4	0.06	16,438	6
	2. LBNO design	478,500	DP	130,500	23	211,000	DP	1000	12,000	4	1000	125,000	42
	3. Part-PCIe design	321,125	3,903	130,500	23	162,188	16	0.06	12,000	4	0.06	16,438	6
	4. Felix design	496,438	6,240	130,500	23	337,500	24	0.06	12,000	4	0.06	16,438	6
C. Logic at flange, LC and TC at CUA	1. RCE design	262,063	2,865	130,500	23	84,375	7	0.06	35,188	12	0.06	12,000	4
	3. Part-PCIe design	321,125	2,590	130,500	23	63,750	8	1125	114,875	39	0.06	12,000	4
	4. Felix design	496,438	2,490	130,500	23	56,250	8	1125	297,688	100	0.06	12,000	4
D. Logic at flange, LC at CUA and TC at Surface	1. RCE design	250,063	2,865	130,500	23	84,375	7	0.06	18,750	7	0.06	16,438	6
	3. Part-PCIe design	309,125	2,590	130,500	23	63,750	8	1125	98,438	33	0.06	16,438	6
	4. Felix design	484,438	2,490	130,500	23	56,250	8	1125	281,250	94	0.06	16,438	6
E. Logic at flange, LC and TC at Surface	1. RCE design	262,063	2,865	130,500	23	84,375	7	0.06	12,000	4	1125	35,188	12
	3. Part-PCIe design	321,125	2,590	130,500	23	63,750	8	1125	12,000	4	1125	114,875	39
	4. Felix design	496,438	2,490	130,500	23	56,250	8	1125	12,000	4	1125	297,688	100
F. Logic, LC and TC at CUA	1. RCE design	318,313	2,740	130,500	23	75,000	2	1125	100,813	34	0.06	12,000	4
G. Logic and LC at CUA, TC at Surface	1. RCE design	306,313	2,740	130,500	23	75,000	2	1125	84,375	29	0.06	16,438	6
H. Logic at CUA,, LC and TC at surface	1. RCE design	306,313	2,740	130,500	23	75,000	2	1125	65,625	22	1125	35,188	12

General conclusions

- The total power estimates are in the broad range from 320kVA to 490kVA per 10kt module.
- The ones that put more processing in computers rather than FPGAs generally offer more flexibility but use more power. All of them will satisfy the physics requirements.
- The network capacity to the surface is either (a) rather modest, if trigger selection is underground or (b) rather demanding (approx. 1000 fast links) for the options where we do all at surface.
- The power needed at each flange port is either, 2.9kVA (no processing) or 3.1kVA (RCE), 4.5kVA (LBNO-style) or 6.3kVA (Felix style computer processing at flange).
- These numbers are still very approximate and the appropriate level of contingency has not been considered yet (none included). This is one of several warnings needed – these are all back-of-envelope numbers.

The main choices can be captured in three questions to the technical board.

Technical board question #1 Network capacity in shafts

- There are only two possible ways this requirement can go, either we trigger underground or we trigger on surface.
- If we trigger underground, the data rate is about 60MB/s/10kt, and the proposed provision of 96 strands (16 per 10kt and 32 for CF) with backup in the other shaft is OK. Could we have slightly more, e.g. 72 fibers (18 fibers per 10kt) and 24 fibers for CF? [Added after discussion with colleagues: Or another 96 bundle?]
  - The 18 fibers would then be four fibers for GPS (1PPS up/down, 10MHz up/down), or (transmit L1/L2 r.f. down for two receivers), seven-pairs for 10Gb Ethernet (3 pairs data, 1 pair low-latency trigger, 1 pair slow control and 2 pairs spare).
- If we trigger on surface, the data rate vastly exceeds the capacity of 96 strands using todays sure off-the-shelf solutions (we need about 1000 strands per 10kt at 10Gbit/s), but there may be big improvements on the 5-year horizon. **It would be too risky to propose that in a CD3A review now**. Scope for a group to investigate – so allow room in the interface document if possible for this change.

Propose interface agreement something like:

The CF shall provide 72 single-mode fibers for the DAQ between underground central area and the surface, capable of running 10Gbit/s Ethernet and GPS signals (i.e. similar to Ethernet ones but no Ethernet-specific repeaters) and will have 24 more strands for their own use. There shall be at least a duplicate set of 96 fibers in the other shaft for backup. The exact fiber type (wavelength, single mode etc.) selection shall be made shortly before procurement by a joint LBNF-DUNE discussion to allow selection of the most future-proof option apparent at that time, and could incorporate findings of DUNE collaborators for running much faster Ethernet to get the un-triggered data out, if progress has been made in the intervening time.

Technical board question #2
Top of cryostat power/cooling

- Summary is per flange/port (=2 APAs). 2.8kW min, 3.2kW with RCEs on flange, 4.5kVA for LBNO style, 6.0kW with Felix type computer on flange.
- In all cases, power exceeds capacity of a rack with local cooler, so strongly recommend we look for another solution. CF group also think a centralized system where the heat can be deposited directly in the airflow up to the surface is better. Can be distributed with air ducting or water. Lots of experience at CERN with building water cooling sturdily into racks. The racks can still be enclosed if that reduces humidity/dust.
- Suggestion is to specify 4.5kW power provision here. That may be modified with question #3. If we still need the individual cooler this becomes 9kW/rack

Propose interface agreement:

Ask CF group to provide wording to describe the central cooling and the interface to it. It would be nice if we can be vague about water/air choice at CD3A so we can get future unrushed work on designing it. It would be even better if we can add to schedule a joint LBNF/DUNE VE exercise on this aspect of the design..

There shall be one 42U 19-inch NEMA-12 rack, with power capacity for 4.5kW at each of 75 port locations in the first detector hall. There will be doors both at the front and back of the rack and access will be needed to both, including min 3.5ft in front of the rack for equipment loading. The subsequent detector halls may have different TPC readout structures, so different arrangements of power, but the total will be the same.

Technical board question #3
Where to locate all the computers?

Both the Felix design and the LBNO design currently foresee reading all the data into computers. This is about 1TB/s/10kt. It uses about 125kVA/10kt in the LBNO design (some processing takes place in FPGAs) and 300kW/10kt in the Felix design.

Technical board to choose which options to be considered further by LBNF and DAQ-group (numbers are for 40kt, all identical): []=scheme from slide 7 ‘System power estimates’

- In underground area LBNO style (4.5kVA at each port + 500kVA 160 rack = 4500ft^2 underground + 50kVA 16 rack = 450ft^2 facility on surface) [A2]
- In underground area Felix style (4.5kVA at each port + 1200kVA 400 rack = 11200ft^2 underground + 50kVA 16 rack = 450ft^2 facility on surface) [C4]
- On surface (see question #1, best not to say now at review we rely on getting data to surface – but can we avoid precluding it now? (e.g. [B2] or even [E4])
- Distributed on top of cryostat [A4] vs [A1]/[A2] changes power per port from 6.2kVA to 4.5 kVA, is this much?
- Don’t allow it, or technical board can place a constraint on power budget – will not preclude much physics, but it is a political question. [A4]/[C4] vs [A1]

If we can avoid giving division of power between central area and cryostat, choice now for TB boils down to

- Do we build a 400 rack or a 160 rack or a 32 rack facility underground (compare [C4] vs [A2] vs [A1])?
- Or do we leave enough capacity for Felix by including the extra 1200kVA/40kt to distribute over the cryostats, but don’t dig the extra space centrally? [A4] vs [C4]

Propose interface agreement: Pick from [A1], [A2], [A4], [C4] depending on technical board decision.

We should include a risk that a collaborator that can provide funding for hardware needs a changed requirement from LBNF – this is to cover that this aspect of these DAQ designs are still very preliminary.

From management

(1) We think that we should allow for 450 kW of total power for each 10-kton far detector module.

(2) We think we should limit the number of racks in the central utility cavern to no more than 50 with a total power budget of 450 KW

(3) No specific limits for the surface at this time

Giles's Google Docs link: (Also in Email)
https://docs.google.com/spreadsheets/d/12PjSeKaawC0JCCJbJjtMw5yifUvONiDeK66eWYUdTTI/edit?usp=sharing

The next slide is a copy of the page on this google-docs that we discussed in the meeting.

		Power in cavern	At CUA	
			DAQ power/10k t at CUA	#racks at CUA
C. Logic at flange, LC and TC at CUA	4. Felix design	186,750	297,688	100
C. Logic at flange, LC and TC at CUA	1. RCE design	214,875	35,188	12
A. Logic and LC at flange TC at CUA	1. RCE design	233,625	16,438	6
A. Logic and LC at flange TC at CUA	3. Part-PCIe design	292,688	16,438	6
A. Logic and LC at flange TC at CUA	2. LBNO design	252,438	125,000	42
F. Logic, LC and TC at CUA	1. RCE design	336,000	100,813	34
F. Logic, LC and TC at CUA	1b. RCE design local cooling	411,000	100,813	34
C. Logic at flange, LC and TC at CUA	3. Part-PCIe design	194,250	114,875	39
C. Logic at flange, LC and TC at CUA	1b. RCE design local cooling	429,750	35,188	12
A. Logic and LC at flange TC at CUA	1b. RCE design local cooling	467,250	16,438	6
A. Logic and LC at flange TC at CUA	4. Felix design	468,000	16,438	6

Backup slides

Component Power estimates aggressive targets

I have attempted to estimate power for each type of item in racks, including non-DAQ ones from Terri.

Dark blue = less certainty.

Crucial ones:

ATCA shelf– New values today from Matt, based on measurements at 35t and scaled up.

Felix receiver computers – A much more aggressive estimate, but may be possible.

Trigger nodes – same calculation as before

Max W/rack CUA&surfac 3000

Power factor 0.8

Element	Power per 10kt	Multiplier (racks)	Power per rack	Space in U (per rack)	Multiplier (per rack)	element in U	Power in VA per element	Power estimate (W)
Ethernet switch	18750	75	250	2	2	1	125	100
SSP PD readout +calib	36750	75	490	7	7	1	70	56
DUNE timing	4688	75	63	4	1	4	63	50
MPOD LoV for preamps	70313	75	938	9	1	9	938	750
Rack protection	0	75	0	1	1	1	0	0
Totals for Det	130500	75	1740	23				
ATCA shelf	65625	75	875	5	1	5	875	700
Boardreader CPU SSP	18750	75	250	2	1	2	250	200
Boardreader CPU RCE	18750	75	250	2	1	2	250	200
Trigger node	9375	1	9375	20	10	2	938	750
Server node	4500	1	4500	16	8	2	563	450
Central 1G switch	500	1	500	5	5	1	100	80
Central 10G switch	500	1	500	1	1	1	500	400
Data storage node	1563	1	1563	10	5	2	313	250
Totals for RCE DAQ	119563		1125	52				
SubTotal move things	16438							
Felix receiver comp	70313	75	937.5	8	2	4	469	375
Copper2Optic at flange	56250	75	750	8	4	2	188	150
Less RCE stuff	-84375	75	-1125	-7				
Delta for Felix	42188		563	9				
Part-PCIe receiver comp	98438	75	1312.5	8	2	4	656	525
Copper2Ethnet at flange	56250	75	750	8	4	2	188	150
Local switch	7500	75	100	1	1	1	100	80
Less RCE stuff	-84375	75	-1125	-7				
Delta for Part-PCIe	77813		1038	10				
DualPhase uTCA	180000	20	9000	72	12	6	750	600
DualPhase FE	6000	20	300	0	12	0	25	20
DualPhase BE	25000	20	1250	0	1	0	1250	1000
DualPhase Online	125000	1	125000	0	1	0	125000	100000
Total for DualPhase	336000		10550	72				
Delta for Dual Phase	216438							
Copper2Ethnet at flange	56250	75	750	8	4	2	188	150

System power estimates - aggressive targets

Considered 4 design philosophies: ('RCE' = reference design from CDR but with optical coupling removed, 'LBNO'=alternate deign, a.k.a. scaled up WA105, Felix = CERN suggestion to read all data into computers, part-PCIe =some mix of the other philosophies).

Considered combinations of 'on cryostat (flange), in central utility or surface.

Power in VA (not including A/C)			At Flange						At CUA			At Surface	
		Total power/10kt	Total power per rack at SP flange	Non-DAQ power/10kt at flange	Non-DAQ U/flange at flange	DAQ power/10kt at flange	DAQ U/flange at flange for SP	Data transfer/10kt cavern to CUA in GB/s	DAQ power/10kt at CUA	#racks at CUA	Data transfer /10kt CUA-surface in GB/s	DAQ power/10kt at surface	#racks/10kt at surface
A. Logic and LC at flange TC at CUA	1. RCE design	262,063	3,115	130,500	23	103,125	9	0.06	16,438	6	0.06	12,000	4
	2. LBNO design	478,500	DP	130,500	23	211,000	DP	1000	125,000	42	0.06	12,000	4
	3. Part-PCIe design	321,125	3,903	130,500	23	162,188	16	0.06	16,438	6	0.06	12,000	4
	4. Felix design	285,500	3,428	130,500	23	126,563	24	0.06	16,438	6	0.06	12,000	4
B. Logic and LC at flange TC at Surface	1. RCE design	262,063	3,115	130,500	23	103,125	9	0.06	12,000	4	0.06	16,438	6
	2. LBNO design	478,500	DP	130,500	23	211,000	DP	1000	12,000	4	1000	125,000	42
	3. Part-PCIe design	321,125	3,903	130,500	23	162,188	16	0.06	12,000	4	0.06	16,438	6
	4. Felix design	285,500	3,428	130,500	23	126,563	24	0.06	12,000	4	0.06	16,438	6
C. Logic at flange, LC and TC at CUA	1. RCE design	262,063	2,865	130,500	23	84,375	7	0.06	35,188	12	0.06	12,000	4
	3. Part-PCIe design	321,125	2,590	130,500	23	63,750	8	1125	114,875	39	0.06	12,000	4
	4. Felix design	285,500	2,490	130,500	23	56,250	8	1125	86,750	29	0.06	12,000	4
D. Logic at flange, LC at CUA and TC at Surface	1. RCE design	250,063	2,865	130,500	23	84,375	7	0.06	18,750	7	0.06	16,438	6
	3. Part-PCIe design	309,125	2,590	130,500	23	63,750	8	1125	98,438	33	0.06	16,438	6
	4. Felix design	273,500	2,490	130,500	23	56,250	8	1125	70,313	24	0.06	16,438	6
E. Logic at flange, LC and TC at Surface	1. RCE design	262,063	2,865	130,500	23	84,375	7	0.06	12,000	4	1125	35,188	12
	3. Part-PCIe design	321,125	2,590	130,500	23	63,750	8	1125	12,000	4	1125	114,875	39
	4. Felix design	285,500	2,490	130,500	23	56,250	8	1125	12,000	4	1125	86,750	29
F. Logic, LC and TC at CUA	1. RCE design	318,313	2,740	130,500	23	75,000	2	1125	100,813	34	0.06	12,000	4
G. Logic and LC at CUA, TC at Surface	1. RCE design	306,313	2,740	130,500	23	75,000	2	1125	84,375	29	0.06	16,438	6
H. Logic at CUA,, LC and TC at surface	1. RCE design	306,313	2,740	130,500	23	75,000	2	1125	65,625	22	1125	35,188	12

Backup

- The following slides are some screen grabs of the quote from amazonws.com
- Remember, this site is intended to show that their AWS solution is attractive, so they add some quite high costs to the alternative ‘in-house’ solution. Nevertheless, it is a useful confirmation we are not too far off the mark.
- The notes that go along with the quotes document what they have assumed very well, and the notes and links are instructive.

On-Premises - Server Costs

Server Hardware Costs

Server Hardware Costs							
# of Servers	**# of Cores**	**RAM (GB)**	**Units (U)**	**Power (KW)**	**Unit Cost**	**Unit Discount**	**Total Cost**
1200	16	64	2400	660	$ 7,776	25%	$ 6,998,400
20	16	64	40	11	$ 7,776	25%	$ 116,640
1220			**2440**	**671**			**$ 7,115,040**

Total Server Hardware cost	$	7,115,040
Server Hardware Maintenance cost for 3 Yrs. (@15%/Yr.)	$	3,201,768
Total number of Racks required (1 Rack=42U, 28U occupied by servers, 4U by ToR switches and PDUs)		88
Total Peak power consumed (kW)		671

Rack Infrastructure Costs

Rack Chassis with PDU (@$3500/rack) cost	$	308,000
PDUs, dual 280V per rack (@$540 each, 2/rack for HA) cost	$	95,040
Top of Rack Switch (48-port 10/100/1G, $5,000 each, 2/rack for high availability)	$	880,000
Rack and Stack one-time deployment cost ($250/server)	$	305,000
Provision for spare servers for 3 Yrs. (@5% spare capacity/Yr.)	$	1,547,521
Total Rack costs (rack infrastructure and server hardware)	**$**	**13,452,369**

Server Software Costs (Host OS)

Total number of Windows licenses required		0
Windows license list price (unit cost for 2 licenses)	$	4,810
Windows license discounted price (unit cost for 2 licenses)	$	3,608
Windows licenses cost	$	-
Windows Software Assurance cost (3 Yrs.)	$	-
Windows Licenses and Software Assurance (3 Yrs.)	**$**	**-**
Total Server Cost (Hardware and Software) - 3 Yr.	**$**	**13,452,369**

AWS - EC2 Costs

EC2 Instance Costs (3 Yr.) – On-Demand and Reserved Instances

3 Yr. Partial Upfront Reserved Instance			
AWS Instance	**Upfront**	**Hourly**	**Total Costs**
r3.2xlarge	$ 5,330	$ 0.136	$ 10,696,646
r3.2xlarge	$ 5,330	$ 0.136	$ 178,277
Total Cost:			**$ 10,874,923**

Total costs = (upfront cost + hourly cost*8,784 hours/yr.*3 years)* # of instances (Applied to the whole term whether or not you're using the Reserved Instance)

On-Demand			
AWS Instance	**Upfront**	**Hourly**	**Total Costs**
r3.2xlarge	$ -	$ 0.780	$ 12,332,736
r3.2xlarge	$ -	$ 0.780	$ 82,218
Total Cost:			**$ 12,414,954**

Total costs = (hourly cost*8,784 hours*3 years*utilization)* # of instances (Hourly usage fee charged for each hour you use the instance)

Lowest Priced Instance		
Instance	**Cost**	**Type**
r3.2xlarge	$ 10,696,646	3 Yr. Partial Upfront RI
r3.2xlarge	$ 82,218	On-Demand
Total Cost:	**$ 10,778,864**	

EC2 Costs (3 Yr.) **$ 10,778,864**

EC2 Reserved Instances discounts (If Applicable)

Facilities Costs (data center space, power and cooling)) - On-Premises

Total Power consumed by servers (kW)		671
Metered cost per kWH	$	0.28
Estimated power cost/month	$	135,273.60
Monthly cost to operate a rack	$	1,800.00
Total rack costs/month	$	158,400.00
Total monthly Facilities costs	$	293,673.60
Facilities costs - On-Premises (3 Yr.)	**$**	**10,572,250**

Server cost break-down

Server cost break-down		
Category	**Cost**	**% of Total Cost**
Hardware	$ 13,452,369	56%
Software	$ -	0%
Operating Costs (3 Yrs.)	$ 10,572,250	44%
Total	**$ 24,024,619**	**100%**

Total server costs, including operational cost (3 Yr.) **$ 24,024,619**

EC2 Reserved Instances

AWS Instance	Instance type	# Instances	Upfront fee	Hourly	Total Cost
r3.2xlarge	3 Yr. Partial Upfront RI	1,200	$ 5,330	$ 0.136	$ 10,696,646
r3.2xlarge	On-Demand	20	$ -	$ 0.000	$ -

Total fee	**$**	**10,696,646**
Discount Tier Applicable		10%
AWS Business Support (EC2)	**$**	**417,794**
EC2 Costs (3 Yr.) after discount and support	**$**	**10,118,772**

On-Premises - Storage Costs

IOPS specified in addition to raw capacity; mix of HDD and SSD used

NAS Costs

Starting capacity/raw capacity (TB)		1,000
Starting capacity/raw capacity (GB)		1,024,000
Capacity after OS Penalty (~7%, capacity OS recognizes) (GB)		952,320
Usable capacity based on RAID (RAID 10 assumed) configuration (GB)		476,160
$/raw GB purchase price	$	6.60
Discounted $/raw purchase price (50% storage hardware discount applied)	$	3.30
Acquisition Cost of NAS storage	**$**	**3,379,200**

Storage backup cost

Total amount of storage to be backed up (TB)		1,000.00
Total amount of storage to be backed up (GB)		1,024,000
Type of Tape Library used		LTO-5
Max uncompressed speed (MB/s) for Tape Library		140
Max uncompressed speed - TB/day		11.54
Backup Window Time(hr.)		8
TBs processed/drive for backup window		3.85
Number of Tape drives required		261
Tape Library price/drive	$	1,800
Backup cost (3 Yr.)	**$**	**469,800**

AWS - Storage Costs

EBS Storage - General Purpose (SSD) or PIOPS volumes used depending on number of IOPS and capacity

EBS Costs - Equivalent to On-Premises NAS environment

Starting capacity (GB)		476,160
Equivalent EBS storage volume		General Purpose (SSD)
IOPS used for equivalent storage volume		48000
Number of EBS volumes required		477
EBS volumes cost/month		52,377.60
EBS IOPS cost/month		-
Initial snapshot cost(one-time)		45,235.20
EBS incremental snapshots cost/month		-
Total EBS cost /month		52,378
EBS Costs (3 Yr.) with IOPS & Capacity	**$**	**1,930,829**
EBS Costs (3 Yr.)	**$**	**2,047,426**
AWS Business Support (EBS)	**$**	**116,597**
Total AWS Storage Costs (3 Yr.) including support		**$2,047,426.17**

Storage Overhead cost (data center space, power, cooling, storage administrator)

Typical TB managed by a storage admin/Yr.		1000
Storage Admin Costs (3 Yr.)	**$**	**450,000**
Amount of TBs hosted by a single rack (TB)		1000
Number of racks required		1
Monthly cost to operate a rack	$	1,500
Total data center space, power, cooling costs (3 Yr.)	**$**	**54,000**

Storage cost break-down

Storage cost break-down		
Category	**Cost**	**% of Total Cost**
Raw Capacity (Incl. IOPS)	**$ 3,379,200**	78%
Backup	**$ 469,800**	11%
Overhead (excl. storage admin)	**$ 54,000**	1%
Storage Admin	**$ 450,000**	10%
Total	**$4,353,000**	**100%**

Total Storage Costs (3 Yr.) **$ 4,353,000**

On-Premises - Networking Costs

Networking Hardware and Software Costs

Network overhead cost as a % of server hardware acquisition cost	20%
Network hardware and software cost	$ 2,690,473.84
Network hardware and software maintenance/Yr.	15%
Maintenance cost (3 Yr.)	$ 1,210,713.23
Total Network Hardware and Software costs (3 Yr.)	**$ 3,901,187**

Bandwidth Costs (On-Premises)

Size of Network Pipe (Mbps)	10000
Peak/Avg. Ratio	1
Average Bandwidth	10,000.00
On-premises Bandwith costs/Mbps	$ 11.00
Bandwith costs/month	$ 110,000.00
Avg. data transferred per month (TB)- Inbound + Outbound	3089.9
Avg. data transferred per month (TB)- North/South	617.98
Avg. data transferred per month (TB) - Outbound	308.99
Bandwith costs - On-Premises (3 Yr.)	**$ 3,960,000**

Network Admin Costs

Network admin effort as % of total IT admin effort	8%
Avg. burdened salary for your Network Admin	$ 150,000
IT labor cost (1 Yr.)	$ 366,000
Network admin costs (1 Yr.)	$ 29,280
Network admin costs (3 Yr.)	**$ 87,840**

Total Networking Costs (3 Yr.) $ 7,949,027

AWS - Data Transfer Costs

Monthly Data Transfer Out (TB) 308.99

Data Transfer Costs

	EU (Ireland)	Tier(GB)	Monthly Cost
First 1 GB per month	$ -	1	$ -
Up to 10 TB per Month	$ 0.09	10240	$ 921.60
Next 40 TB per Month	$ 0.09	40960	$ 3,481.60
Next 100 TB per Month	$ 0.07	102400	$ 7,168.00
Over 350 TB per Month	$ 0.05	162805	$ 8,140.26

Total monthly data transfer costs	$ 19,711.46
AWS Business Support (data transfer)	**$ 42,852**
Data Transfer Costs (3 Yr.) including support	**$ 752,464**