

RACF/SDCC Network Operations Status Report (Apr 26, 2018)


Alexandr ZAYTSEV
alezayt@bnl.gov

Project: RACF/SDCC Network Operations & Interventions

Chart

This chart shows the number of issues **created** vs the number of issues **resolved** in the last 224 days.

Previous Report




Project: Post RHIC Migration BCF & CDCE Cleansing & Service Normalization (2018Q1-3)

Chart

Previous Report

This chart shows the number of issues **created** vs the number of issues **resolved** in the last 224 days.




Recent Developments (1)

- Replacement of s805 (subnet 6) and s780/s781 (10.42.34.x) switches
 - **DONE**; legacy patch panel and underfloor copper uplinks removal is still in progress
- Replacement of the legacy s246 switch serving subnet 59 (1 GbE)
 - **DONE**; legacy patch panel and underfloor copper uplinks removal is still in progress
- Retirement and repurposing of s780/s781 switch pair (1 GbE distribution for subnet 10.42.34.x)
 - **DONE**; one of the switches is reused as a “warm” backup for HPSS tape silo control network switch in CDCE (r41-7)
- IC cluster extension (2 racks)
 - **DONE**; new model HPE 1920S switches kept
- New SF installation in the BGL room
 - 3 new compute node racks are fully deployed
 - All 6x 7060SX2 BoR switches are deployed and configured; 2 of them will have to move once the 4th rack arrives (second half of May'18)
 - Several issues with SF Bright deployment system as used via SF VPN were solved in the process and should not affect the new SF racks deployment
 - The decommissioned Dell blade chassis and the former SF FE switch pair are yet to be picked by SF (during this week)

Recent Developments (2)

- Conversion of compute node racks to spine-and-leaf infrastructure in CDCE
 - All 21x ToR switches needed for replacement of legacy Dell/Force10 Exascale s705 (row 50S) and Brocade SX-1600 systems (row 48N) are received, installed and configured
 - 5 interventions scheduled by Chris Hollowell and agreed with ITD to complete the conversion of ATLAS Tier-1 and ATLAS Tier-3 sites (NETOPS-104):
 - Apr 26, 2018 (DONE) – ATLAS Tier-1 racks in row 50S
 - May 8, 2018 – All ATLAS Tier-3 racks in rows 47N, 50S, 50N and 51S (to coincide with the physical migration of the ATLAS Tier-3 GPFS rack from 4-6 to 41-2, and rack ATLAS Tier-3 rack 48-21 move to position 47-22)
 - May 21, 2018 – ATLAS Tier-1 racks in row 47N
 - May 24, 2018 – ATLAS Tier-1 racks in row 51S
 - Thus, both s705 and Brocade SX-1600 switch systems are expected to be gone by the end of May'18
- Fiber infrastructure deployment in CDCE rows 41-43, 46 to accommodate the following new systems to be deployed in May-June 2018 timeframe:
 - STAR GPFS (rows 42-43; position 43-2 still need to be liberated)
 - Home directory NFS appliance (row 41)
 - Shared RHEV NFS appliance (row 41)
 - Belle II dCache storage upgrade (row 46)

Legacy Network Equipment Replacement Timeline



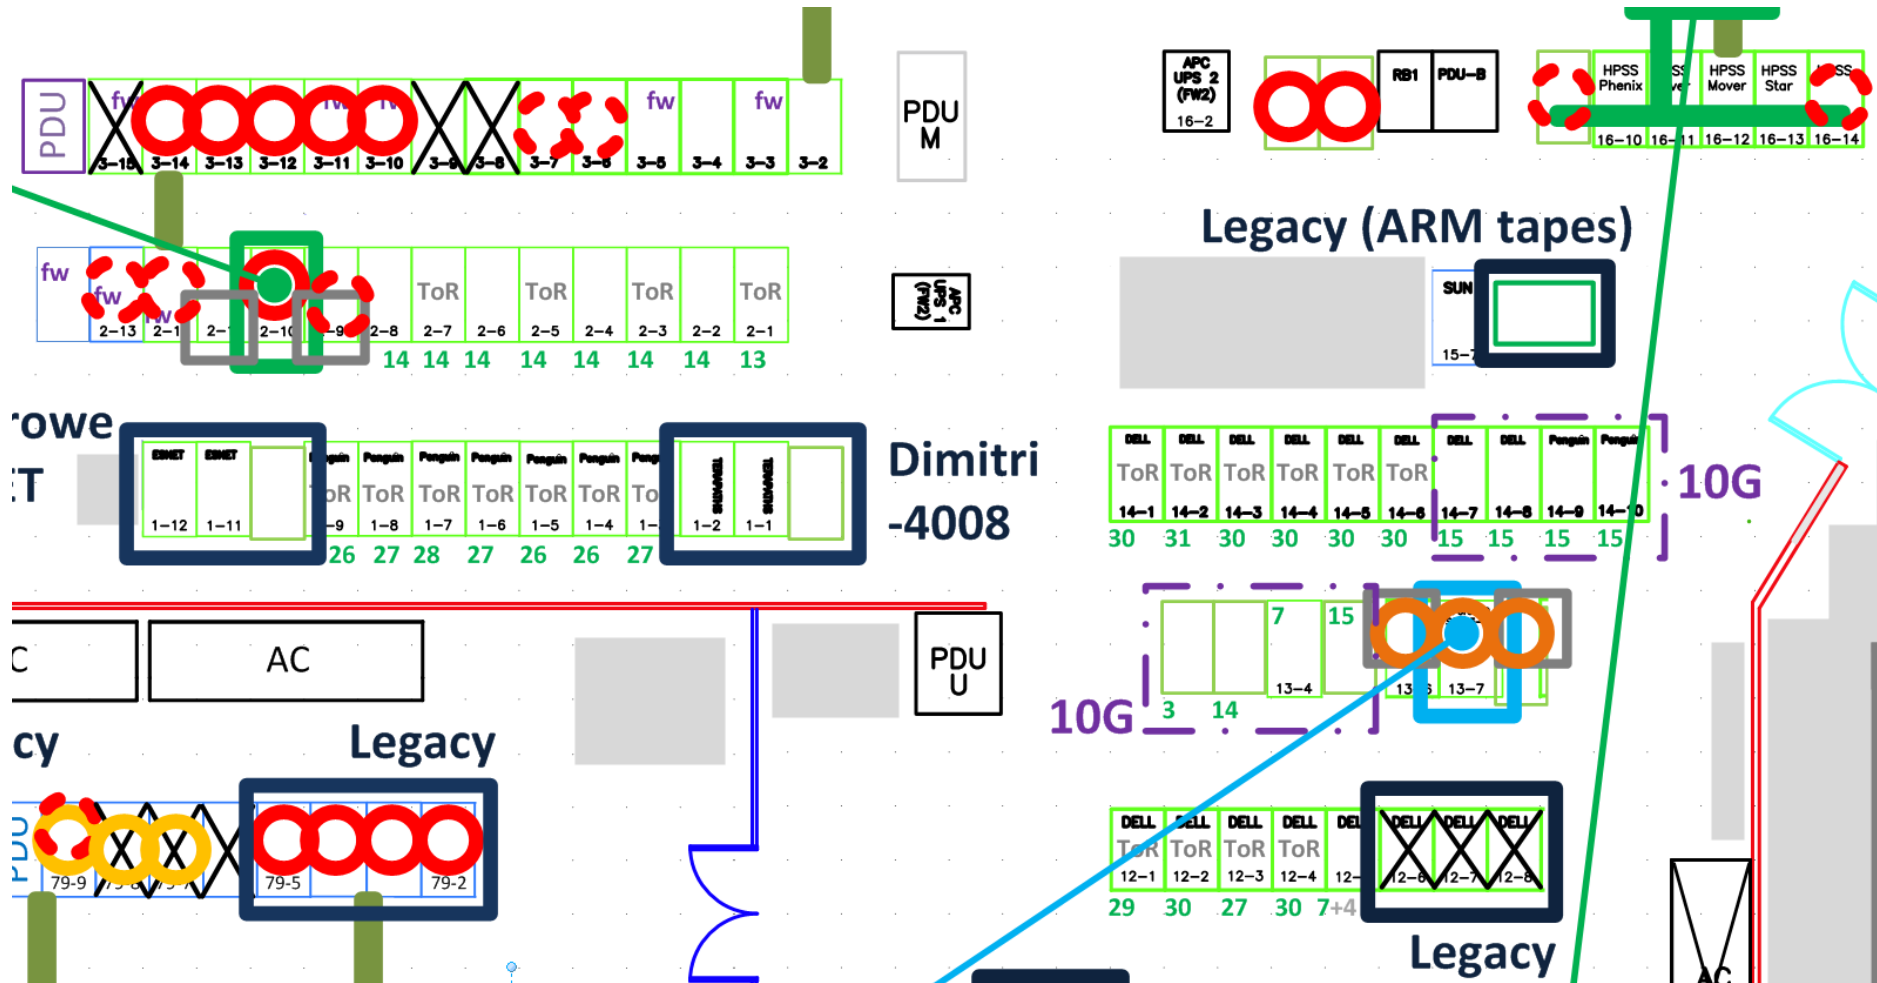
The RHIC Farm Networking

- The issue of the Dell/Force10 Exascale switch systems s803 (RCF area) and s804 (main BCF area) going EOS 2020
 - The existing two Exascale systems (s803 and s804) are already the source of significant operational concern due to the frequent partial failure of the 90-port HDE line cards, and an increased ETA for receiving the replacement components from the vendor as of 2017
 - Stretching the production life of these systems beyond EOS, until FY21 is highly undesirable
 - The solution we are converging on is to replace both s803 and s804 systems with ToR switches connected to the existing spine group in the CDCE much like it is done for the “ATLAS like” racks if compute nodes in CDCE (with the difference that the racks with 2U nodes are going to be wired in pairs to one leaf in order to avoid wasting the port capacity)
 - 4 ToRs are needed to handle row 21 (RCF – STAR) on s803
 - 2 ToRs are needed to handle row 25 (RCF – Small Experiments) on s803
 - 22 ToRs are needed to handle rows 1, 2, 12, 14 (main BCF – mix of STAR/PHENIX) on s804, with the loss of 9 nodes in rows 12, 14 to the /27 (29 nodes per rack) IPv4 limit
 - The aim is to perform (STEP 1) the s803 replacement with 6 ToR in June 2018 in parallel to the cut-in of the new Home Directories NFS intervention (so, no special intervention day need) and (STEP 2) the s804 replacement with 22 ToRs sometime in July-December 2018 timeframe, with one day of downtime for both STAR and PHENIX parts of the Farm. The procurement is planned to be executed in FY18 for both sets of switches (about 15-18 of these ToR are expected to be reused for the new RHIC Farm racks in FY19-21 period, as their life will stretch until FY23)
 - The extension of life for the Arista MXP line cards beyond 2021 EOL seems doable, so the 10 GbE RHIC Farm switches can be kept operational until 2020 when the current 10 GbE connected RHIC Farm nodes are expected to retire in place

The RHIC Farm Networking (RCF area)



The RHIC Farm Networking (main BCF area)



Questions & Comments