




Systems Engineering Update

Matt Pounsett, OARC 31
Austin, USA




DNS-OARC

Domain Name System Operations Analysis and Research Center

See The Written Report

- For most subjects, please see the written report at <https://indico.dns-oarc.net/event/32/contributions/736/attachments/702/1162/201910-OARC31-Systems-Report.pdf>




DNS-OARC

Domain Name System Operations Analysis and Research Center

Looking for Assistance

- There are a long list of engineering projects that need to get done. Not nearly enough time to complete them all.
- OARC is going to be looking for assistance on two fairly straight-forward, self-contained items in order to reduce draws on Engineering's time.
- See Matt or Keith for more information




DNS-OARC

Domain Name System Operations Analysis and Research Center

File Server Hardware Issues

- Two issues that need detailed troubleshooting
- Fs1 crashes on a regular basis
 - Usually load-related (e.g. research or generating checksums of its files)
 - Also seems to go down on an irregular schedule (approximately monthly) with no obvious trigger
- Fs2's new disks in a new external JBOD (both one year old) generate ZFS checksum errors on high write load




DNS-OARC

Domain Name System Operations Analysis and Research Center

Network Config Audit/Cleanup

- OARC's routers at Fremont & Ottawa need some review and cleanup
- Half-complete IPSEC config needs careful removal
- Peering configs need review
- Old ACLs with no annotations need review, update, removal, docs

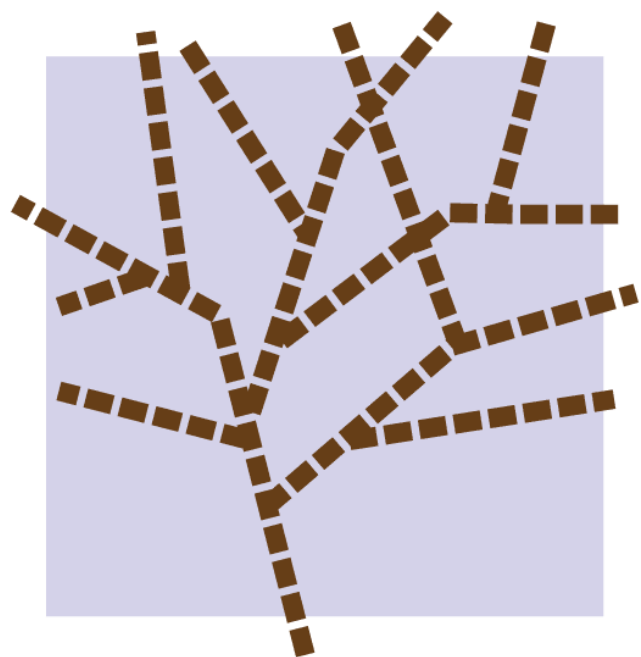


DNS-OARC

Domain Name System Operations Analysis and Research Center

OARC Needs New Secondary DNS


- Public benefit service previously provided by ISC is going away as they phase out the platform it is on, later this year
- Anycast, *XFR, mastered/signed by OARC
- OARC production zones (fwd/rev), AS112 production zones (fwd/rev), plus some quiet sub zones for testing, research, etc.
- No traffic stats available
- Email matt@dns-oarc.net for more information



DNS-OARC

Domain Name System Operations Analysis and Research Center

New Data Storage Architecture




DNS-OARC

Domain Name System Operations Analysis and Research Center

See the Written Proposal

- In the Indico timetable at: <https://indico.dns-oarc.net/event/32/contributions/736/attachments/702/1194/Filesystem-Clustering.pdf>




DNS-OARC

Domain Name System Operations Analysis and Research Center

Existing Infrastructure is Aging

- OARC has been replacing disks every year, but not any other components
- Disk replacement current takes 8 years, which is much longer than the average replacement period for hardware in the industry
- Growth in total storage has not been accompanied by necessary growth in processor and memory




DNS-OARC

Domain Name System Operations Analysis and Research Center

Existing Infrastructure is Broken

- Significant physical damage to one server
- Another reboots at random intervals, approximately monthly, and on heavy read load
- Another generates mass CRC errors on writes, on effectively brand new disks




DNS-OARC

Domain Name System Operations Analysis and Research Center

Cost to Repair

- US\$80,000 expense required to update the infrastructure




DNS-OARC

Domain Name System Operations Analysis and Research Center

Existing Architecture is Expensive

- Vertically scaled
- Inefficient use of storage
- High labour cost to maintain data
- Unnecessary complexity for consumers of the data




DNS-OARC

Domain Name System Operations Analysis and Research Center

Clustered Filesystem


- Horizontally scaled
- Better performance
- More efficient use of storage
- Reduced manual labour maintaining data sets
- Less complexity and surprise for data consumers



DNS-OARC

Domain Name System Operations Analysis and Research Center

Budgeting for a new Architecture




DNS-OARC

Domain Name System Operations Analysis and Research Center

Settled on Ceph

- Provides a networked filesystem, an object store, and block devices
- Excellent redundancy and high availability properties
- Horizontally scales well



DNS-OARC

Domain Name System Operations Analysis and Research Center

Budgeting is Complicated


Maximum Total Storage

Minimum Possible Budget

Spindles per Server

Balanced Bandwidth per Server

Scaling RAM with Storage




DNS-OARC

Domain Name System Operations Analysis and Research Center

Budgeting Approach

- Pick a platform
- Separate base server cost, from cost of RAM, from cost of disks
- Calculate all possible disk combinations and required RAM
- Calculate cost of an infrastructure based on number of servers required for each disk combination
- Repeat




DNS-OARC

Domain Name System Operations Analysis and Research Center

Pick a Platform / Base Cost

Description	Units	Price	Additions	Subtractions	Total
Base Cost	1	\$2529.00	\$2529.00		
Intel Xeon Silver 4210 2.2G 10C/20T 85W	1	\$251.33	\$251.33		
Remove base RAM	-1	\$238.13		-\$238.13	
PERC H740P Controller	1	\$0.00	\$0.00		
480GB SSD Read Intensive (boot)	1	\$250.70	\$250.70		
240GB SSD Mixed Use (journal)	1	\$294.68	\$294.68		
Remove default disk	-1	\$137.60		-\$137.60	
iDRAC9 Enterprise	1	\$307.25	\$307.25		
Riser Config 1	1	\$0.00	\$0.00		
Broadcom 57416 Dual Port 10Gb BASE-T	1	\$0.00	\$0.00		
QLogic FastLinQ 41112 Dual Port SFP+ Full Height	1	\$300.97	\$300.97		
SFP+ SR Optic	1	\$150.17	\$150.17		
Single Hot-Plug PSU (495W)	1	\$0.00	\$0.00		
NEMA 5-15P cable	1	\$0.00	\$0.00		
Standard Bezel	1	\$30.79	\$30.79		
Rails + Cable Management	1	\$118.75	\$118.75		
			\$4,233.64	-\$375.73	\$3857.91




DNS-OARC

Domain Name System Operations Analysis and Research Center

Memory and Disk

Size (GB)	Rank	Frequency (MT/s)	Price	Cost per GB
8	Single	2666	\$238.13	\$29.77

Size (TB)	Technology	Price	Cost per TB
8	SATA	\$690.53	\$86.32
12	SATA	\$1061.24	\$88.44
12	SAS	\$1092.65	\$91.05
14	SATA	\$1318.85	\$94.20




DNS-OARC

Domain Name System Operations Analysis and Research Center

Calculate Possible Configurations

Storage Size	Base Server	HDD Cost	RAM Cost	Total Cost
12	\$3857.91	\$1061.24	\$952.52	\$5871.67
24	\$3857.91	\$2122.48	\$952.52	\$6932.91
36	\$3857.91	\$3183.72	\$1428.78	\$8470.41
48	\$3857.91	\$4244.96	\$1905.04	\$10007.91
60	\$3857.91	\$5306.20	\$2381.30	\$11545.41
72	\$3857.91	\$6367.44	\$2381.30	\$12606.65
84	\$3857.91	\$7428.68	\$2857.56	\$14144.15
96	\$3857.91	\$8489.92	\$3333.82	\$15681.65
108	\$3857.91	\$9551.16	\$3810.08	\$17219.15
120	\$3857.91	\$10612.40	\$3810.08	\$18280.39




DNS-OARC

Domain Name System Operations Analysis and Research Center

Scale it Out


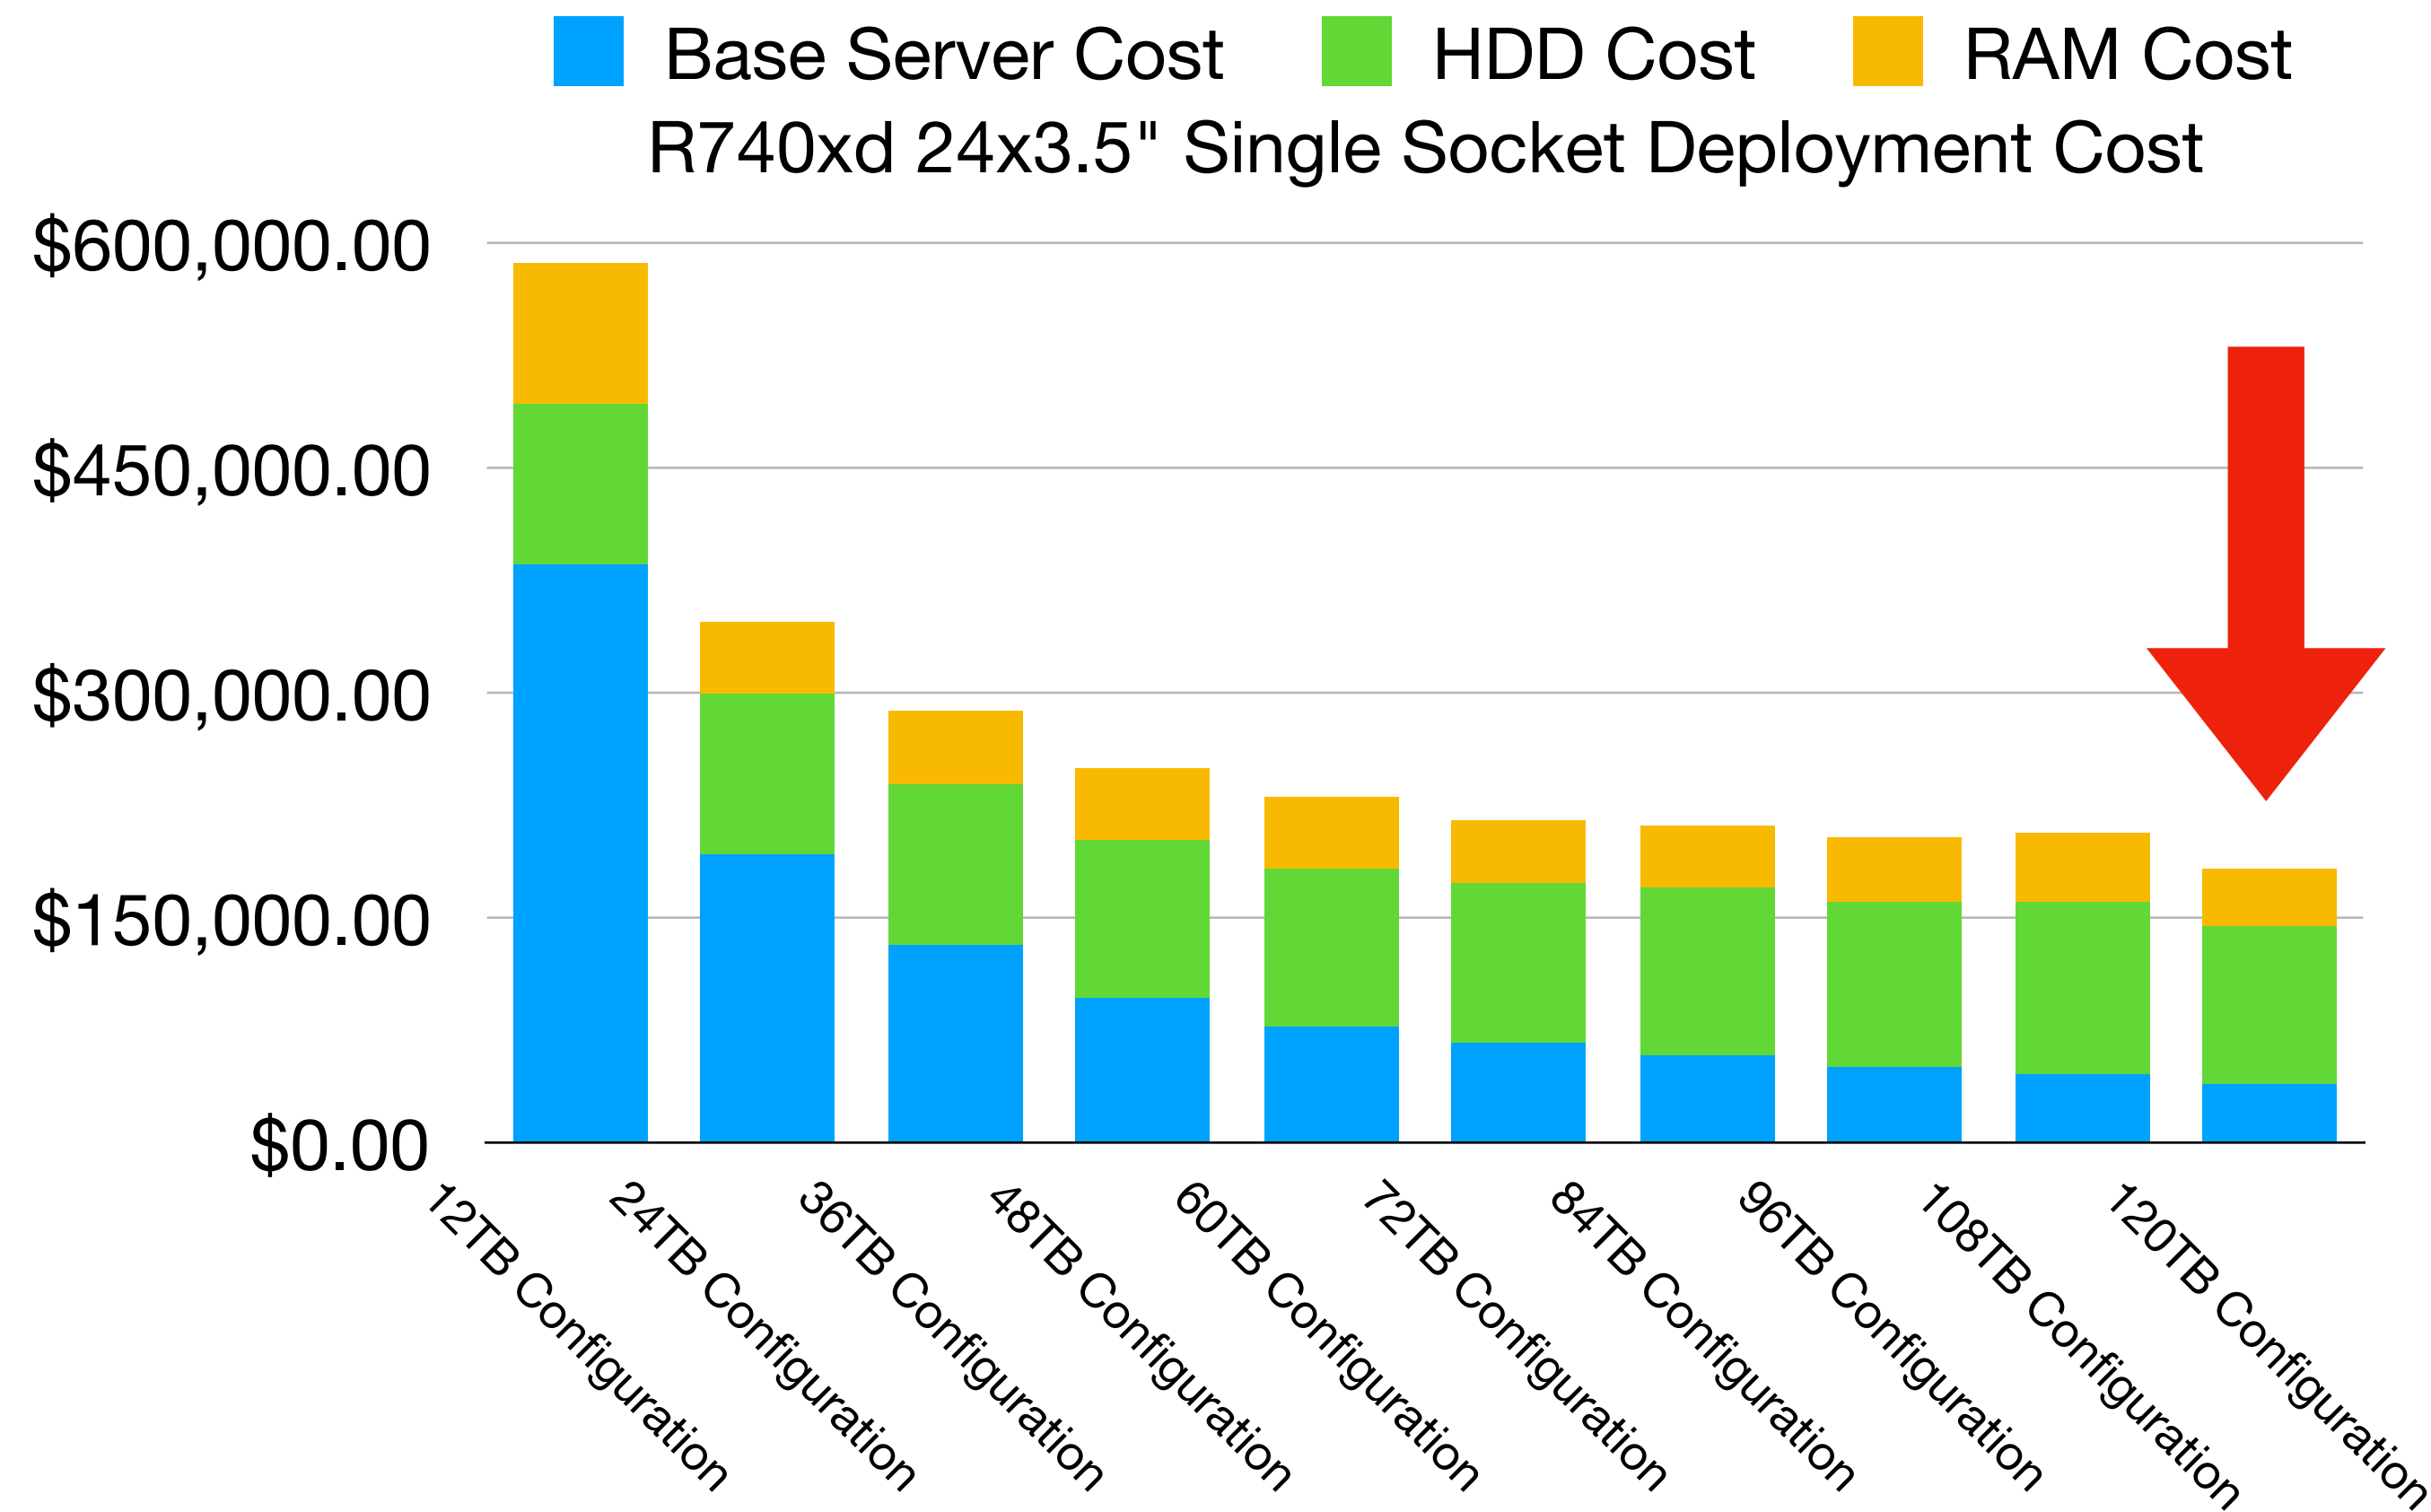
Configuration Name	Total Server Storage	Servers Required	Base Server Cost	HDD Cost	RAM Cost	Total
12TB Configuration	12	100	\$385,791.00	\$106,124.00	\$95,252.00	\$587,167.00
24TB Configuration	24	50	\$192,895.50	\$106,124.00	\$47,626.00	\$346,645.50
36TB Configuration	36	34	\$131,168.94	\$108,246.48	\$48,578.52	\$287,993.94
48TB Configuration	48	25	\$96,447.75	\$106,124.00	\$47,626.00	\$250,197.75
60TB Configuration	60	20	\$77,158.20	\$106,124.00	\$47,626.00	\$230,908.20
72TB Configuration	72	17	\$65,584.47	\$108,246.48	\$40,482.10	\$214,313.05
84TB Configuration	84	15	\$57,868.65	\$111,430.20	\$42,863.40	\$212,162.25
96TB Configuration	96	13	\$50,152.83	\$110,368.96	\$43,339.66	\$203,861.45
108TB Configuration	108	12	\$46,294.92	\$114,613.92	\$45,720.96	\$206,629.80
120TB Configuration	120	10	\$38,579.10	\$106,124.00	\$38,100.80	\$182,803.90



DNS-OARC

Domain Name System Operations Analysis and Research Center



Find the Cheapest Deployment



DNS-OARC

Domain Name System Operations Analysis and Research Center



Find the Cheapest Deployment



DNS-OARC

Domain Name System Operations Analysis and Research Center

Repeat




DNS-OARC

Domain Name System Operations Analysis and Research Center

Dell is ~~Expensive~~ Over Engineered

- Dell disks are 4x the cost of everyone else
- Server hardware similarly much more expensive than some other options
- The architecture provides redundancy, so we don't need the same kind of expensive redundancies and complex engineering in each component
- Repeat everything with Dell servers but 3rd party HDD sales



DNS-OARC


Domain Name System Operations Analysis and Research Center

Budgeting

- Multiple server types with different capacity and redundancy requirements
- Scaled out for anticipated 2020 data size

US\$125k

- Will be obtaining quotes from cheaper, more customizable vendors



DNS-OARC


Domain Name System Operations Analysis and Research Center

A night view of a city skyline, likely Los Angeles, with many lights and buildings visible against a dark sky. The text "LOS ANGELES NOVEMBER, 2019" is overlaid in white, serif font.

LOS ANGELES
NOVEMBER, 2019



Discussion



DNS-OARC

Domain Name System Operations Analysis and Research Center