Available online www.ejaet.com

European Journal of Advances in Engineering and Technology, 2022, 9(11):52-58

Research Article ISSN: 2394 - 658X

Optimizing Expansions at Data Center

Anurag Reddy[1], Anil Naik[2], Sandeep Reddy[3]

[1]Head of Infrastructure Planning, Public CDN and Cloud, UC Berkeley, CA
Email: anurag_reddy@berkeley.edu; Ph: 408-784-5668
[2]Product Lead, Telecom & AR/VR, UC Berkeley, CA
Email: anilgnaik9@gmail.com; Ph: 408-219-5670
Senior Program Manager, Consumer Technology, University of Rochester, NY"
Email: sandeepreddy4488@gmail.com; Ph: 5124122747

ABSTRACT

This paper proposes a novel approach to data center expansion planning and execution, aiming to streamline the process and enhance responsiveness to dynamic compute demands. The traditional model, involving Commercial Readiness, Cage Readiness, Network HW Readiness, and Compute Provisioning, often leads to long lead times and high variability in short-term planning. The suggested paradigm introduces an intermediate "Medium Term Planning" phase to bridge the gap between site preparation and short-term demand, dissociating Cage/Network Readiness from compute expansions.

Key components of the proposed model include a Colocation Model for cage readiness forecasting, Compute Model for prioritizing Network Readiness work orders, and a Heat Map to trigger server execution based on CPU utilization. The Compute Execution Trigger identifies three factors influencing server deployment lead time, facilitating a Just-in-Time approach. Additionally, a Network Readiness Heat Map is proposed to monitor the readiness headroom for server deployments.

The benefits of this approach include mitigating cage/network bottlenecks ahead of compute demands, reducing server deployment lead time, and enabling a reactive capacity expansion strategy. A Total Cost of Ownership (TCO) analysis highlights the advantages of a Just-in-Time strategy over traditional batch-based deployments.

To implement this model, a recommended ticket structure is proposed, emphasizing dedicated BOM for Cage/Network Readiness post-Commercial closure. The BOM Scope in Network Readiness is designed to achieve a zero-touch BOM review for server expansions, further reducing lead time in the first compute batch.

Overall, this paradigm shift in data center expansion planning and execution promises increased agility, reduced deployment lead times, and improved responsiveness to evolving compute demands.

Key words: data center expansion, just-in-time model, Medium-Term Planning, reduced lead times, responsiveness, Colocation Model, Compute Model, Heat Map, Total Cost of Ownership (TCO) analysis, proactive monitoring, streamlined processes, network readiness SLAs, operational timelines, optimization, computational requirements.

INTRODUCTION

This paper introduces a transformative paradigm for data center expansion, departing from traditional processes and embracing a just-in-time (JIT) model. The current sequential approach of Commercial Readiness, Cage Readiness, Network HW Readiness, and Compute Provisioning often results in high variability and extended lead times. The proposed model advocates for a "Medium Term Planning" phase, effectively decoupling Cage and Network Readiness from compute expansions to enhance responsiveness. Key components include a

Colocation Model, Compute Model, and a Heat Map triggering server execution based on real-time CPU utilization. The paper explores the methodology, benefits, and presents a Total Cost of Ownership (TCO) analysis, contributing to the discourse on optimizing data center operations.

BACKGROUND

The existing operational workflow revolves around the progression of the Site Commercials until the completion of colocation commercial terms. Subsequent to this phase, the site attains the status of "Ready for Production" and enters a standby mode, anticipating short-term demand. The release of short-term demand traditionally occurs in batches, encompassing activities such as cage preparation, connection establishment, and the installation/provisioning of network equipment, all coinciding with compute deployment.

However, a notable challenge arises in the form of cage readiness, a crucial aspect of short-term operations that deviates from standard short-term planning norms due to its extended lead times and high variability. Recognizing this, we propose the implementation of an intermediate phase termed "Medium Term Planning." This phase is strategically designed to proactively prepare a site in advance, thereby mitigating the prolonged lead times associated with short-term planning execution.

The core principle behind this proposed shift is to eliminate the conventional approach of batch releases and embrace a just-in-time (JIT) methodology. Work orders for compute work orders would be triggered by CPU utilization reaching approximately the target threshold, ensuring that the site is consistently prepared to absorb racks. This transition is envisioned to bring about a substantial reduction in compute deployment lead times, transforming the timeline from months to mere days. The objective is to ensure that all compute expansions transpire within the month of release.

A pivotal advantage of this proposed model lies in its ability to proactively identify and address potential bottlenecks or challenges on the cage side well in advance of the actual compute demand. By catching these issues ahead of time, the operational framework gains the flexibility to devise and implement solutions before they escalate into urgent concerns during periods of high compute demand. This proactive and JIT-oriented approach is poised to significantly enhance the efficiency and responsiveness of the overall data center expansion process.

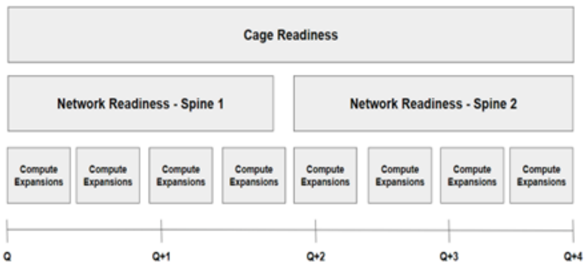

OPERATIONAL TRANSFORMATION

This strategic initiative aims to liberate operational processes by disentangling Cage and Network Readiness from short-term planning, fostering increased agility and responsiveness in data center expansions. The current operational challenge lies in the heavy reliance on integrating Cage Readiness and Network Readiness into short-term planning for executing server-related tasks.

Server only expansion:

- short lead time
- low variability

Cage readiness/Network Readiness:

- long lead time
- high variability

Server-only expansion excels in short lead times and minimal variability, in stark contrast to the extended lead times and high variability associated with Cage and Network Readiness.

Strategic Shift Medium-Term Planning advocates for a strategic transition by relocating Cage and Network Readiness work orders to medium-term planning. This shift aims to fully disconnect these readiness phases from short-term compute expansions. Below are the anticipated Benefits of shifting to medium term planning:

1) Enhanced Readiness: Ensures the completion of signing and commercial closure for all sites, achieving a 100% readiness state for deployment.
2) Reduced Planning Variability: Mitigates prolonged lead times and variability in short-term planning, enabling swift responses to demand signals and facilitating server-only expansion.

Real-world scenarios vividly demonstrate the significant consequences of project delays, specifically stemming from challenges like the intricacies of DF Turn-up, complexities in spines/DCI provisioning, and protracted design discussions embedded in the existing operational framework.

For heightened stakeholder engagement, a meticulously crafted rack design template is extended to each site. This collaborative effort, coupled with the utilization of long-term compute forecasts, serves as a comprehensive guide for strategic planning. The dynamism introduced by dynamically triggering racks requests, based on real-time CPU Utilization data, enhances adaptability and responsiveness in the operational approach.

In optimizing server execution, there is a deliberate focus on achieving a substantial reduction in compute timelines, envisioning a streamlined process that shortens the duration to a matter of days. In this context, post-arrival operations for integrated racks are simplified, requiring only the straightforward connection of ToR Uplinks to the spines. This nuanced approach ensures a responsive, efficient, and agile operational framework.

A. Integrated Operational Framework

1) Colocation Model: The model assumes a pivotal role, orchestrating inputs crucial for the cage readiness forecast. Executing on a quarterly basis, it not only anticipates future requirements but also issues tickets for the meticulous implementation of Cage Readiness & Network Readiness.
2) Compute Model: the model takes charge of the monthly landscape, strategically allocating priorities to Network Readiness Work Orders. This intricate process involves identifying sites on the verge of compute deployment shortages, ensuring a proactive coverage spanning 2 quarters of compute demand.
3) Capacity Heat Map: The dynamic Heat Map emerges as the linchpin, serving as the trigger for server execution. It operates by monitoring CPU utilization, and with the integration of a novel adjustment deployment lead time, swiftly responds to fluctuations. Upon reaching a predefined threshold, it autonomously releases a rack-equivalent ticket for seamless expansion. This finely-tuned mechanism not only optimizes resource utilization but also ensures a responsive and agile data center ecosystem.

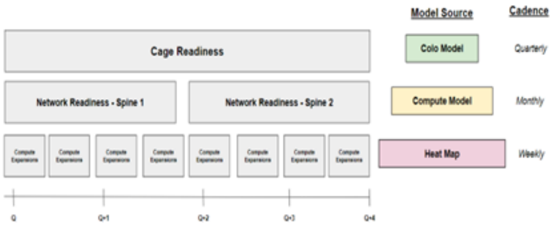

B. Efficient Server Deployment Activation Strategy

In the orchestrated sequence, server Bill of Materials (BOMs) become available for consultation through Cage Readiness and, with Network Readiness physically in place on-site, poised to promptly absorb compute executions. The lead time for server deployments is intricately governed by three pivotal factors:

1) Server Racks Production Lead Time (Team Responsible: Rack Integration)
2) Transportation Lead Time (Team Responsible: Logistics)
3) Installation (Team Responsible: DC OPs)

To streamline operations, a committed lead time is sought from each responsible team. This commitment aligns with determining the CPU utilization trigger point that signifies the initiation of a server work request execution. For instance, in US sites where transportation spans 2/3 days and on-site installation, facilitated by metrotechs, takes less than 1 week, synchronization with server racks production at integrators meeting these criteria enables a server deployment execution in under 2 weeks. This strategic approach ensures that CPU utilization hovers around 60% until the release of an order, averting the deployment of larger server executions based on forecasted growth demands that might take longer to materialize.

C. Navigating Network and Compute Deployments

While Network Compute will be the source to trigger an expansion, we need a way to monitor the network readiness to absorb compute work orders requests as well as what is the timely coverage which that network readiness provides us.

1) *Network Readiness Execution Dashboard:* This dynamic dashboard meticulously showcases the current headroom on sites, courtesy of Network Readiness work orders, to seamlessly absorb compute work requests.
2) *Network Readiness Heat Map:* This heat map serves the goal of comparing Cage Readiness and Network Readiness Coverage concerning the compute forecast.

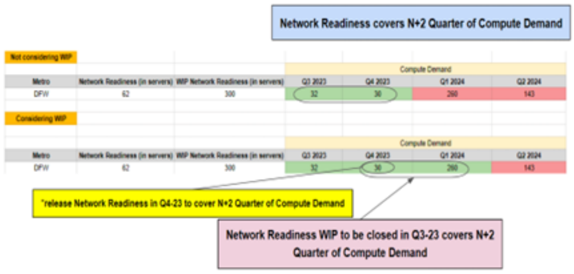

D. Benefits of Enhanced Readiness Monitoring

1) Anticipation and mitigation of Cage/Network bottlenecks ahead of compute demands
2) Facilitation of seamless server execution and reduction in deployments lead time
3) Adoption of a Just-in-Time approach for capacity demands
4) Elimination of capacity expansion driven solely by long-term forecasted growth rates, embracing a new reactive approach
5) Specialized teams for Network Readiness & Compute Readiness Execution
6) Minimization of the undeployed capital horizon, optimizing resource allocation.

E. Total Cost of Ownership (TCO) Analysis

To substantiate our perspective, we meticulously conducted a comprehensive Total Cost of Ownership (TCO) analysis, delving into various deployment strategies. Our chosen approach involved breaking down the deployment into smaller batches and considering critical factors such as colocation (colo), depreciation, shipping, and installation costs. The findings underscored a significant spike in costs, notably observed between 1 Batch and 8 Batches.

This detailed analysis not only illuminated the financial implications but also provided nuanced insights into the operational dynamics. It showcased the intricate relationship between batch sizes and associated costs, prompting a strategic reflection on the optimal balance

Deployment Strategy	Total Cost	Colo Cost	Depreciation Cost	Shipping Cost	Installation Cost
1 Batch	$4,412,000	$2,592,000	$1,250,000	$510,000	$60,000
2 Batches	$3,471,500	$1,944,000	$937,500	$520,000	$70,000
3 Batches	$3,171,333	$1,728,000	$833,333	$530,000	$80,000
4 Batches	$3,031,250	$1,620,000	$781,250	$540,000	$90,000
5 Batches	$2,955,200	$1,555,200	$750,000	$550,000	$100,000
6 Batches	$2,911,167	$1,512,000	$729,167	$560,000	$110,000
7 Batches	$2,885,429	$1,481,143	$714,286	$570,000	$120,000
8 Batches	$2,871,125	$1,458,000	$703,125	$580,000	$130,000
9 Batches	$2,864,444	$1,440,000	$694,444	$590,000	$140,000
10 Batches	$2,863,100	$1,425,600	$687,500	$600,000	$150,000
11 Batches	$2,862,726	$1,411,855	$680,871	$610,000	$160,000
12 Batches	$2,871,083	$1,404,000	$677,083	$620,000	$170,000
13 Batches	$2,878,277	$1,395,360	$672,917	$630,000	$180,000
14 Batches	$2,888,214	$1,388,571	$669,643	$640,000	$190,000
15 Batches	$2,899,067	$1,382,400	$666,667	$650,000	$200,000
16 Batches	$2,911,063	$1,377,000	$664,063	$660,000	$210,000

As we surpass the threshold of 8 Batches, the incremental gains in dollar terms exhibit a diminishing trend. Yet, the emphasis on adopting Just-in-Time (JIT) strategies becomes even more compelling, unveiling a multitude of advantages. The strategic shift toward JIT methodologies not only enhances reactivity but also encompasses benefits such as streamlined operational efficiency, reduced lead times, minimized excess inventory costs, and the ability to promptly respond to dynamic market demands.

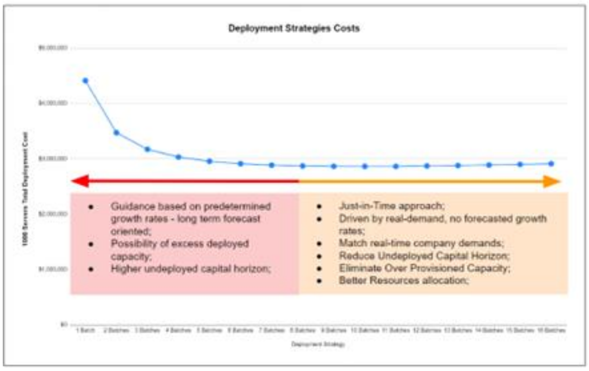

F. Suggested Work Order Framework

In implementing these innovative workflow configurations, it is imperative to fully separate Cage and Network Readiness tickets from compute expansions. Introducing a specialized Work Order type is essential, scheduled immediately after the closure of Site Commercials. This dedicated work order serves the purpose of initiating the tracking process for cage readiness planning and network readiness execution.

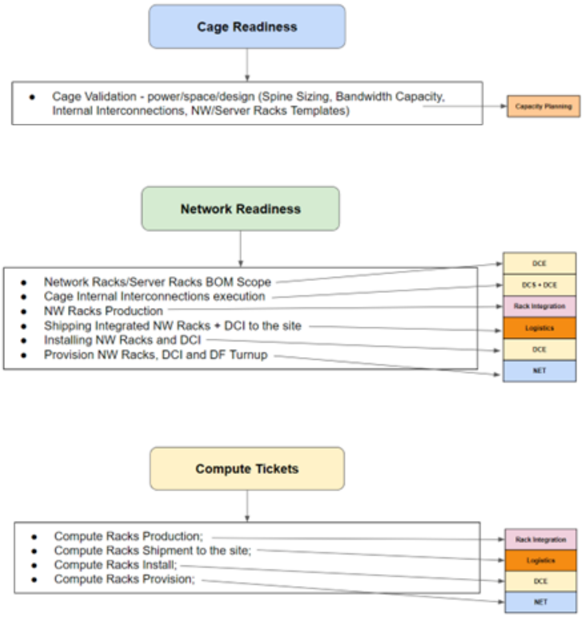

G. Streamlining Processes and Enhancing Readiness

Enhancing deployment efficiency through strategic streamlining: Tailoring Bill of Materials (BOM) processes and accelerating readiness measures, from Network Readiness to the first compute batch, while extending optimizations to Edge Site integration.

1) Bill of Materials (BOM) Scope for Server Racks in Network Readiness: The objective is to achieve a streamlined BOM review process for server expansions, where rack integration seamlessly obtains BOMs through the examination of Network Readiness tickets. The BOM scope for Network Readiness Server Racks will be predefined, considering cable lengths required for each position of the server racks within the cage.
2) Lead Time Reduction in the 1st Compute Batch: Initiating Network Readiness ahead of compute (shipping/deploying/pre-provision spines and DCI) highlighted the potential for further enhancements in expediting the first compute deployment on-site. Recognizing that certain tasks impeded the desired

speed within the current Network Readiness setup—specifically, DM set and the creation of a new colo—an agreement was reached with the Data Center Engineering (DCE) teams. Alongside Network Racks, 2/3 metals would be shipped (for resiliency in case of failure). One of them would be configured as DM, facilitating the creation of the new colo and spine in prodash. All metals would be maintained in "P" status with anycast disabled. This strategic approach enables additional pre-provisioning steps, aligning the lead time for the first batch with subsequent ones.

H. Network Readiness Service Level Agreements (SLAs)

Precision in operational timelines is paramount; therefore, all Network Readiness SLAs will be meticulously crafted, aligning with the timelines associated with long-term projected compute demands. These SLAs will be categorized and prioritized according to the following benchmarks:

Priority	INFOPs Timelines	Net Timelines	Total Time
P1 High	2 weeks	2 weeks	4 weeks
P2 High	3 weeks	3 weeks	6 weeks
P3 High	4 weeks	4 weeks	8 weeks
P4 Normal	6 weeks	6 weeks	12 weeks

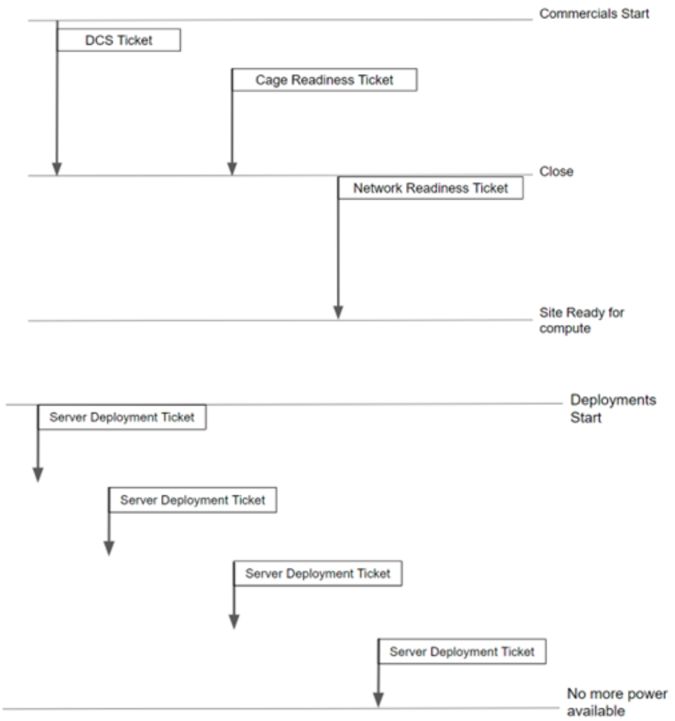

CONCLUSION

In conclusion, the proposed paradigm shift in data center expansion planning, adopting a just-in-time (JIT) model, presents a transformative approach. The introduction of Medium-Term Planning, dissociating Cage and Network Readiness from compute expansions, promises increased agility, reduced lead times, and improved responsiveness to dynamic compute demands. The strategic components, such as Colocation Model, Compute Model, and Heat Map, together with efficient server deployment strategies and Total Cost of Ownership (TCO) analysis, contribute to a holistic framework. The focus on proactive monitoring, streamlined processes, and network readiness SLAs adds precision to operational timelines. This innovative model offers a strategic blueprint for optimizing data center operations in response to evolving computational requirements.

REFERENCES

[1]. Dongmei Huang, "DATA CENTER INFRASTRUCTURE MANAGEMENT," in Data Center Handbook: Plan, Design, Build, and Operations of a Smart Data Center, Wiley, 2021, pp.627-644, doi: 10.1002/9781119597537.ch33.

[2]. L. Jiadi, H. Yang, L. Huan, Z. Xinli and L. WenJing, "Research on Data Center Operation and Maintenance Management Based on Big Data," 2020 International Conference on Computer Engineering and Application (ICCEA), Guangzhou, China, 2020, pp. 124-127, doi: 10.1109/ICCEA50009.2020.00033.

[3]. M. Wiboonrat, "Energy Management in Data Centers from Design to Operations and Maintenance," 2020 International Conference and Utility Exhibition on Energy, Environment and Climate Change (ICUE), Pattaya, Thailand, 2020, pp. 1-7, doi: 10.1109/ICUE49301.2020.9307075.

[4]. Y. Cui, C. Ingalz, T. Gao and A. Heydari, "Total cost of ownership model for data center technology evaluation," 2017 16th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, USA, 2017, pp. 936-942, doi: 10.1109/ITHERM.2017.7992587.

[5]. P. Martinez-Julia, V. P. Kafle and H. Asaeda, "Using the Total Cost of Ownership to Decide Resource Adjustment in Virtual Networks," 2019 22nd Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN), Paris, France, 2019, pp. 329-335, doi: 10.1109/ICIN.2019.8685844.

[6]. A. Vafamehr, M. E. Khodayar, S. D. Manshadi, I. Ahmad and J. Lin, "A Framework for Expansion Planning of Data Centers in Electricity and Data Networks Under Uncertainty," in IEEE Transactions on Smart Grid, vol. 10, no. 1, pp. 305-316, Jan. 2019, doi: 10.1109/TSG.2017.2738444.