SYMBEEO

PUBLIC

DESIGN STUDY

Colocation move planned: 40 racks into an Equinix hall without a maintenance window

Complete design on current platforms

CAPABILITY	INDUSTRY	SETTING
Data Center	Media & Gaming	Games studio
DOCUMENT	**VERSION**	**DATE**
DC-005	1.3	Oct 10, 2026

WRITTEN BY

Steve Phillips, Principal Data Center Architect

© 2026 Symbeeo. All rights reserved.

Document control

Document	DC-005	Version	1.3
Issued	Oct 10, 2026	Status	Issued by Symbeeo
Classification	Public	Supersedes	1.2, Sep 27, 2026
Capability	Data Center	Industry	Media & Gaming
Author	Steve Phillips	Role	Principal Data Center Architect
Credentials	CCSP, CCIE, AWS Certified Security – Specialty		
Approval	Approved by Edwin Ozer, Founder, Oct 10, 2026. Issued by Symbeeo, Oct 10, 2026. Design study: every figure is vendor-published, a design target or a marked assumption.		
Sign-off	"An engagement under this design is signed by the practice lead before any build proceeds."		

Contents

1 · Scope 3
1c · Public implementations this design follows 4
2 · Requirements register 5
3 · Architecture overview 6
4 · Design by subsystem 6
5 · Failure behavior 8
6 · Limits 8
7 · Lessons learned 8
8 · The author 9

© 2026 Symbeeo. All rights reserved. · Public

CONVENTIONS

Source tags. [S] figure read from the vendor's own current page on Sep 18, 2026 · [S*] figure from the vendor's page summary where the page itself could not be opened · [A] design assumption, to be confirmed by the signing engineer · [T] design target: the record's outcome figure and anything derived from it. No figure comes from any engagement, file or asset. A row with more than one tag is counted once: as a published figure if any tag is [S] or [S*], otherwise as an assumption if any tag is [A], otherwise as a design target.

Table 1. Revision history

Date	Version	Change
Sep 18, 2026	1.0	First issue.
Sep 20, 2026	1.1	Public implementations added; figures reconciled with published sources.
Sep 27, 2026	1.2	Figures and platform details rechecked against published sources.
Oct 10, 2026	1.3	Approval line now reads: approved by the Founder, issued by Symbeeo. Steve Phillips remains the author.

1 · Scope

Who it is for. The head of technology of a studio with live game services that cannot go dark, two server rooms whose cooling and power are past their limits, and a lease that ends before a like-for-like rebuild could.

Typical buyer [A]. A games studio consolidating two aging server rooms. Today: 52 racks of mixed-age hardware across the two rooms, 1,100 physical and virtual hosts, build farm, asset storage, matchmaking and game-session services; peak draw 6 kW per rack. Target: 40 racks in one Equinix IBX hall in the same metro, new compute for 70 % of the estate, 30 % moved as-is.

What is in. The Equinix cabinet and power order, the network fabric in the hall, the new compute, the rack-pair staging and burn-in method, the pairwise cutover without a service window, and the per-rack first-power-on record. **What is out.** Cloud burst for game sessions; the studio's CDN; the decommissioning of the old rooms after the last pair moves.

Assumptions. Equinix is the colocation provider, Dell the server vendor, Cisco the network vendor [A]. A 2 × 100 GbE metro link between the old rooms and the hall exists for the move period [A]. Services are stateless behind load balancers or replicate at the application layer; no service depends on a single rack [A]. The hall is in the US, so NEC 210.20(A) applies and the circuits come from Equinix's AMER (240/415 V) table [A]; the hall's actual offer at order time governs.

Related documents, read Sep 18, 2026.

- Equinix, secure cabinet with kVA-based power — https://docs.equinix.com/colocation/secure-cabinets/colo-secure-cabinets-kva/
- Equinix, colocation power (read again Sep 27, 2026: AMER 240/415 V circuits drawable to 100 %, 120/208 V to 80 %; installed primary circuits up to twice the draw cap) — https://docs.equinix.com/colocation/about/colo-power/
- Dell PowerEdge R7725 technical guide — https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/poweredge-r7725-technical-guide.pdf
- Cisco Nexus 9300-FX3 series switches data sheet — https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-744052.html

© 2026 Symbeeo. All rights reserved. · Public

- Cisco Nexus 9300-GX2 series fixed switches data sheet — https://www.cisco.com/site/us/en/products/collateral/networking/switches/nexus-9000-series-switches/nexus-9300-gx2-series-fixed-switches-data-sheet.html
- Cisco Validated Design, data center networking blueprint for AI/ML applications (read Sep 20, 2026) — https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cvd-for-data-center-networking-blueprint-for-ai.html
- Cisco Nexus 9000 VXLAN BGP EVPN fabrics, design and implementation guide (read Sep 20, 2026) — https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-vxlan-bgp-evpn-design-and-implementation-guide.html
- Equinix SV16 IBX data center page (read Sep 20, 2026) — https://www.equinix.com/data-centers/americas-colocation/united-states-colocation/silicon-valley-data-centers/sv16
- NEC 210.20(A), continuous-load rule — a standard overcurrent device is listed for 80 % of its rating on a load of three hours or more (read Sep 20, 2026) — https://ecalpro.com/blog/nec-80-percent-rule-explained

1c · Public implementations this design follows

Deployment rows are real public deployments, each dated. Where no public deployment matches, the row is Vendor documentation: the vendor's document that this design, Symbeeo's own, is built on.

Table 2. Public implementations this design follows

Implementation	Who published it · date	What it settles for this design	URL
Vendor documentation. Cisco Validated Design fabric: four Nexus 9332D-GX2B spines with 20 of 32 × 400G ports used and 12 left for growth, ten leaves, every leaf to every spine, NX-OS 10.3.2.F	Cisco · Jun 6, 2023	R8 and §4.2: the 9332D-GX2B as a fixed 400G spine with spare ports is a published pattern; this design uses 10 of 32 per spine	https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cvd-for-data-center-networking-blueprint-for-ai.html
Vendor documentation. VXLAN BGP EVPN spine-leaf design guide with the Nexus 93180YC-FX3 as the worked leaf example on 48 host ports; "the more spines a leaf attaches, the lower your oversubscription ratio is"	Cisco · updated Jun 13, 2025	R7 and §4.2: the 93180YC-FX3 as a 48-port leaf dual-attached to two spines	https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-vxlan-bgp-evpn-design-and-implementation-guide.html
Deployment. A live IBX hall listing 1,420 m² of colocation space, 5 kVA minimum cabinet density, N+1 power and cooling, 30 hours or more of generator fuel at full load	Equinix · updated Mar 16, 2026 (page metadata, read Sep 27, 2026)	R3: an 8 kVA cabinet cap sits above a published 5 kVA minimum density in a real hall; N+1 power and cooling are the hall's stated redundancy. The hall lists 120/208 V distribution, so it is not the circuit basis of R4	https://www.equinix.com/data-centers/americas-colocation/united-states-colocation/silicon-valley-data-centers/sv16

© 2026 Symbeeo. All rights reserved. · Public

2 · Requirements register

Table 3. Requirements register

#	Requirement	Value	Source
R1	Racks	40 in the hall, from 52 today	[T]
R2	Cabinet	47U at 2200 mm or 52U at 2422 mm; 600 or 800 mm wide; 1500 kg static load	[S] Equinix secure cabinet (the page's published cabinet table, listed for EMEA; the US hall's cabinet confirmed at order [A])
R3	Power per cabinet	8 kVA contracted draw cap; **design budget 80 % of the cap = 6.4 kVA sustained**; A+B redundant circuit pair, the load on the pair never above what one circuit carries	[A] cap and budget; [S] "redundant pairs of AC circuits", load not above a single circuit
R4	Circuit size	one 20 A, 415 V three-phase pair per cabinet: √3 × 415 × 20 = 14.38 kVA per circuit, drawable to 100 % on AMER 240/415 V circuits; 0.8 × 14.38 = 11.50 kVA if the NEC 210.20(A) 80 % continuous-load rule is applied to the circuit as well	[S] Equinix AMER circuit table and draw rule; [S] NEC 210.20(A); [A] circuit choice
R5	Draw rule	draw cap measured in kVA; server and switch draw in kW converted at a power factor of 0.95	[S] Equinix draw cap; [A] power factor
R6	Server	2U, 2 × AMD EPYC 9005, 800 W to 3200 W Titanium PSUs, 2 × OCP 3.0 to 400 GbE	[S] PowerEdge R7725
R7	Top of rack	48 × 1/10/25G SFP28 + 6 × 40/100G QSFP28, 1RU, 325 W typical	[S] Nexus 93180YC-FX3
R8	Spine	32 × 400G QSFP-DD, 25.6 Tbps, 1RU, 638 W typical; 4 × 100G breakout per port	[S] Nexus 9332D-GX2B
R9	Cutover unit	one rack pair at a time, 20 pairs	[T]
R10	Service window	none; every service stays up through its pair's move	[T]
R11	Burn-in	48 h powered and loaded per rack before any service moves onto it	[A]
R12	Power-on record	first-time power-on pass or fail recorded per rack, 40 records	[T]
R13	Move pace	2 pairs per week, 10 weeks	[A]

Thirteen rows; seven carry a vendor figure, two are design assumptions, four are design targets.

© 2026 Symbeeo. All rights reserved. · Public

3 · Architecture overview

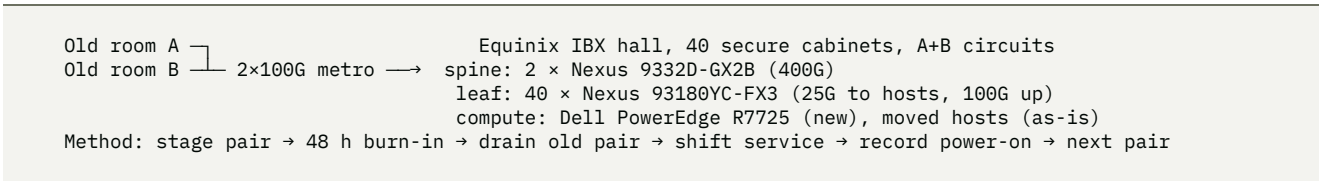


Figure 1. Colocation move by rack pair

Twenty rack pairs, each pair a self-contained unit of compute, its own leaf switches and its own circuit pairs. A pair goes live only after its burn-in, and the old pair it replaces is drained by the load balancer before anything is unplugged. No service is ever on one pair only.

Table 4. Architecture overview

Role	Product (exact model)	Qty	Capacity used	Why this unit
Cabinet	Equinix Secure Cabinet, 47U, 800 × 1200 mm	40	24U of servers in 47U per rack [A]; 12 servers per cabinet, not 16 (6.4 kVA design budget, 80 % of the 8 kVA cap, §4.1)	1500 kg static load, kVA-based power [S]
Circuits	Equinix three-phase AC pair per cabinet	40 pairs	8 kVA cap on a 20 A/415 V three-phase pair, 14.38 kVA per circuit [A][S]	redundant A+B pairs [S]
Spine	Cisco Nexus 9332D-GX2B	2	40 leaf uplinks × 100G = 10 of 32 ports at 4 × 100G breakout [S][A]	25.6 Tbps in 1RU, 638 W [S]
Leaf	Cisco Nexus 93180YC-FX3	40 (pair per rack pair, one per rack)	24 × 25G host ports of 48 in a new-compute pair: 24 hosts, each dual-homed [A]	48 × 25G + 6 × 100G, 325 W [S]
New compute	Dell PowerEdge R7725	240 (12 per rack, 20 racks)	70 % of the estate: 770 of 1,100 hosts at 3.2 per server [A] (§4.3)	2U, 2 × EPYC 9005, up to 192 cores each [S]
Moved compute	existing hosts, as-is	20 racks	30 % of the estate [A]	not rebuilt before the lease ends [A]

4 · Design by subsystem

4.1 Cabinets and power.

Secure cabinets at 2200 mm are 47U; 2422 mm are 52U; widths 600 or 800 mm, depths 1000 or 1200 mm, 1500 kg static load [S] (Equinix publishes this table for EMEA; the US hall's cabinet is confirmed at order [A]). The design orders 800 × 1200 mm at 47U for cable room [A]. Power is delivered "through redundant pairs of AC circuits"; the usable power of a pair is set by the primary circuit, and the load spread across the pair must not exceed the capacity of a single circuit [S].

One basis, in kVA. The contract sets an 8 kVA draw cap per cabinet [A], and the design holds sustained draw at 80 % of it, 6.4 kVA, so the full-load burn-in and later growth have room [A]. Server and switch draw in kW is converted at a power factor of 0.95 [A]. The circuit is one 20 A, 415 V three-phase pair per cabinet from Equinix's AMER (240/415 V) table [A]: √3 × 415 V × 20 A = 14.38 kVA per circuit, drawable to 100 % on AMER 240/415 V circuits [S]. Applying NEC 210.20(A) as well, which rates a standard overcurrent device for 80 % of its rating on a continuous load of three hours or

© 2026 Symbeeo. All rights reserved. · Public

more [S], one feed presents 0.8 × 14.38 = 11.50 kVA. Both circuit figures are above the 8 kVA cap, so the cap, not the circuit, sets the budget, and either feed alone carries the whole cabinet. The installed primary capacity, 14.38 kVA, is within Equinix's limit of twice the draw cap, 16 kVA [S].

Servers per cabinet. Load per new-compute cabinet = n × 450 W per R7725 [A] + one leaf at 325 W typical [S] (each rack pair's two leaves sit one per rack, §4.2). Twelve servers: 12 × 0.45 + 0.325 = 5.725 kW ÷ 0.95 = **6.03 kVA, 75 % of the cap and inside the 6.4 kVA budget**; 42 % of one 14.38 kVA feed and 52 % of the 11.50 kVA NEC figure. Thirteen: 6.175 kW ÷ 0.95 = 6.50 kVA, over the budget. Sixteen: 7.525 kW ÷ 0.95 = 7.92 kVA, 99 % of the cap, rejected because a full-load burn-in would take it over [A]. So the cabinet takes **twelve** R7725, and the 450 W, an assumed average draw, is replaced by the measured draw from the §4.4 burn-in before the next cabinet pair is ordered. Moved-as-is racks sit under the same cap: the old rooms' 6 kW peak [A] plus a leaf is 6.325 kW ÷ 0.95 = 6.66 kVA, 83 % of the cap, inside the cap but over the budget at peak, so §4.6 watches those racks first. Rejected: 6 kW racks at the old rooms' density — 52 racks do not fit R1.

4.2 Fabric.

Two Nexus 9332D-GX2B spines, 32 × 400G QSFP-DD each, 25.6 Tbps, 1RU [S], with 4 × 100G breakout per port [S]; 40 leaves at 2 × 100G each consume 80 breakout lanes, 10 ports of 32 on each spine [A]. Leaf pairs per rack pair: Nexus 93180YC-FX3, 48 × 25G SFP28 and 6 × 100G QSFP28 [S]; each host dual-homed at 25G to both leaves [A]. Rejected: 9364D-GX2A 64-port spine — 10 of 64 ports used, 2RU and 1380 W typical [S], over-provisioned for 40 racks (R1).

4.3 New compute.

PowerEdge R7725: 2U, two EPYC 9005 processors of up to 192 cores, 24 DIMM slots, two OCP 3.0 NIC slots to 400 GbE, PSUs from 800 W to 3200 W Titanium [S]. Specified build: 2 × 64 cores, 768 GB RAM, dual 25 GbE OCP, 1400 W redundant PSUs [A]; twelve per rack at 2U each is 24U, leaving room for the leaf, PDU and patching in 47U [A]. Coverage: 70 % of the 1,100 hosts is 770 hosts on 240 servers, an average of 3.2 hosts per server, which gives each migrated host about 40 cores and 240 GB on average (128 ÷ 3.2 and 768 ÷ 3.2) [A]. The inventory confirms that density before the first order; where it cannot be reached, the moved-as-is share rises, not the rack count (R1) [A].

4.4 Rack-pair staging and burn-in.

Each pair is racked, cabled and powered in the hall before any service touches it. Burn-in: 48 hours at full CPU and memory load with fans, PSU redundancy and both circuits exercised by pulling one side for one hour [A]. The first power-on is recorded per rack: date, engineer, pass or fail on the first attempt, and the fault if any [T]. Forty records are the rate; nobody types the rate.

4.5 Pairwise cutover without a window.

Per pair: services on the old pair are added to the new pair's hosts across the 2 × 100G metro link [A]; the load balancer drains the old pair over 30 minutes; stateful services replicate and switch primaries one at a time [A]; the old pair is powered off only after 24 hours with no traffic [A]. Moved-as-is racks follow the same rule with their twin: the twin carries the service while the rack is in the truck [A]. Rejected: a weekend big-bang move — fails R10.

4.6 Monitoring and lifecycle.

Per-cabinet kVA read from the Equinix portal against the 8 kVA cap [A]; **the alarm threshold is the 6.4 kVA design budget, 80 % of the cap [A]** (the NEC continuous-load rule applies to the circuit, §4.1, not to the contract cap), and a rack over it for a week gets its next host placed elsewhere [A]; spine and leaf telemetry to the studio's NMS.

© 2026 Symbeeo. All rights reserved. · Public

5 · Failure behavior

Table 5. Failure behavior

Failure	Behavior	Source
One circuit of a pair lost	rack runs on the other circuit (6.03 kVA on a 14.38 kVA circuit); PSUs are 1+1	[S] redundant pairs; [S] R7725 redundant PSU options
One leaf down	hosts stay up on the second leaf at half bandwidth	[A] dual-homed hosts
One spine down	fabric at half capacity; no service loss	[A]
Rack fails burn-in	pair does not go live; fault logged in the power-on record; cutover date slips a week	[A] §4.4; [T] R12
Old pair fails during drain	new pair already carries the service; drain completes early	[A] §4.5
Metro link down during a move	in-flight pair pauses; services on both sides keep running	[A]

6 · Limits

- The 8 kVA cap, the 80 % budget, the 0.95 power factor and 450 W per server are assumptions; the burn-in load test sets the real per-rack draw before the next pair is ordered. **The 6.4 kVA budget (80 % of the contract cap) is what sets twelve servers per cabinet rather than sixteen; the circuit, 14.38 kVA per feed or 11.50 kVA under the NEC 80 % rule, is not the limit.**
- Coverage of 70 % of the estate by 240 servers rests on an average of 3.2 hosts per server [A]; the inventory confirms it before the first order.
- The pairwise method needs every service to exist on two racks; a service found on one rack is fixed before its pair moves, which can slow the 10-week pace.
- Moved-as-is racks are not burned in like new ones; their power-on record notes the difference.
- Circuit values come from the published Equinix table; the hall's actual offer at order time governs.
- 30 % of the estate arrives without refresh and carries its own end-of-life dates into the hall.

7 · Lessons learned

None yet — this is a design study; populated after the first build against it.

© 2026 Symbeeo. All rights reserved. · Public

8 · The author

WRITTEN BY

Steve Phillips

Principal Data Center Architect

CCSP, CCIE, AWS Certified Security – Specialty

Signed off by Edwin Ozer, Founder.

© 2026 Symbeeo. All rights reserved. · Public

SYMBEEO

Each vendor's tool shows you their option. We show you all of them against one requirement.

Compared. Then Delivered.

hello@symbeeo.com · symbeeo.com

© 2026 Symbeeo. All rights reserved.

Product and company names and logos are their owners' trademarks, shown only to identify the lines Symbeeo sources. Naming a vendor or showing its logo does not mean the vendor sponsors or approves Symbeeo's comparisons or any recommendation Symbeeo makes.

Figures are as published by their sources on the dates given in this document; figures marked [A] are assumptions.

DC-005 · Version 1.3 · Oct 10, 2026