MICROSOFT AZURE LOCAL

Network Design Decision Framework

Eleven decisions from connectivity to software-defined networking — across the Hyperconverged and Disaggregated architecture paths.

Cristian Edwards Sabathe
Principal Product Manager – Azure Local Networking Lead

MICROSOFT AZURE LOCAL

Network Design Decision Framework

Eleven decisions from connectivity to software-defined networking — the Hyperconverged and Disaggregated paths.

1 Connectivity mode

APPLIES TO BOTH ARCHITECTURES

- Connected
- Disconnected

2 Architecture

- Hyperconverged Path
- Disaggregated Path

3 Cluster topology

HYPERCONVERGED PATH

- Single Rack
- Rack-Aware (2 zones)

DISAGGREGATED PATH

- Single Rack
- Multi-Rack (1–8 racks)

4 Cluster size

HYPERCONVERGED PATH

- 1–16 nodes

DISAGGREGATED PATH

- Up to 64 nodes (1–8 racks)

5 Storage connectivity

HYPERCONVERGED PATH

- Switchless (1–4)
- Switched (5+)
- Hybrid Storage (S2S + SAN)

DISAGGREGATED PATH

- Fibre Channel SAN
- iSCSI SAN

6 Network adapter ports

HYPERCONVERGED PATH

- 2 / 4 / 6 / 8 ports
- VLAN IDs and Speed

DISAGGREGATED PATH

- 6 ports
- VLAN IDs and Speed

7 Network traffic intents

HYPERCONVERGED PATH

- Group all
- Mgmt + Compute
- Compute + Storage
- Custom

DISAGGREGATED PATH

- Mgmt + Compute
- Cluster
- iSCSI
- Backup

8 Mgmt IPs & Infra network

APPLIES TO BOTH ARCHITECTURES

- Static / DHCP
- Default / Custom VLAN
- Infra CIDR + IP Pool + GW
- DNS
- ARB reserved subnets

9 Backup network

HYPERCONVERGED PATH

- No backup network
- Dedicated: 2nd compute intent (Custom)
- +2 NIC ports / node

DISAGGREGATED PATH

- Fibre Channel: 6-port backup
- iSCSI 6-port: optional backup
- iSCSI 4-port: not supported

10 Outbound connectivity

APPLIES TO BOTH ARCHITECTURES

when Azure Local is in Connected mode

OUTBOUND MECHANISM

- Arc gateway
- Proxy
- Firewall

OUTBOUND PATH

- via Public internet
- via ExpressRoute
- via S2S VPN

11 Software-Defined Networking (SDN)

On HCI, SDN requires a Group-all or Mgmt+Compute intent (Step 7)

HYPERCONVERGED PATH

- SDN enabled by Azure Arc
- Logical networks (LNETs) + NSGs

DISAGGREGATED PATH

- Fabric-based SDN (VXLAN EVPN)
- Leaf-spine overlay (VRF)

How to read this framework

- Hyperconverged path
- Disaggregated path
- Applies to both architectures

Steps 1, 8 and 10 apply to both architectures.
Steps 2–7, 9 and 11 follow the path you choose.

Network Design decision 1: Connectivity mode

The connectivity mode applies to both the hyperconverged and disaggregated architectures

Connected

The standard cloud deployment: nodes reach Azure over the internet for registration, billing, and lifecycle. Your job is to guarantee a reliable outbound path — you design it in Decision 10.

Disconnected

Air-gapped mode for sovereign or regulated sites that can't reach public Azure. A local Autonomous Cloud replaces the public Azure endpoints, deployed as a dedicated management cluster plus workload clusters.

Default to Connected; only choose Disconnected when a regulatory/air-gap mandate requires it

Network Design decision 1: Connectivity mode

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	Connected deployments need outbound access to Azure for Arc registration, billing, and lifecycle management. Plan the outbound topology in Decision 10.	**Both**
2	Disconnected deployments use Azure Local disconnected operations with a local Autonomous Cloud endpoint and don't require public Azure endpoints.	**Both**
3	Disconnected deployments require a dedicated three-node management cluster plus one or more workload clusters.	**Both**
4	The connectivity mode applies to both the hyperconverged and disaggregated architectures you choose in Decision 2.	**Both**

Default to Connected; only choose Disconnected when a regulatory/air-gap mandate requires it

Network Design decision 2: Architecture

Hyperconverged vs Disaggregated

Hyperconverged

Storage lives inside the cluster on local drives, pooled with Storage Spaces Direct over RDMA. Compute and storage scale together – the right fit for most deployments.

Disaggregated

Storage lives outside the cluster on an external Fiber Channel or IP based SAN. Compute and storage scale independently, to much larger clusters.

The architecture you choose determines the storage connectivity, network adapter ports, and network traffic intents available in later steps.

Network Design decision 2: Architecture

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	Hyperconverged (HCI) uses Storage Spaces Direct with local drives. Compute and storage scale together, up to 16 nodes.	HCI
2	Disaggregated (DA) uses an external FC or IP based SAN with no S2D. Compute and storage scale independently across 1–8 racks, up to 64 nodes.	DA
3	The architecture you choose determines the storage connectivity, adapter ports, and traffic intents available in later steps.	Both
4	Hybrid storage is HCI-only: S2D with an external SAN side by side, attached as a day-2 operation, not during first deployment.	HCI
5	HCI supports a rack-aware variant across two zones; DA scales across multiple racks. You define the layout in Decision 3.	Both

The architecture you choose determines the storage connectivity, network adapter ports, and network traffic intents available in later steps.

Network Design decision 3: Cluster topology

Hyperconverged single rack vs Hyperconverged rack aware | Disaggregated single rack vs Disaggregated multiple racks

Hyperconverged

Single Rack | Rack Aware

Grows within the storage cluster: deploy a single rack, or go rack-aware across two rooms or zones for rack-level resiliency.

Disaggregated

Single Rack | Multi rack (1-8 racks)

Grows by adding racks: a single rack for smaller clusters, or a multi-rack leaf-spine fabric to scale out.

Choose Disaggregated only if you already run a SAN or must scale compute and storage independently. Plan hybrid SAN adapters up front — it's a day-2 attach.

Network Design decision 3: Cluster topology

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	A single rack with a pair of ToR switches supports up to 16 nodes for both architectures.	Both
2	HCI single-rack clusters run all host traffic over the same MLAG-configured ToR pair. Don't use a separate storage-only switch pair.	HCI
3	HCI rack-aware spans two zones, an even number of nodes up to 8, and needs <1 ms latency. Not supported with external SAN (hybrid).	HCI
4	For HCI rack-aware, choose one of four uplink options; dedicated storage links keep RDMA at the ToR layer with the lowest latency.	HCI
5	DA clusters span 1–8 racks on a leaf-spine (Clos) fabric, with up to 16 nodes per rack and 64 nodes per cluster.	DA
6	For more than 16 DA nodes, use a leaf-spine fabric across multiple racks. Only fabric-based SDN is supported.	DA

Choose Disaggregated only if you already run a SAN or must scale compute and storage independently. Plan hybrid SAN adapters up front — it's a day-2 attach.

Network Design decision 4: Cluster size

Up to 16 nodes in a single rack | Up to 64 nodes in up to 8 racks.

Hyperconverged

1 to 16 nodes

1 to 16 nodes. The node count is what drives your storage choice — only small clusters can run switchless.

Disaggregated

Up to 64 nodes distributed in up to 8 racks

Up to 64 nodes across racks. Compute and storage are sized independently because capacity comes from the external SAN.

Any chance of growing past 4 nodes? Deploy switched storage now — switchless scale-out means manual re-cabling Microsoft doesn't validate.

Network Design decision 4: Cluster size

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	HCI clusters with more than 4 nodes that use S2D require a physical switch for the storage network traffic.	HCI
2	If you intend to scale out the cluster using the orchestrator, you must use a physical switch for storage traffic.	HCI
3	DA scales compute and storage independently (1–8 racks, up to 64 nodes). Storage capacity is independent of node count because it comes from the external SAN.	DA

Any chance of growing past 4 nodes? Deploy switched storage now — switchless scale-out means manual re-cabling Microsoft doesn't validate.

Network Design decision 5: Storage connectivity

Hyperconverged Switched, Switchless or hybrid | Disaggregated SAN based FC storage or iSCSI storage

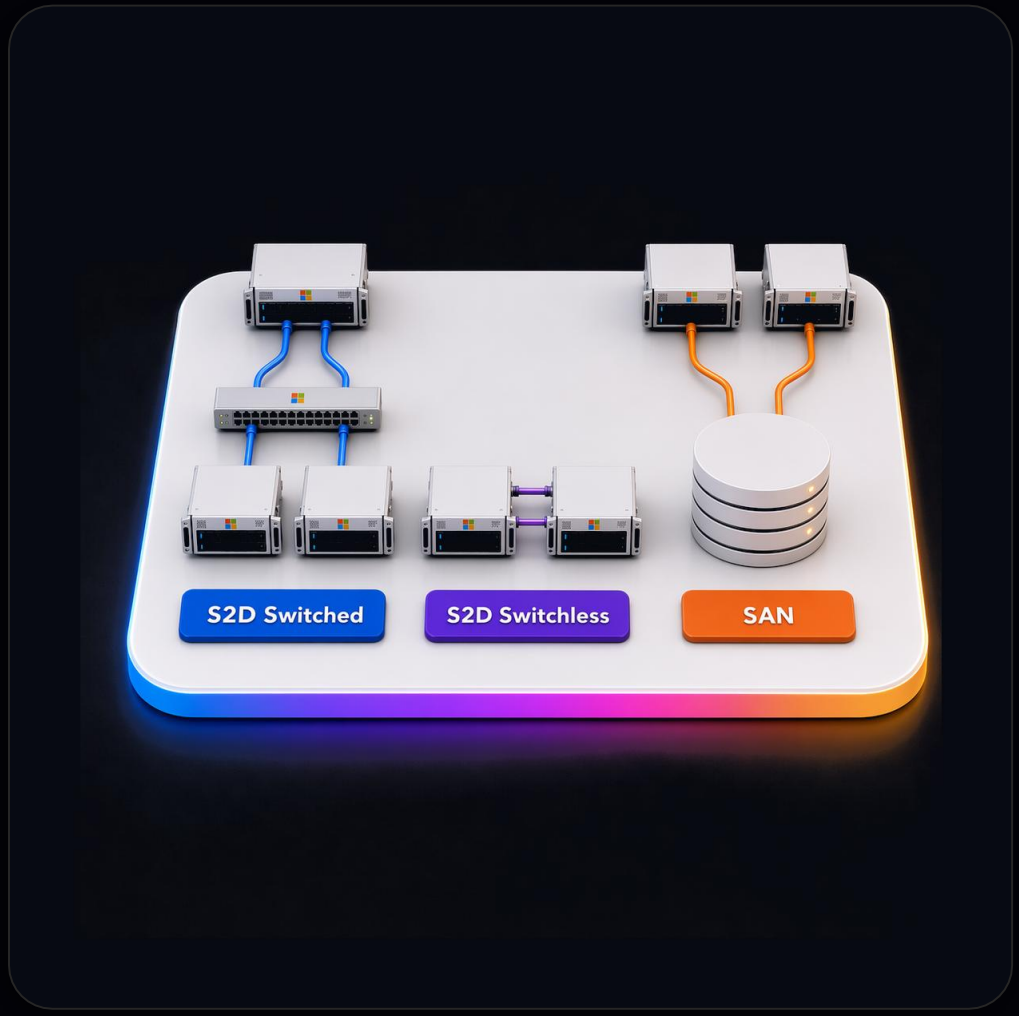

Hyperconverged

Switchless | Switched | Hybrid (S2D & SAN)

Internal Storage Spaces Direct – switched, or switchless for small clusters – with an optional hybrid SAN attached on day 2.

Disaggregated

FC SAN | IP based SAN

An external SAN reached over Fiber Channel or IP address. There's no Storage Spaces Direct and no RDMA storage intent.

Keep the storage intent on the same ToR pair as management/compute to avoid split-brain; pre-define storage IPs in ARM for 3–4-node switchless.

Network Design decision 5: Storage connectivity

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	Switchless via the Azure portal is supported for 1–2 node clusters. 3- and 4-node switchless clusters require Resource Manager templates.	HCI
2	Scale-out isn't supported with switchless deployments. Any change to node count after deployment requires manual reconfiguration.	HCI
3	A network switch for storage works with 1–16 nodes, and a single intent can carry all traffic types.	HCI
4	On HCI, the storage intent must share the same ToR pair as management and compute. A storage-only switch pair can cause cluster split-brain.	HCI
5	DA uses FC or IP based to an external array. There's no S2D, and compute and storage scale independently up to 64 nodes.	DA
6	iSCSI uses 6-port dedicated-path (optional backup) pattern.	DA
7	iSCSI requires per-target /32 host static routes and MPIO so storage stays on the storage adapter. iSCSI runs over TCP, so PFC/lossless aren't required.	DA

Keep the storage intent on the same ToR pair as management/compute to avoid split-brain; pre-define storage IPs in ARM for 3–4-node switchless.

Network Design decision 6: Network Adapter ports

Hyperconverged 2, 4, 6 or 8 network ports | Disaggregated 4 or 6 network ports

Hyperconverged

2,4,6 and 8 network ports

RDMA

The port count follows how you group intents – 2, 4, 6, or 8 ports per node.

Disaggregated

6 ports

At least 10 Gbps

Host traffic sits on the onboard adapter, with the cluster and SAN networks on separate standalone ports – 6 per node.

Keep the storage intent on the same ToR pair as management/compute to avoid split-brain; pre-define storage IPs in ARM for 3–4-node switchless.

Network Design decision 6: Network Adapter ports

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	Use the default Network ATC configurations as much as possible.	Both
2	Physical switches must be configured to match the network adapter configuration.	Both
3	Work with your OEM to determine which network adapters are supported and qualified for each intent on Azure Local.	Both
4	With defaults, Network ATC auto-configures storage adapter IPs and VLANs (Storage Auto IP). Some cases require ARM-declared storage IPs.	HCI
5	At least 10 Gbps interfaces are required to support RDMA storage traffic. 25 GbE or higher is recommended.	Both
6	At least 10 Gbps interfaces are required for disaggregated deployments with iSCSI	DA
7	It is recommended to use access mode for the cluster network VLANs on the physical switch ports.	DA
8	It is recommended to use access mode for the iSCSI network VLANs on the physical switch ports.	DA
9	Disaggregated storage and cluster traffic runs over TCP, so PFC and host-side ETS/DCB aren't required. Host ETS only affects the host-to-leaf link and isn't honored across the routed leaf-spine fabric; engineer fabric capacity to protect cross-rack traffic.	DA

Keep the storage intent on the same ToR pair as management/compute to avoid split-brain; pre-define storage IPs in ARM for 3–4-node switchless.

Network Design decision 7: Network traffic intents

Hyperconverged all intents managed by Network ATC | Disaggregated only Mgmt + compute intent managed by Network ATC

Hyperconverged

Group all | Mgmt & compute | custom

Network ATC manages all host traffic, with flexible intent groupings for management, compute, and storage.

Disaggregated

Mgmt & compute | Cluster | iSCSI | Backup

Network ATC manages only management and compute; the cluster and iSCSI networks run outside ATC.

Keep Network ATC defaults; go Custom only for a genuine second compute/backup intent, and confirm the grouping is supported for your storage type.

Network Design decision 7: Network traffic intents

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	Switchless HCI clusters require at least two intents (management and compute, plus storage).	HCI
2	The Group-all and Compute+Storage intents require a physical switch for storage and aren't available for switchless clusters.	HCI
3	Use at least two network adapter ports per intent to ensure high availability.	Both
4	DA deployments use a management and compute intent.	DA
5	For Cluster Networks, the paths are standalone ports outside Network ATC	DA
6	For iSCSI, the iSCSI paths are standalone ports outside Network ATC	DA

Keep Network ATC defaults; go Custom only for a genuine second compute/backup intent, and confirm the grouping is supported for your storage type.

Network Design decision 8: Management network

Applies to both Hyperconverged and Disaggregated architectures

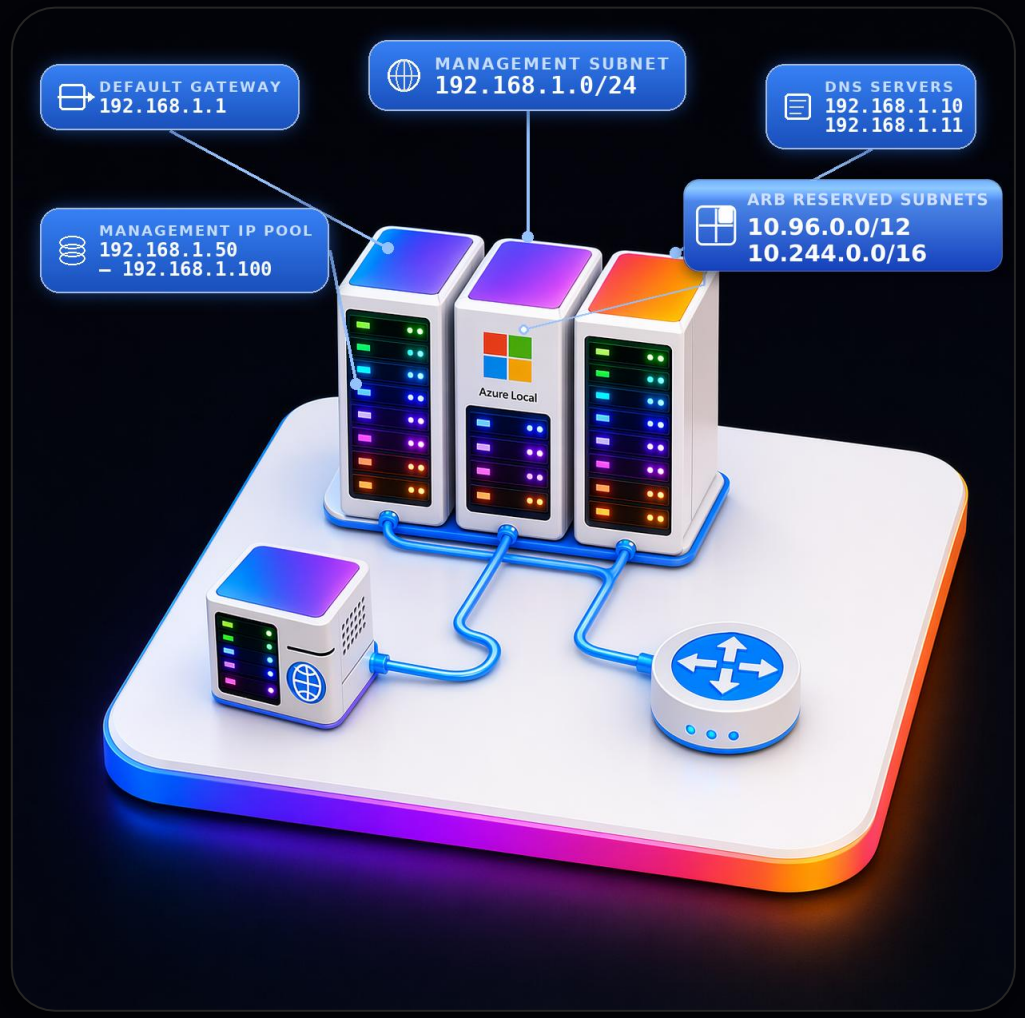

The Mgmt IP Pool

Static / DHCP | Mgmt subnet CIDR | At Least 6 IPs | DNS | Default / Custom VLAN | ARB reserved subnets

- Identical for both paths — and locked in before you deploy.
- Reserve a stable management IP pool and set the management VLAN before Arc registration; neither can change afterward.
- Use external DNS that resolves Active Directory and the internet, and steer clear of the reserved Arc Resource Bridge ranges.

Lock the management VLAN and IP pool before Arc registration — neither can change afterward. Verify nothing overlaps 10.96.0.0/12 or 10.244.0.0/16.

Network Design decision 8: Management network

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	Reserve a management IP pool of at least six consecutive, available IPs. This range can't be changed after deployment.	**Both**
2	Keep the pool on the node subnet but exclude node IPs and default gateway. Use DHCP exclusions and reservations when nodes use dynamic IPs.	**Both**
3	The default gateway must allow ICMP and provide outbound internet connectivity.	**Both**
4	Set the management VLAN ID on the physical NICs before Arc registration. It can't change after deployment and is inherited by the ARB VM.	**Both**
5	Infrastructure services (ARB, Network Controller) always use static IPs from the management pool, even when nodes use DHCP.	**Both**
6	DNS servers must resolve Active Directory and the internet. Public resolvers such as 8.8.8.8 aren't supported and can't change after deployment.	**Both**
7	At least one configured DNS server must run outside the cluster, so name resolution works during boot, shutdown, and recovery.	**Both**
8	Never overlap the reserved Arc Resource Bridge ranges 10.96.0.0/12 and 10.244.0.0/16.	**Both**

Lock the management VLAN and IP pool before Arc registration — neither can change afterward. Verify nothing overlaps 10.96.0.0/12 or 10.244.0.0/16.

Network Design decision 9: Backup network

Hyperconverged dedicated compute intent | Disaggregated dedicated or shared network ports

Hyperconverged

No backup network

Dedicated compute intent for backups

Optional. Add a dedicated backup network as a second compute intent – effectively the 8-port layout.

Disaggregated

6 network ports required for guest backups

Determined by the port layout:

- 6-port Fiber Channel can use dedicated ports with an additional compute intent and a vNIC on the host.
- 6-port iSCSI layouts use the Management and Compute intent. Create a vNIC manually for Guest Backups network on the hosts.

Need a dedicated backup path? Budget the extra 2 ports/node (HCI 8-port or DA 6-port) at purchase — the iSCSI 4-port layout can never add one.

Network Design decision 9: Backup network

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	A dedicated backup network requires two extra network adapter ports per node. Use at least two ports for high availability.	Both
2	On HCI, add the backup network as a second compute intent using the Custom configuration intent model.	HCI
3	On DA, use the 6-port Fiber Channel or 6-port iSCSI dedicated-path layout to add a guest backup network.	DA
4	The iSCSI 6-port pattern manually adds a vNIC on the Management & Compute intent for Guest Backup network on the host	DA

Need a dedicated backup path? Budget the extra 2 ports/node (HCI 8-port or DA 6-port) at purchase — the iSCSI 4-port layout can never add one.

Network Design decision 10: Outbound connectivity

With or without Arc gateway | Public vs Private Path | Applies to Hyperconverged and Disaggregated architectures

Outbound Mechanism

Arc gateway | Proxy | Firewall

How traffic reaches Azure: an Arc gateway and/or an enterprise proxy, or a firewall alone. Proxy + Arc gateway is recommended and shrinks the firewall list.

Outbound path

Public internet | ExpressRoute /S2S VPN

Where it flows: the public internet, ExpressRoute, or a site-to-site VPN. A fully private path over ExpressRoute or S2S VPN is not available yet.

Use Proxy + Arc gateway for new production (cuts the firewall list to < 30) and always deploy the initial registration over a public path.

Network Design decision 10: Outbound connectivity

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	Five outbound topologies exist: direct, enterprise proxy, Arc gateway, proxy + Arc gateway (recommended), and private path.	**Both**
2	Deploying infrastructure registration over a fully private path isn't supported yet. Use a public-path topology for deployment first.	**Both**
3	The Arc gateway reduces required public endpoints from 100+ to under 30 (about 23 FQDNs remain on the allowlist).	**Both**
4	Only non-authenticated proxies are supported. PAC files and .local proxy domains aren't supported.	**Both**
5	Do not configure proxy manually. Since Azure Local 2506, the Arc registration script configures the proxy automatically across WinINET, WinHTTP, and environment variables.	**Both**
6	Plan the proxy bypass list before deployment (machine/cluster IPs and names, infra subnet, private endpoints). It can't change afterward.	**Both**
7	SSL inspection isn't supported for Arc gateway endpoints, and firewall configuration must be done before registering nodes.	**Both**
8	When AKS is present, its subnet needs line of sight to the infra subnet on TCP 22, 6443, 40343, 55000, and 65000.	**Both**

Use Proxy + Arc gateway for new production (cuts the firewall list to < 30) and always deploy the initial registration over a public path.

Network Design decision 11: SDN

Hyperconverged SDN enabled by Azure Arc and Microsoft Network Controller | Disaggregated fabric based SDN

Hyperconverged

Microsoft SDN managed by Arc

LNETs and NSGs

Software-defined networking is managed by Azure Arc and scoped to logical networks (LNETs) and NSGs only.

Disaggregated

Fabric based SDN (VXLAN EPVN)

Leaf and Spine overlay

Software-defined networking is delivered by the external leaf-spine fabric (VXLAN EVPN), not by Azure Arc.

Decide SDN up front — HCI Arc-managed SDN needs a Group-all or Mgmt+Compute intent (Decision 7). Need VNETs/SLB/GRE? Use Windows Server SDN instead.

Network Design decision 11: SDN

Design considerations — validate these before you deploy

#	Design consideration	Applies to
1	SDN is optional and layered on top of the host network after deployment. Decide whether workloads need LNETs or NSGs.	Both
2	Arc-managed SDN on Azure Local supports LNETs and NSGs only. VNETs, SLB and RAS/GRE gateways aren't supported.	HCI
3	Unmanaged VMs on Azure Local whose SDN is managed by on-premises tools can't be hydrated as Azure Local VMs.	HCI
4	LNETs are backed by VLANs. Configure and trunk the corresponding VLANs on the physical switches for each logical network.	HCI
5	If workloads require VNETs, SLB, GRE, or the full Microsoft SDN gateway stack, use a Windows Server SDN deployment instead.	HCI
6	DA uses external fabric-based SDN (VXLAN EVPN) on the leaf-spine fabric. The Microsoft SDN Network Controller isn't used.	DA

Decide SDN up front — HCI Arc-managed SDN needs a Group-all or Mgmt+Compute intent (Decision 7). Need VNETs/SLB/GRE? Use Windows Server SDN instead.

Realize the framework with ODIN

aka.ms/ODIN — the Optimal Deployment and Infrastructure Navigator

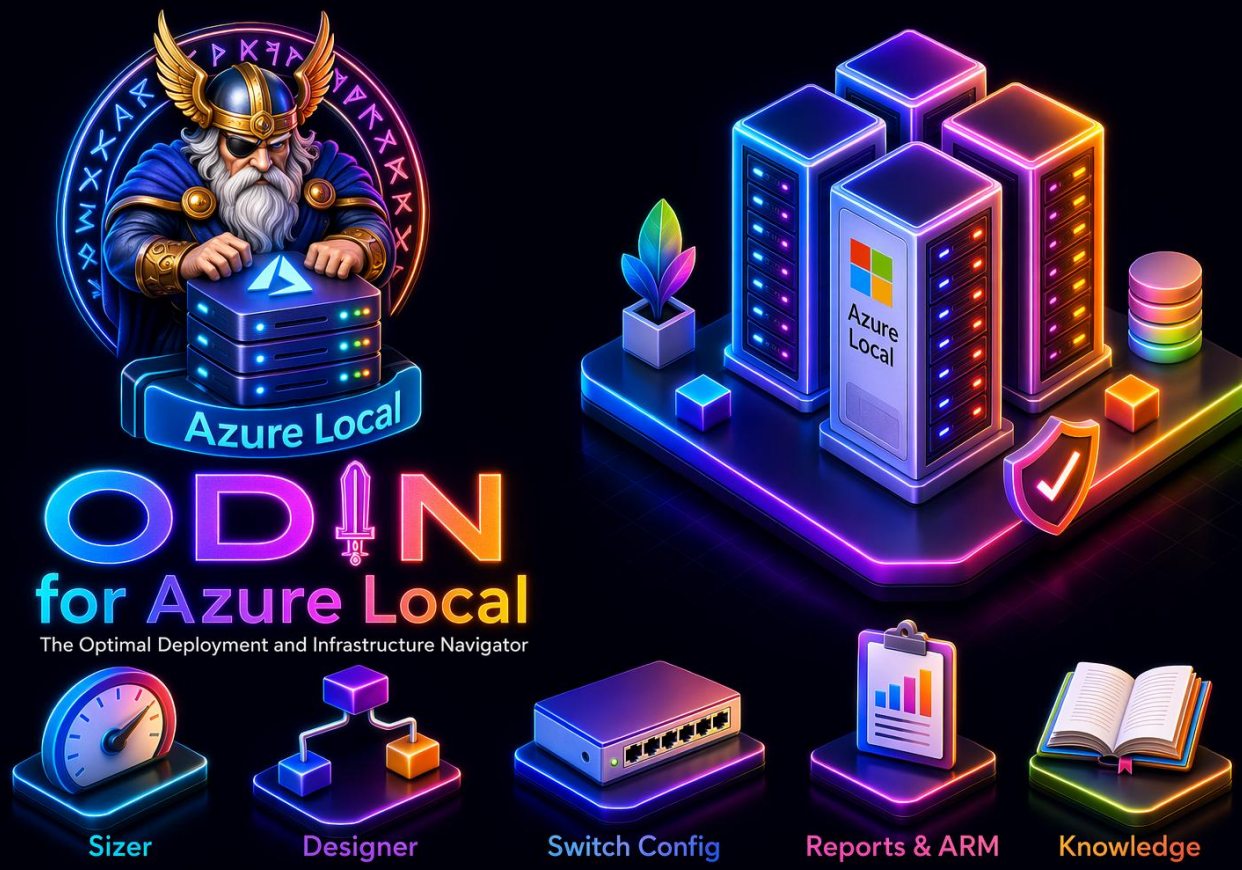

ODIN turns this 11-decision framework into a guided, validated Azure Local design. It's a browser-based decision-tree tool that walks you through every decision and produces validated architecture and network configuration you can deploy.

- **Sizer** — Convert your VM / AKS / AVD workload profile into a supported node count (Decision 4).
- **Designer** — Walk the decision tree for Connected, Disconnected, Multi-Rack, or M365 Local and design a validated instance.
- **Switch Config** — Generate and QoS-audit top-of-rack switch configuration for your intents and storage fabric (Decisions 5–7).
- **Reports & ARM** — Export the validated design and deploy it with Azure Resource Manager templates (Decisions 8 & 10).
- **Knowledge** — Jump to validated architecture and network guidance for every decision in the framework.

Experimental and provided as-is (not Microsoft-supported). Ideal for skilling up and validating a design before you deploy — everything runs locally in your browser.

Recommended ODIN workflow

From business purpose to a deployment-ready design — in three steps

STEP 1

Shape the requirements

Knowledge ▸ Sovereign Private Clouds reference architectures

Map your business purpose to a starter reference architecture and export a 10,000-foot design deck.

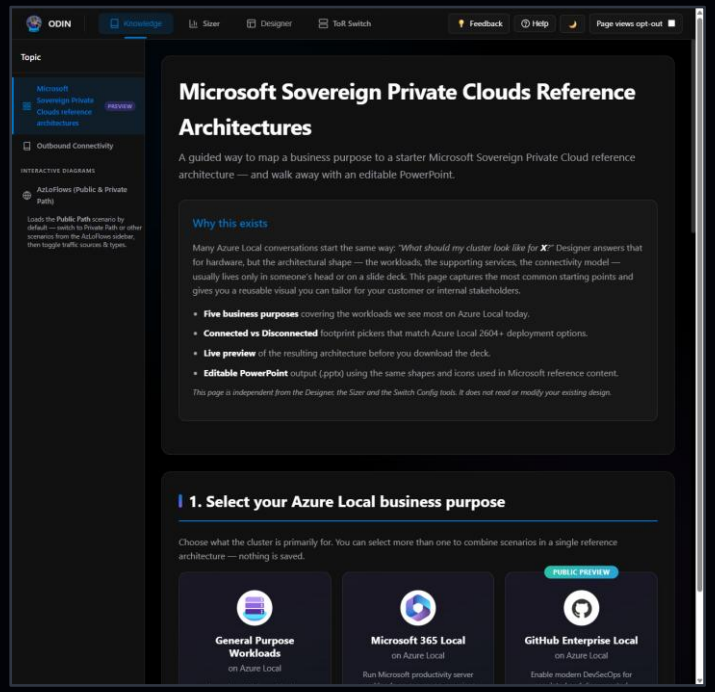

STEP 2

Size the hardware

Sizer

Turn your VM, AKS, and AVD workloads into the node count and hardware you need — then hand off to the Designer.

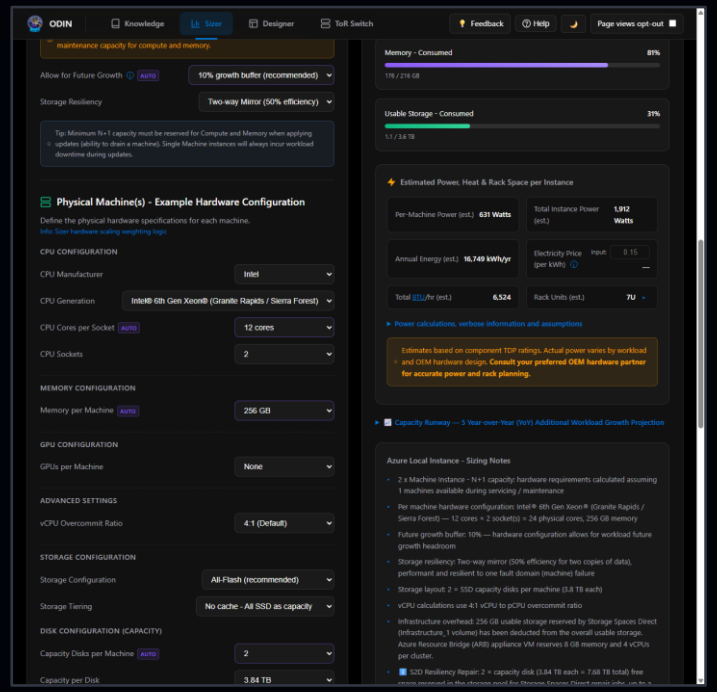

STEP 3

Finalize the design

Designer

Complete every network design decision and generate the documentation and ARM templates, ready to deploy.

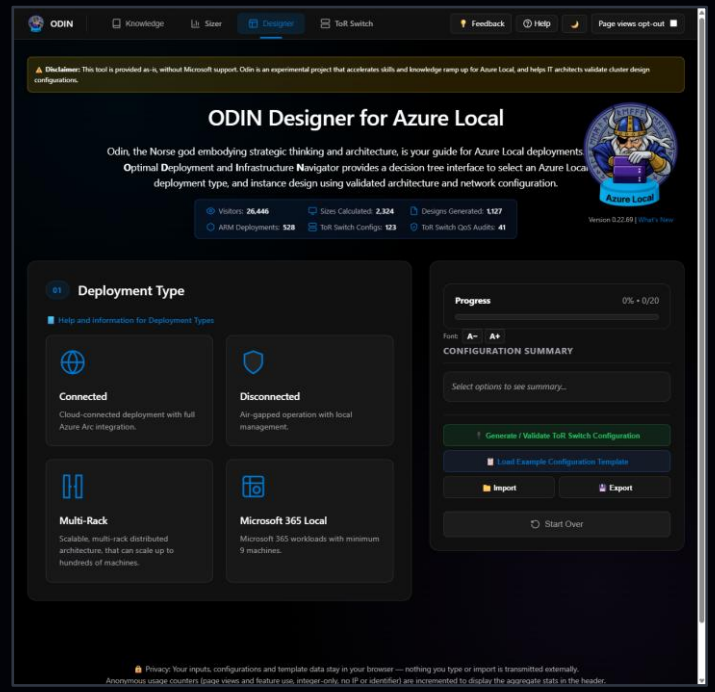

Recommended flow: start with the business purpose, size the hardware, then finalize in the Designer — each step feeds the next.

What ODIN generates for you

Example outputs from the workflow — ready to share and deploy

10,000-ft wizard ▸ reference architecture (editable PPTX)

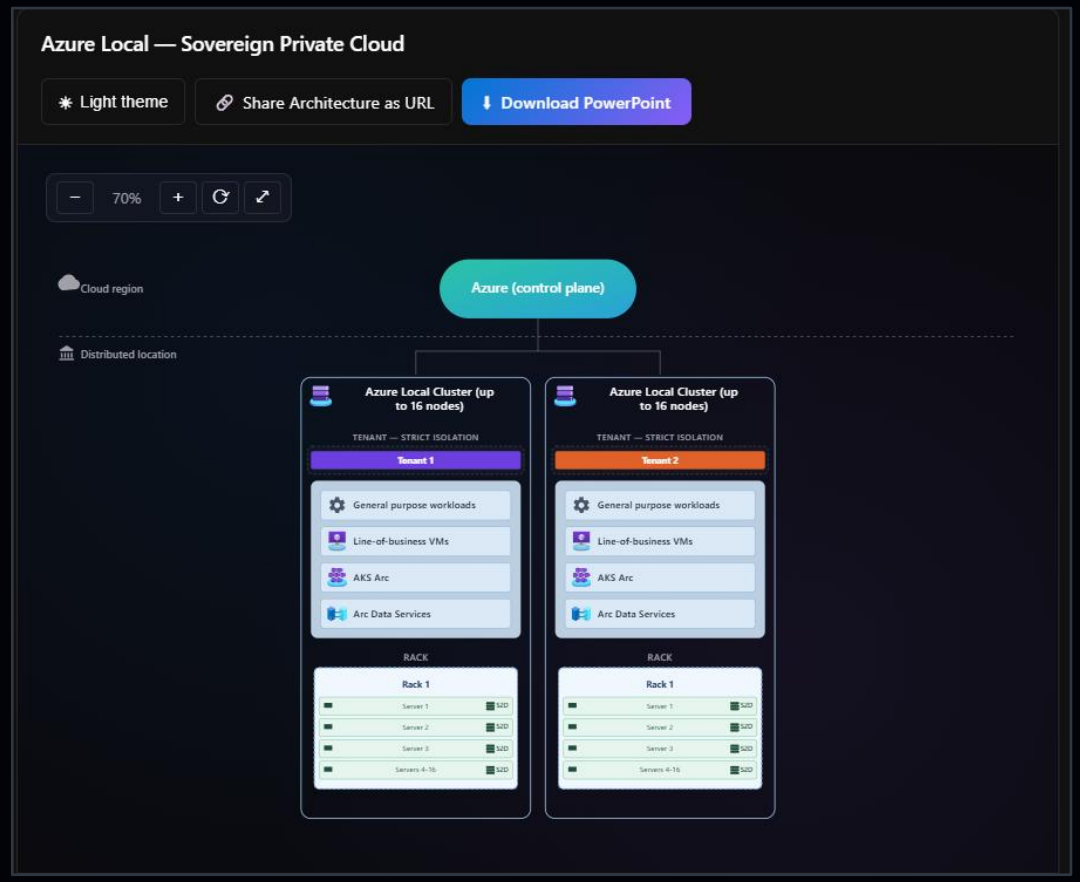

Designer ▸ topology & rack layout diagrams

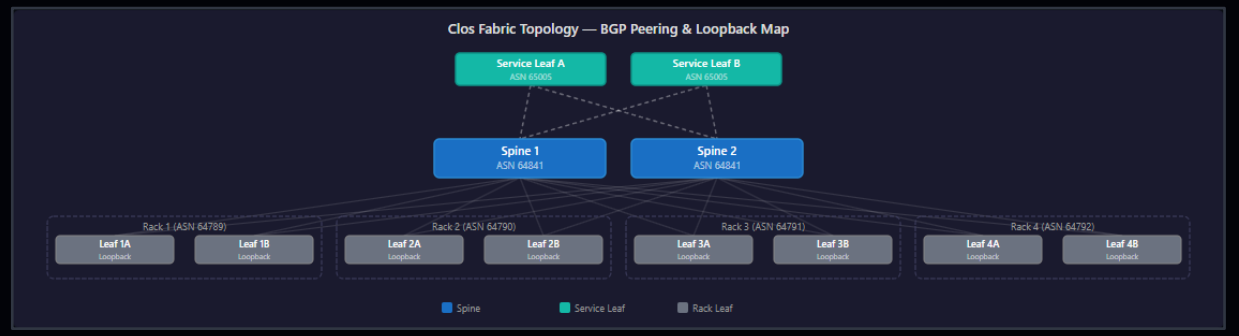

Also generated:

- ARM deployment templates
- ToR / BMC switch configuration
- Cluster design document

Every step produces shareable, deployment-ready output — architecture decks, validated topology and rack diagrams, switch configs, and ARM templates.

ODIN across the 11 decisions

How each design decision maps to an ODIN capability

Design decision	How ODIN helps you realize it
1 · Connectivity mode	The Designer decision tree starts with the deployment type — Connected, Disconnected, Multi-Rack, or Microsoft 365 Local.
2 · Architecture	Designer forks the design into hyperconverged or disaggregated and enforces valid architecture/storage/intent combinations.
3 · Cluster topology	Designer selects single-rack, rack-aware, or multi-rack and maps it to a validated network reference pattern.
4 · Cluster size	The Sizer turns your VM, AKS, and AVD profile into a supported node count and deployment shape.
5 · Storage connectivity	Designer picks switched/switchless S2D or FC/iSCSI SAN and validates the supported deployment shape.
6 · Adapter ports	Designer lays out ports per intent; Switch Config maps them to the ToR fabric with the right QoS.
7 · Traffic intents	Designer applies the Network ATC intent grouping that's supported for your storage type.
8 · Management IPs	Designer captures the IP pool, VLAN, and DNS, and flags overlaps with the reserved ARB ranges.
9 · Backup network	Designer adds the backup intent and ports where the chosen port layout supports it.
10 · Outbound connectivity	Designer sets the proxy / Arc gateway topology; Reports & ARM carry it into the deployment template.
11 · SDN	Designer records Arc-managed SDN (LNET/NSG) for HCI or fabric-based SDN for disaggregated.