Cisco Cyber Vision Performance and Scale Guide, Release 5.5.x

Cisco Cyber Vision
Updated April 28, 2026

Topics included

1 Performance and Scale Overview

Topics:

Introduces the Cisco Cyber Vision Release 5.5.x performance and scale guide, planning scope, audience, and performance and scale updates.

Performance and scale guide overview

This guide helps architects, deployment engineers, and performance owners plan Cisco Cyber Vision Release 5.5.x deployments for scale, recovery, and validation.

The Cisco Cyber Vision Performance and Scale Guide provides planning guidance for Cisco Cyber Vision Release 5.5.x Centers, Global Centers, CVSM, Sensors, VM and cloud deployments, bandwidth, recovery behavior, retention, and validation.

Use this guide to size deployments with performance margin and to identify conditions that can affect performance before production rollout.

Planning areas

This guide covers these planning areas:

- Center sizing.
- Global Center sizing.
- CVSM sizing.
- Sensor selection.
- Retention planning.
- Bandwidth planning.
- Performance validation.

Release 5.5.x planning scope

Use this guide as Cisco Cyber Vision Release 5.5.x planning guidance.

This guide is scoped to Cisco Cyber Vision Release 5.5.x.

- Apply numeric limits, platform references, VM tiers, deployment-model assumptions, and compatibility statements to the Release 5.5.x release family.
- Use cloud sizing guidance only for deployments on the specified cloud platform.
- For VMware ESXi, Hyper-V, or Nutanix AOS deployments, use the Release 5.5.x installation and compatibility documentation before final sizing.
- For other Cyber Vision release families, use the documentation set for that release.

Release 5.5.x performance and scale updates

Release 5.5.x introduces a published M8N Center scale point, keeps Global Center scale on the published M6N figures, and adds planning inputs for FIPS, Snort, Center DPI, and backup and restore.

Use these Release 5.5.x updates when revising performance and scale plans.

Table 1: Release 5.5.x planning updates

Planning area	Release 5.5.x planning guidance
Published M8N Center scale	Release 5.5.x introduces a new scale for a Center running on the CV-CNTR-M8 or an equivalent VM: 70,000 components, 400 Sensors, and 21 million stored flows.

Planning area	Release 5.5.x planning guidance
Global Center scale	Global Center scale remains 150,000 synced components and 20 registered Centers on CV-CNTR-M6N or VM resources equivalent to M6N.
FIPS	Release 5.5.x adds FIPS-compliant Center packages and FIPS-specific Sensor constraints.
Snort and Center DPI	Account for Snort IDS and Center DPI load when either feature is enabled.
Backup and restore	Plan backup and restore workflows as part of Center storage and recovery sizing.

Audience and scope

This topic identifies the intended audience and the performance planning scope for the Release 5.5.x guide.

Use this scope statement to determine whether the guide applies to your deployment planning activity.

Audience

- Architects who design Cisco Cyber Vision deployments.
- Deployment engineers who size Center and Sensor infrastructure.
- Teams who manage performance, storage, and growth.

Scope

- Cisco Cyber Vision Release 5.5.x performance and scale planning.
- Standalone Center and Local Center sizing.
- Global Center scale considerations.
- CVSM scale considerations.
- Sensor packet-rate and flow planning.
- Center DPI impact.
- VM and cloud deployment baselines.
- FIPS and cloud constraints that affect deployment planning.
- Retention, data growth, outage recovery, and replay behavior.

Limitations

This guide does not replace Cisco release notes, support matrices, or product installation guides. For non-Release 5.5.x deployments, use the documentation set for that release.

2 Platform and Center Sizing

Topics:

Provides platform, Center, CVSM, and Global Center sizing guidance.

Scale drivers

Cisco Cyber Vision scale depends on component count, flow volume, Sensor activity, Center DPI, retention, synchronization, and recovery behavior.

Multiple workload dimensions drive Cyber Vision scale. Two deployments with the same component count can have different CPU, storage, and WAN behavior.

Core scale drivers

- Discovered devices and components.
- Stored flows.
- Managed Sensors.
- Center DPI usage.
- Packet rate at each Sensor.
- Protocol mix and metadata volume.
- Retention period.

Feature and deployment scale drivers

- Synchronization volume to a Global Center.
- Replay bursts after outages.
- Snort IDS usage.
- Cloud VM, storage, and network profile.

Release 5.5.x platform scale reference

Release 5.5.x platform planning includes the CV-CNTR-M6N and CV-CNTR-M8 Center appliance classes, with Global Center scale kept on the published M6N figures.

Use these platform values as the Release 5.5.x Center and Global Center hardware scale reference.

Table 2: M6N and M8N hardware scale reference

Platform	CPU class	RAM	Storage	Published scale highlights
CV-CNTR-M6N	24-core AMD	128 GB	Two or four 1.6 TB NVMe drives	50,000 components, 300 Sensors, and 16 million stored flows
CV-CNTR-M8	32-core AMD	192 GB	Two or four 1.6 TB NVMe drives	70,000 components, 400 Sensors, and 21 million stored flows

Use the CV-CNTR-M8 or an equivalent VM when a Center requires the 70,000-component, 400-Sensor, 21-million-flow scale point.

Global Center scale values

For Global Center deployments, Release 5.5.x stays on the already published M6N figures. Use CV-CNTR-M6N or VM resources equivalent to M6N to support up to 20 registered Centers and 150,000 synced components.

Compatibility notes

The Release 5.5.x release notes include compatibility entries for CV-CNTR-M5S5 and CV-CNTR-M5S3 appliance classes. Use CV-CNTR-M8 or equivalent resources for new larger Center deployments, and use M6N-class resources for Global Center planning.

Center sizing factors and tiers

Size Cyber Vision Centers by considering component count, stored flows, Sensor count, Center DPI, retention, replay behavior, growth, FIPS mode, and VM or cloud requirements.

Use the sizing factors and starting points in this topic to define the initial deployment size.

Workload sizing factors

- Component count.
- Flow volume.
- Sensor count.
- Center DPI usage.
- Snort IDS usage.
- Retention duration.
- Outage replay behavior.

Deployment sizing factors

- Expected growth over 12 to 24 months.
- FIPS or non-FIPS deployment mode.
- Cloud-specific sizing requirements.

Table 3: Recommended Center sizing tiers

Center tier	Scale point	CPU	RAM	Storage and platform
Mid-size Center VM	Up to 20,000 components and 150 Sensors without Center DPI	16 cores	64 GB	1 TB NVMe storage with validated or provisioned IOPS
Large Center or Center DPI deployment	Up to 70,000 components, 400 Sensors, and 21 million stored flows, or any Center deployment with Center DPI	32-core AMD or equivalent CPU resources	192 GB	CV-CNTR-M8 with two or four 1.6 TB NVMe drives, or an equivalent VM storage profile with NVMe performance

Center VM equivalency

For a non-cloud Center VM, use only the Center sizing points in this topic unless the target deployment is covered by a cloud-specific sizing table.

A VM that is equivalent to the CV-CNTR-M8 must provide the same CPU resources, RAM, and disk performance as the appliance. NVMe-class storage is mandatory for the 70,000-component Center scale point.

CVSM sizing

Size CVSM separately from Center and Global Center workloads.

A CVSM VM with 8 vCPUs and 16 GB RAM can manage up to 100 enrolled Centers. The maximum supported limit is 100 enrolled Centers per CVSM VM; allocating additional CPU or memory does not increase this limit.

Global Center scope

For Global Center deployments, use CV-CNTR-M6N or VM resources equivalent to M6N to support up to 20 registered Centers and 150,000 synced components. Do not use Center component-count tiers to size Global Center.

Cloud-specific scope

Cloud tier values are cloud-specific. Do not use an AWS, Azure, or Google Cloud table as a universal VM sizing table.

For VMware ESXi, Hyper-V, Nutanix AOS, or cloud deployments, use the Release 5.5.x installation and compatibility documentation before final sizing.

Center DPI planning

A Center deployed with the CV-CNTR-M8N profile, or an equivalent CV-CNTR-M8-sized Center VM, supports up to four DPI interfaces.

Each DPI interface supports up to 300,000 packets per second (pps). Size DPI capacity per interface and keep steady-state load below 70 to 80 percent of the validated capacity.

Headroom planning

Do not design Cyber Vision production deployments to operate continuously at published maximum values.

Keep steady-state load below 70 to 80 percent of validated component, flow, and Sensor capacities.

- Reserve capacity for outage replay and recovery bursts.
- Reserve storage and write performance for retention growth.
- Revisit headroom after you enable Center DPI, Snort IDS, packet capture, or additional registered Centers.
- Treat this value as a planning recommendation, not as a Cisco-published support threshold.

CVSM scale planning

Size CVSM separately from Center and Global Center workloads.

CVSM sizing is based on the number of enrolled Centers that CVSM manages. Do not use Center component, Sensor, stored-flow, or Center DPI sizing tiers to size CVSM.

CVSM VM sizing

A CVSM VM with 8 vCPUs and 16 GB RAM can manage up to 100 enrolled Centers.

Related sizing

Size the managed Centers with the Center sizing guidance, and size any Global Center with the CVSM sizing guidance.

Global Center scale planning

Size a Global Center separately from a Local Center because synchronization activity and aggregate component count create different workload patterns.

A Global Center provides a central view of multiple Centers. It must be sized for synchronization volume, registered Center count, synced component count, storage growth, and replay behavior.

Planning inputs

- Number of registered Centers.
- Aggregate synchronization activity.
- Total synced component count.
- Storage growth over time.
- WAN behavior during normal operation and recovery.

Use 20 registered Centers and up to 150,000 synced components as Release 5.5.x planning inputs.

Scale planning

Use CV-CNTR-M6N or VM resources equivalent to M6N for Release 5.5.x Global Center deployments that support up to 20 registered Centers and 150,000 synced components. Size Global Center for synchronization scale and recovery behavior, not only for steady-state visibility.

Center-to-Global-Center synchronization is driven primarily by registered Center count, synced component count, activity changes, and recovery backlog. Do not size Global Center resources only from aggregate Sensor ingress traffic.

3 Sensor, VM, Cloud, and Network Planning

Topics:

Provides Sensor, FIPS, VM, cloud, bandwidth, and recovery planning guidance for Cisco Cyber Vision deployments.

Sensor performance planning

Plan Sensors by using observed packet rate, mirrored traffic scope, feature usage, and supported platform compatibility.

Sensor performance depends on the traffic that the Sensor receives and the features that the deployment enables. Site size and switch model alone do not provide enough information to select a Sensor.

Sensor selection factors

- Observed packet rate.
- Expected flow churn.
- Packet capture needs.
- Store-and-forward needs.
- Mirrored traffic scope.
- FIPS or non-FIPS deployment mode.

Traffic scope

If the monitored traffic scope is broader than the use case requires, reduce the traffic scope before you scale hardware.

If observed packet rate or DPI load exceeds the selected Sensor platform capacity, reduce mirrored traffic by selecting only the required ports or VLANs, or by applying capture filters.

Keep the traffic that is required to monitor the intended zones and conduits. Here's a list of max packets processed by the respective platforms:

Platform	Max packets per second
IC3000	12,000
IE3300 / IE3400	12,000
IE3500	13,000
IR1101 / IR1835	13,200
IR8300	15,000
IE9300	13,000
Catalyst 9300 / 9400	30,000
Center DPI	300,000

FIPS Sensor constraints

Use only FIPS-supported Center and Sensor combinations for Release 5.5.x FIPS deployments.

For FIPS deployments, Sensors that run a FIPS build can enroll only in Centers that run the FIPS build.

- Do not plan updates or downgrades between non-FIPS Centers and FIPS Centers.
- Do not use IC3000 Sensors in FIPS deployments.
- Do not use Docker Sensors or Virtual Machine Sensors in FIPS deployments.
- Do not rely on the Sensor management extension or Sensor self-update in FIPS deployments.

- Perform Sensor deployment and Sensor updates manually or with supported automation tools.

VM and cloud deployment planning

Release 5.5.x Center compatibility includes VMware ESXi, Hyper-V, Nutanix AOS, AWS, Azure, Google Cloud, and CVSM VM deployment planning.

Use this topic to separate Release 5.5.x Center compatibility from cloud-specific sizing guidance.

Release 5.5.x Center compatibility

- VMware ESXi 6.x and later.
- Microsoft Hyper-V Server 2016 or later.
- Nutanix AOS 6.10 or later.

Cloud sizing scope

Cyber Vision cloud deployment guidance includes AWS, Azure, and Google Cloud. Use the sizing values only for the cloud platform that the sizing table names.

For VMware ESXi, Hyper-V, Nutanix AOS, AWS, Azure, or Google Cloud deployments, use the Release 5.5.x installation and compatibility documentation before final sizing.

Center VM sizing

For a non-cloud Center VM, use these Release 5.5.x sizing points:

- Up to 20,000 components and 150 Sensors without Center DPI: 16 cores, 64 GB RAM, and 1 TB NVMe storage with validated or provisioned IOPS.
- Up to 70,000 components, 400 Sensors, and 21 million stored flows, or any Center deployment with Center DPI: resources equivalent to the CV-CNTR-M8, including equivalent CPU resources, 192 GB RAM, and equivalent NVMe disk performance.

Do not use lower or intermediate Center VM sizing points for Release 5.5.x planning.

CVSM VM sizing

A CVSM VM with 8 vCPUs and 16 GB RAM can manage up to 100 enrolled Centers.

VM and cloud design rules

- Use NVMe-class storage for Center VM deployments and high-performance cloud disk tiers for cloud deployments.
- Validate sustained write IOPS and throughput, not only raw capacity.
- Avoid storage oversubscription for larger Centers.
- Treat snapshots, backups, and maintenance windows as possible performance events.

Cloud sizing guidance

Use each cloud sizing table only for the Cisco Cyber Vision cloud platform that the table names.

Use this sizing guidance for Cisco Cyber Vision Center deployments on AWS, Azure, and Google Cloud. Size CVSM and Global Center deployments separately.

Cloud sizing scope

Cloud sizing values are platform-specific. Do not translate AWS, Azure, or Google Cloud sizing values from one cloud platform to another.

The Center tables list only the reviewed Release 5.5.x Center sizing points. CVSM and Global Center sizing are separate from these Center component-count tiers.

Before final sizing, use the cloud installation and compatibility documentation for the target release and cloud platform.

Table 4: AWS sizing tiers

Size point	CPU	RAM	Storage
Up to 20,000 components and 150 Sensors without Center DPI	Intel Xeon, 16 cores	64 GB minimum	1 TB SSD minimum with provisioned IOPS
Up to 70,000 components, 400 Sensors, and 21 million stored flows, or with Center DPI	Intel Xeon, 32 cores	192 GB minimum	2 TB SSD minimum with provisioned IOPS

Table 5: Azure sizing tiers

Size point	VM size	RAM	Storage
Up to 20,000 components and 150 Sensors without Center DPI	D16s_v4, 16 vCPUs	64 GB minimum	1 TB managed disk with provisioned IOPS
Up to 70,000 components, 400 Sensors, and 21 million stored flows, or with Center DPI	D48s_v4, 48 vCPUs	192 GB minimum	1 TB managed disk with provisioned IOPS

Table 6: Google Cloud sizing tiers

Size point	CPU	RAM	Storage
Up to 20,000 components and 150 Sensors without Center DPI	Intel Xeon, 32 vCPUs	128 GB minimum	1 TB Hyperdisk with provisioned IOPS
Up to 70,000 components, 400 Sensors, and 21 million stored flows, or with Center DPI	Intel Xeon, 64 vCPUs	256 GB minimum	1 TB Hyperdisk with provisioned IOPS

Google Cloud performance constraints

- Use Hyperdisk for Center storage. Hyperdisk provides higher IOPS and throughput limits than Persistent Disk.
- Use N2 or C3 machine series for production environments.
- Account for AMQP traffic over TCP port 5671 between Centers, Global Centers, and Sensors that exchange data.

FIPS and cloud constraints

Do not assume that public cloud marketplace images or non-standard cloud environments are FIPS-compliant.

Use the standard release unless the organization requires FIPS compliance.

- Do not assume that Cisco Cyber Vision products published in public cloud marketplaces are FIPS-compliant.
- Do not plan Cisco Cyber Vision deployments on AWS GovCloud or other non-standard cloud environments unless the Release 5.5.x documentation explicitly supports the environment.
- Do not use the report management extension in FIPS deployments.
- Validate FIPS package selection, Sensor compatibility, and cloud deployment model before production design approval.

Bandwidth and recovery behavior

Bandwidth demand can increase during recovery because Sensors and Centers may replay stored data after connectivity returns.

Bandwidth demand between Sensor and Center, and between Center and Global Center, depends on monitored traffic scope, protocol mix, synchronization activity, and recovery behavior.

A link that looks adequate during steady-state monitoring can fail during recovery if replay traffic is not included in the design.

These stages describe the recovery behavior to plan for:

1. A Sensor or Center loses connectivity to its destination.
2. The disconnected component stores data locally while connectivity is unavailable.
3. When connectivity returns, stored data can replay faster than steady-state collection.
4. WAN links, Center storage, and Global Center synchronization can experience recovery load.

Size remote links for normal operation and replay, identify where QoS or rate limiting is required, and test recovery after controlled outages before production sign-off.

4 Performance Validation and Design Guardrails

Topics:

Provides performance validation, troubleshooting, retention, and performance design guardrails for Cisco Cyber Vision performance planning.

Retention and data growth

Retention settings, stored-flow growth, and backup requirements affect storage pressure and recovery planning.

Retention is one of the main drivers of storage growth and database pressure. Backup and restore planning also affects storage and recovery design.

Retention guidance

- Keep flow retention conservative unless measured data supports a longer retention period.
- Measure storage growth before changing retention.
- Revisit IOPS headroom after any retention increase.
- Use data management and backup controls as storage and recovery controls.

Data-growth warning signs

- Rapid increase in stored flows.
- Rising write IOPS.
- Sustained storage growth after outages.
- Component growth that approaches the alert threshold.

Backup and restore planning

Release 5.5.x appliance documentation describes Center backup and restore workflows for migration and recovery.

Plan enough free space to generate backup archives locally and move the archives to secure storage.

Restore only onto targets with matching network interface and interface mode assumptions. Licenses and report extension packages are not restored automatically.

Validate performance before production

Validate normal load, peak behavior, recovery, storage growth, and FIPS constraints before production sign-off.

Use this task to verify that the deployment has enough margin for normal operation and recovery behavior.

Validation must include steady-state monitoring and failure recovery because replay behavior can stress links, storage, and synchronization paths.

Procedure

1. Validate normal monitoring load.
2. Validate packet-rate peaks on Sensors.
3. Validate Center DPI behavior if Center DPI is enabled.
4. Validate Snort IDS behavior if IDS is enabled.
5. Disconnect and reconnect a Sensor to validate replay behavior.
6. Validate Center-to-Global-Center recovery after a controlled synchronization interruption.
7. Measure storage growth over a representative window.
8. Validate backup export and restore procedures for recovery-critical Centers.
9. Validate FIPS Center and Sensor enrollment constraints if FIPS is required.

Results

The deployment is ready for production only when normal operation, recovery behavior, and storage growth stay within the planned margin.

Troubleshoot performance degradation

Check Sensor load, Center resources, synchronization backlog, WAN replay load, and feature-related load when performance degrades.

Use this task to isolate common causes of Cyber Vision performance degradation.

Performance degradation can originate at the Sensor, Center, Global Center, storage layer, or WAN path. Check the highest-risk areas first.

Procedure

1. Check Sensor packet rate and packet drops.
2. Check Center CPU, RAM, storage usage, and write IOPS.
3. Check synchronization backlog.
4. Check WAN utilization during replay.
5. Check mirrored traffic scope.
6. Check Snort or Center DPI load if either feature is enabled.
7. Check for FIPS compatibility mismatches if the deployment uses FIPS packages.

Results

The likely root cause is identified from Sensor oversubscription, storage latency, Center DPI load, unscoped Snort traffic, WAN congestion during recovery, or FIPS package mismatch.

Performance design guardrails

Use these guardrails to reduce performance risk in Cisco Cyber Vision deployments.

Design Cyber Vision deployments for sustained operation, recovery behavior, and future growth.

Use these capacity guardrails:

- Do not treat validated limits as design targets.
- Size for recovery, not only normal operation.
- Keep flow retention conservative until growth is measured.
- Plan Global Center storage and synchronization early.
- Reduce traffic scope before escalating hardware.
- Revisit sizing after adding Center DPI, packet capture, or more registered Centers.

Use these deployment guardrails:

- Do not treat monitoring visibility as a substitute for Center or Global Center scale planning.
- Do not mix FIPS and non-FIPS Center or Sensor package assumptions.
- Do not assume that cloud marketplace images are FIPS-compliant.
- For cloud deployments, use supported machine series or VM sizes and the reviewed high-performance storage tier for the target cloud.

5 Sizing Examples and Decision Flows

Topics:

Provides sizing examples and decision flows for Cisco Cyber Vision performance planning.

Sizing examples

Use these planning examples as starting points for Cisco Cyber Vision Release 5.5.x sizing discussions.

The examples summarize starting points for the reviewed Release 5.5.x Center and Global Center sizing tiers.

Table 7: Sizing examples

Scenario	Starting point	CPU and RAM	Storage
Up to 20,000 without Center DPI	Mid-size Center VM or matching target cloud sizing tier	16 cores and 64 GB, or the target cloud table values	1 TB NVMe or provisioned-performance cloud storage; use Hyperdisk on Google Cloud
Up to 70,000, 400 Sensors, and 21 million stored flows, or with Center DPI	CV-CNTR-M8, equivalent VM, or matching target cloud sizing tier	32-core AMD and 192 GB, or the target cloud table values	Two or four 1.6 TB NVMe drives or equivalent provisioned-performance cloud storage
Global Center: up to 150,000 synced components	CV-CNTR-M6N or VM resources equivalent to M6N	24-core AMD and 128 GB, or equivalent CPU and RAM resources	Two or four 1.6 TB NVMe drives or equivalent storage performance

Sizing example scope

Cloud-specific values apply only to the cloud platform that the sizing table names.

CVSM sizing is separate: use 8 vCPUs and 16 GB RAM for up to 100 enrolled Centers.

For VMware ESXi, Hyper-V, Nutanix AOS, AWS, Azure, or Google Cloud deployments, use the Release 5.5.x installation and compatibility documentation before final sizing.

Center sizing decision flow

Use this decision flow to select an initial Center sizing tier.

Center sizing starts with component count, Sensor count, Center DPI usage, deployment platform, and whether the system is a Center, Global Center, or CVSM deployment.

This flow identifies the starting tier. Validate the final design with pilot data and release-specific installation guidance.

Use these stages to select the initial sizing tier:

1. For a cloud deployment, select the AWS, Azure, or Google Cloud sizing tier that matches the target cloud platform.
2. For a non-cloud Center VM up to 20,000 components and 150 Sensors without Center DPI, start with 16 cores, 64 GB RAM, and 1 TB NVMe storage with validated or provisioned IOPS.
3. For a Center with Center DPI, or up to 70,000 components, 400 Sensors, and 21 million stored flows, use CV-CNTR-M8 or equivalent VM resources with equivalent CPU, 192 GB RAM, and NVMe disk performance.
4. For a Global Center, use CV-CNTR-M6N or VM resources equivalent to M6N for up to 20 registered Centers and 150,000 synced components.
5. For CVSM, use 8 vCPUs and 16 GB RAM for up to 100 enrolled Centers.
6. If the deployment is between tiers, use the higher tier and validate with pilot data.

The selected starting tier provides a conservative starting point for detailed design, validation, and support-matrix review.

Sensor selection decision flow

Use this decision flow to select a Sensor platform by traffic behavior, feature requirements, and FIPS compatibility.

Sensor selection requires observed packet-rate data and a clear understanding of mirrored traffic scope.

Choose the Sensor platform after you define traffic scope, packet capture needs, store-and-forward needs, and FIPS requirements.

Use these stages to select a Sensor:

1. Measure the observed packet rate on mirrored traffic.
2. If packet rate exceeds the supported platform limit after you apply design margin, select a higher-capacity Sensor or reduce traffic scope.
3. If packet capture or heavy store-and-forward is required, include the SSD requirement in the platform choice.
4. If the deployment requires FIPS, use only FIPS-supported Center and Sensor combinations.
5. If mirrored traffic is broader than the use case requires, reduce the traffic scope before changing hardware.
6. Validate the final selection in a pilot deployment.

The selected Sensor platform aligns with traffic load, feature requirements, and release-specific compatibility requirements.

Global Center scale decision flow

Use this decision flow to size a Global Center for registered Centers, synced components, and recovery behavior.

Global Center sizing depends on registered Center count, synced component count, WAN behavior, and synchronization recovery.

Plan a Global Center separately from a Local Center because synchronization scale, recovery behavior, and aggregate storage growth create a distinct workload.

Use these stages to evaluate Global Center scale:

1. Identify the number of Centers to register.
2. Estimate the projected synced component count.
3. If WAN paths are constrained or outage-prone, prioritize replay testing and storage headroom.
4. If the design approaches 20 registered Centers or 150,000 synced components, use CV-CNTR-M6N or VM resources equivalent to M6N and validate it separately.
5. If the design does not approach 20 registered Centers or 150,000 synced components, size for growth and replay, not only steady state.

The Global Center design accounts for synchronization scale, recovery stress, and storage growth.