

Chef 360 Migration Guide

Standalone Infra Server to
Chef 360 Platform

WHITEPAPER

Table of Contents

1. Scenario Description [/3](#)
2. Assumptions [/4](#)
3. Chef 360 Sizing Guidelines [/6](#)
4. Prerequisites [/9](#)
5. Migration Steps at a Glance [/11](#)
6. Step 1 – Pre-Migration Assessment [/11](#)
7. Step 2 – Install Chef 360 Platform (Parallel Cluster) [/14](#)
8. Step 3 – Install Migration Tools [/15](#)
9. Step 4 – Export Data from Chef Infra Server [/17](#)
10. Step 5 – Import Data into Chef 360 [/20](#)
11. Step 6 – Verify the Chef 360 Import [/22](#)
12. Step 7 – Re-Point Zones to Chef 360 [/24](#)
13. Step 8 – Enroll Nodes into Chef 360 Node Management (Optional) [/27](#)
14. Step 9 – Zone-by-Zone Parallel Migration Execution [/28](#)
15. Step 10 – Retire the Chef Infra Server [/30](#)
16. Post-Migration Validation Checklist [/31](#)
17. Rollback Plan [/32](#)
18. Troubleshooting [/33](#)
19. Reference: Key CLI Commands [/34](#)
20. Appendix [/36](#)

1. Scenario Description

“Standalone Infra Server deployment without Automate. Zones connect directly or via proxy to the Chef Infra Server. Customer prefers the parallel path for upgrade.” Chef 360 Upgrade Patterns Deck

Infra Server Only Parallel Clusters covers the migration of a standalone Chef Infra Server (no Chef Automate) to Chef 360 Platform. Nodes are distributed across deployment zones that connect directly or via proxy to the Chef Infra Server. The strategy is a **parallel cluster migration**: a new Chef 360 cluster is built alongside the existing server, data is migrated to Chef 360's **Declarative State Management (DSM)** service, zones are re-pointed one at a time, and the old server is retired only after all zones are validated.ted.

What is Chef DSM?

Chef Declarative State Management (DSM) is Chef 360 Platform's equivalent of the traditional standalone Chef Infra Server a central repository for cookbooks, policies, roles, environments, and node metadata, served to Chef Infra Client on managed nodes. DSM runs as a Kubernetes-managed service within Chef 360, providing horizontal scalability, self-healing, and RBAC-integrated security.

No new organizations need to be manually created the import process automatically creates organizations in Chef 360 DSM from the Chef Infra Server backup.

Reference: [Chef DSM Overview](#)

2. Assumptions

2.1 Infrastructure Assumptions

Assumption	Detail
Source platform	Standalone Chef Infra Server (v12 or later), no Chef Automate
Target platform	Chef 360 Platform 1.7, DSM service
Migration style	Parallel cluster new Chef 360 built alongside existing Chef Infra Server
Total node count	Maximum 8,000 nodes across all zones
Number of zones	3 zones (Zone 1, Zone 2, Zone 3)
Nodes per zone	~2,000–3,000 nodes per zone
Node connection	Each zone connects directly or via proxy to the Chef Infra Server

2.2 What “Zone” Means in This Context

A zone is both a physical deployment zone AND a Chef environment (or Policyfile policy group).

Each zone maps 1:1 to a **Chef environment** – or a **Policyfile policy group** for nodes managed with Policyfiles – on the Chef Infra Server. Policy groups are the Policyfile equivalent of environments and are treated the same way for migration sequencing purposes.

Note – complex zone topologies:

If a single physical deployment zone contains multiple environments or policy groups, or if a single environment/policy group spans multiple physical zones, contact your Progress Chef Customer Architect or Professional Services team before proceeding. The zone-by-zone migration approach described in this guide assumes a 1:1 mapping.

Each zone maps as follows:

Dimension	Zone 1	Zone 2	Zone 3
Physical	Network segment / DC / AZ 1	Network segment / DC / AZ 2	Network segment / DC / AZ 3
Chef environment / Policy group	zone1	zone2	zone3
Migration order	First (lowest risk)	Second	Last (most critical)

Nodes in Zone 1 carry chef_environment: zone1 (or belong to policy group zone1) on the Chef Infra Server. Zone membership drives the migration sequencing – start with the least critical zone.

2.3 Node and Data Assumptions

Assumption	Detail
Chef Infra Client version	Chef Infra Client 15 or later
OS platforms	Linux (RHEL / Ubuntu / CentOS) and/ or Windows Server
Org structure	Single or multi-org; all orgs migrate into one Chef 360 tenant
Large backups	Backups > 5 GB require AWS CLI for upload; > 10 GB may take several hours to import

2.4 Migration Window Assumptions

Assumption	Detail
Migration timing	Off-peak hours to minimize impact
Parallel window	Both systems active during zone-by-zone migration
Validation period	Minimum 72 hours per zone before proceeding
Rollback window	Chef Infra Server stays live until final decommission

2.5 Out of Scope

- Chef Automate migration (not applicable Scenario #6 is Infra Server Only)
- Air-gapped environments (contact Progress Chef for separate guidance)
- Merging multiple Chef Infra Server deployments into one DSM

3. Chef 360 Sizing Guidelines

Chef 360 provides an official web-based **Sizing Calculator** to estimate the infrastructure requirements for your deployment. Use this tool to determine CPU, memory, and storage needs before provisioning your Chef 360 cluster.

Calculator URL: <https://sizing-calculator.cloud.chef.io/>

Disclaimer: The sizing calculator provides estimates based on the inputs provided. Actual infrastructure requirements may vary depending on workload characteristics, environment-specific factors, and usage patterns. Final sizing should be validated based on detailed customer requirements and testing.

3.1 How to Use the Sizing Calculator

Follow these steps to generate a sizing estimate for your Chef 360 deployment:

Step 1 – Open the Calculator

Navigate to <https://sizing-calculator.cloud.chef.io/> in a web browser. The calculator automatically loads the latest Chef 360 specification (currently Chef 360 v2026.03). No login is required.

Step 2 – Select Deployment Topology

Under **Chef 360 Deployment Architecture**, choose your deployment topology from the **Deployment Topology** dropdown:

Topology	Description
Hyperconverged Non-HA	Single cluster, no high availability – suitable for dev/test
Hyperconverged HA	Single cluster with high availability – recommended for production
Tiered HA	Separated control-plane and data-plane tiers with HA
Hyperscale HA	Large-scale deployment with high availability for enterprise fleets
BYOK Non-HA	Bring Your Own Kubernetes cluster, no HA
BYOK HA	Bring Your Own Kubernetes cluster with high availability

Step 3 – Configure Nodes

- **Nodes Under Management:** Total number of nodes to be managed by Chef 360
- **Agent Nodes:** Subset of nodes that will run the Node Management Agent; the remainder are agentless
- **Operating System:** Select Linux Standard, Linux Hardened, or Linux Container OS based on your node OS
- **Rolling Upgrades:** Enable if you require rolling upgrade capability for managed nodes

Step 4 – Select Features / Capabilities

Select the Chef 360 features you plan to enable. Each feature adds specific resource requirements:

Feature	Required for This Migration	Purpose
DSM	Yes – core requirement	Manages cookbooks, roles, environments, and node check-ins (replaces Chef Infra Server)
Courier	Optional (Step 9)	Remote job execution, scheduled automation, InSpec compliance scans
Reporting	Optional	Run report storage and dashboards within Chef 360
Habitat	Optional	Service management and continuous deployment using Habitat
Event Management	Optional	Event-driven automation triggers
Content Service	Optional	Centralised content repository for cookbooks and packages

Step 5 – Configure DSM Parameters (shown when DSM is selected)

- **Converge Frequency:** How often nodes converge (default: 30 minutes)
- **Splay:** Random delay before convergence to avoid thundering herd (default: 15 minutes)
- **Converge Duration:** Average convergence time per node (default: 5 minutes)
- **Cookbooks:** Number of cookbooks in the system
- **Search Utilization:** LOW / MED / HIGH based on frequency of knife search usage in your workflows

Step 6 – Generate and Download the Sizing Report

Click Generate Sizing Estimate. The results display required CPU, memory, and storage per component with totals including headroom. Click Download PDF Report to save the estimate for infrastructure provisioning and approval workflows.

3.2 Recommended Starting Inputs for This Scenario

Use the following starting inputs when running the calculator for this scenario (Standalone Infra Server → Chef 360 DSM, up to 8,000 nodes):

Input	Recommended Starting Value	Notes
Nodes Under Management	Your actual node count (up to 8,000)	Use exact count from premigration inventory (Section 6.1)
Agent Nodes	Same as Nodes Under Management	Adjust downward if some nodes will remain agentless
Deployment Topology	Hyperconverged HA (production) or Hyperconverged Non-HA (dev/test)	Select based on your HA and availability requirements
Operating System	Linux Standard	Change to Hardened if nodes run RHEL or hardened Linux
Features	DSM (required) + Courier (if Step 8 is planned)	Add Reporting if Chef 360 reporting dashboards are needed
Converge Frequency	30 minutes (default)	Match your current chef-client interval
Splay	15 minutes (default)	Match your current splay setting
Search Utilization	LOW	Increase to MED or HIGH if your workflows use heavy knife search
Cookbooks	Actual count from pre-migration inventory	Use value from Section 6.1 output

Recommendation: Share the calculator PDF report with your Progress Chef account representative or Chef Support to obtain a validated sizing specification before provisioning infrastructure.

4. Prerequisites

For complete prerequisites details, refer to the **Before you begin** section of the official DSM migration guide: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/

4.1 Chef Infra Server (Source)

Reference: [Chef 360 Platform 1.7.0 Release Notes](#)

Requirement	Detail
Chef Infra Server version	Bundled Kubernetes: v1.32.12+k0s.1 (k0s distribution)
Admin access	Administrative access to the Chef Infra Server to export data
knife-ec-backup	Habitat package version 3.0.6 or greater
Disk space	Sufficient space on the Chef Infra Server or migration workstation to store the backup tar file
Connectivity	Migration workstation must be able to reach the Chef Infra Server API

Note on pivotal key extraction:

The process for extracting the Chef Infra Server pivotal key required for backup is documented in the “Get the pivotal private key” section of the official DSM migration guide. Always follow that guide do not rely on the `/etc/opscode/pivotal.pem` file on disk as it may be outdated. See: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/#get-the-pivotal-private-key

4.2 Chef 360 Platform (Target)

Requirement	Detail
Chef 360 Platform version	1.6.1 or later (1.7 recommended)
Deployed and running	Chef 360 must be installed and accessible before migration begins
FQDN + TLS	Chef 360 FQDN known; reachable from migration workstation and all nodes
Admin credentials	Required to register chef-import-cli with Chef 360

4.3 Migration Workstation

Requirement	Detail	Documentation
chef-import-cli	Downloaded from Chef 360 Platform UI Download Centre	Chef Import CLI reference
habitat (hab)	Installed — used to run knife-ec-backup Habitat package	Install Habitat
knife-ec-backup	v3.0.6 or greater (hab pkg install chef/knife-ec-backup -bf)	knife-ec-backup on GitHub
AWS CLI	Required if backup file is 5 GB	Install AWS CLI
Chef Workstation (knife)	Required to run knife commands throughout this guide	Install Chef Workstation
Network access	Reach both Chef Infra Server and Chef 360 Platform APIs	System requirements

4.4 Node Requirements (Per Zone)

Protocol	Requirement
Linux / SSH	Port 22 open; sudo privileges; RSA or ed25519 SSH key
Windows / WinRM	Ports 5985 and 5986 open; WinRM configured
All	Public DNS or IP (not 127.0.0.1)
All	No IP overlap with 10.244.0.0/16 or 10.96.0.0/12
All	Chef Infra Client installed and running

Reference: [Node Enrollment Requirements](#)

5. Migration Steps at a Glance

The following steps map directly to the Scenario #6 slide progression (Slides 4956):

Step Action Execution Location

1. Pre-Migration Assessment Migration Workstation
2. Install Chef 360 (parallel cluster) Chef 360 K8s host
3. Install migration tools Migration Workstation
4. Export data (backup) from Chef Infra Server Chef Infra Server
5. Import data into Chef 360 DSM Migration Workstation
6. Verify DSM import Chef 360 UI / CLI
7. Re-point zones to Chef 360 DSM /DNS/ per option
8. [Optional] Enroll nodes into Node Mgmt Nodes (via cookbook)
9. Zone-by-zone parallel migration All locations
10. Retire Chef Infra Server Chef Infra Server

Execution location legend used throughout this document:

Chef Infra Server run directly on the Chef Infra Server host – **Migration**

Workstation run on the workstation managing the migration – **Chef 360**

Platform actions in Chef 360 UI or via Chef 360 CLIs – **On Node** run on the managed node (or pushed via knife ssh/winrm)

6. Step 1: Pre-Migration Assessment

Execution location: Migration Workstation

6.1 Inventory All Nodes Across All Zones

Confirm the total node count is within the 8,000 node assumption and capture a counts-only snapshot for post-migration comparison. Save the output to pre-migration-counts.txt – this file is referenced again in Step 6 (Section 11) and the Post-Migration Validation Checklist (Section 16).

This script must be run as the pivotal user so it can enumerate all organizations. Configure knife with the pivotal key before running:

```
node_name 'pivotal'
client_key '<PATH_TO_PIVOTAL_KEY>'
chef_server_url 'https://<CHEF_INFRA_SERVER>'
```

See: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/#get-the-pivotal-private-key

```
CHEF_SERVER_URL="https://<CHEF_INFRA_SERVER>"
PIVOTAL_KEY="<PATH_TO_PIVOTAL_KEY>"
{
  echo "=== Pre-Migration Counts ($(date)) ==="
  echo ""

  ORGS=$(knife org list -s "${CHEF_SERVER_URL}" -u pivotal -k "${PIVOTAL_KEY}")
  echo "Total organizations: $(echo "${ORGS}" | wc -l)"
  echo ""

  for ORG in ${ORGS}; do
    ORG_URL="${CHEF_SERVER_URL}/organizations/${ORG}"
    echo "--- Org: ${ORG} ---"

    echo " Node Counts:"
    echo " Total nodes : $(knife node list -s "${ORG_URL}" -u pivotal -k "${PIVOTAL_KEY}" | wc -l)"

    for ZONE in zone1 zone2 zone3; do
      echo " Zone ${ZONE} : $(knife search node "chef_environment:${ZONE}" -i -s "${ORG_URL}" -u pivotal -k "${PIVOTAL_KEY}" | wc -l) nodes"
    done

    echo " Chef Server Objects:"
    echo " Cookbooks : $(knife cookbook list -s "${ORG_URL}" -u pivotal -k "${PIVOTAL_KEY}" | wc -l)"
    echo " Roles : $(knife role list -s "${ORG_URL}" -u pivotal -k "${PIVOTAL_KEY}" | wc -l)"
    echo " Environments: $(knife environment list -s "${ORG_URL}" -u pivotal -k "${PIVOTAL_KEY}" | wc -l)"
    echo " Data bags : $(knife data bag list -s "${ORG_URL}" -u pivotal -k "${PIVOTAL_KEY}" | wc -l)"
    echo " Policies : $(knife raw /policies -s "${ORG_URL}" -u pivotal -k "${PIVOTAL_KEY}" 2>/dev/null | python3 -c 'import json,sys; print(len(json.load(sys.stdin)))' 2>/dev/null || echo 'N/A')"
    echo ""
  done
} | tee pre-migration-counts.txt
```

Keep pre-migration-counts.txt — you will compare these numbers against Chef 360 DSM in the post-migration validation steps (Sections 11 and 16).

6.2 Review Organization Names

List all organization names on the Chef Infra Server and check for conflicts with any existing organizations in Chef 360. Organizations with the same name will be overwritten during import.

```
# List all organization names identify naming conflicts before import
knife org list
```

Reference: [Organization conflict handling during import](#)

6.3 Record Current chef_server_url from Node client.rb

Document the current chef_server_url used in each zone this is what will change when re-pointing in Step 7:

```
# On a sample node per zone (or via knife ssh)
grep chef_server_url /etc/chef/client.rb
# Example: chef_server_url "https://chef-infra.example.com/organizations/myorg"
```

Zone	Current chef_server_url	Target chef_server_url (after migration)
Zone 1	https://<CHEF_INFRA_SERVER>/organizations/<ORG>	https://<CHEF_360_FQDN>/organizations/<ORG>
Zone 2	https://<CHEF_INFRA_SERVER>/organizations/<ORG>	https://<CHEF_360_FQDN>/organizations/<ORG>
Zone 3	https://<CHEF_INFRA_SERVER>/organizations/<ORG>	https://<CHEF_360_FQDN>/organizations/<ORG>

6.4 Pre-Migration Checklist

- Total node count confirmed <= 8,000
- Zone-to-Chef-environment / policy group mapping documented
- All organization names listed; naming conflicts with Chef 360 identified
- All current chef_server_url values recorded from client.rb
- No CIDR conflicts on zone nodes
- Migration window confirmed (off-peak hours)
- Stakeholder sign-off obtained

7. Step 2: Install Chef 360 Platform (Parallel Cluster)

Execution location: Chef 360 K8s host

“Customer installs Chef 360 in planned target topology” Slide 51

Install Chef 360 Platform as a new, separate cluster running in parallel with the existing Chef Infra Server. The Chef Infra Server continues serving all zones during this step.

Key principles:

Deploy on a **new host / new Kubernetes cluster** never on the same host as the existing Chef Infra Server – Ensure pod CIDR does not conflict with zone node network ranges – Configure ingress/TLS for HTTPS on port 443 – Record the Chef 360 FQDN (e.g., <https://chef360.example.com>)

Follow the official Chef 360 Platform installation guide:

<https://docs.chef.io/360/1.7/install/> For BYOK (Bring Your Own Kubernetes) deployments introduced in 1.7.0, contact Progress Chef for the BYOK installation guide.

8. Step 3 Install Migration Tools

Execution location: Migration Workstation & Chef Server

Follow the “Download the required tools” section of the official DSM migration guide for exact installation steps: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/#download-the-required-tools

The five tools required are:

1. **Habitat (hab) [Migration Workstation]** – Install Habitat: <https://docs.chef.io/habitat/latest/install/>
2. **chef-import-cli [Migration Workstation]** – Download from Chef 360 Platform UI › Download Centre › Chef Platform Bundled Tools › Chef Import CLI
3. **knife-ec-backup gem [Chef Server]** – Required to execute knife ec backup directly on the Chef Infra Server. Reference: <https://github.com/chef/knife-ec-backup>
4. **Chef Workstation (includes knife) [Migration Workstation]** – Required to run knife commands throughout Steps 4, 6, 7, and 9 of this guide. <https://docs.chef.io/workstation/>
5. **AWS CLI [Migration Workstation]** – Required only if backup file is ≥5 GB (expected for 8K node environments). <https://docs.aws.amazon.com/cli/latest/userguide/getting-startedinstall.html>

Register the migration workstation with Chef 360:

Follow the “Register your device and configure authentication” section of the official DSM migration guide: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/#register-your-device-and-configure-authentication

```
# Migration Workstation
chef-import-cli register-device \
  --url "https://<CHEF_360_FQDN>" \
  --profile-name "chef-server-import" \
  --device-name "<DEVICE_NAME>"

chef-import-cli set-default-profile chef-server-import
```

8.1 Network Readiness Checks

With Chef 360 now installed, verify workstation connectivity to Chef 360 DSM and check for node-level CIDR conflicts.

Verify workstation connectivity to Chef 360 DSM

The official method to verify workstation connectivity is to configure knife and run a test command. Note that the Chef 360 DSM endpoint uses **port 31000**.

Reference: [Connect to Chef DSM](#)

```
# Migration Workstation - configure knife to point to Chef 360 DSM
# In ~/.chef/config.rb or ~/.chef/credentials:
# chef_server_url "https://<CHEF_360_FQDN>:31000/organizations/<ORG_NAME>"
# node_name "<DSM_USER_NAME>"
# client_key "~/.chef/<KEY_FILE_NAME>.pem"

# Fetch and trust SSL certificates from Chef 360
knife ssl fetch

# Verify connectivity - a successful node list confirms the workstation can
# reach Chef 360 DSM
knife node list --profile chef360
```

If Chef 360 uses self-signed certificates, retrieve and trust the root CA first:

```
curl -k https://<CHEF_360_FQDN>/platform/system/v1/tenant/root-ca \
| jq -r '.item.rootCa' \
| sed 's/\\n/\\n/g' \
> ~/.chef/trusted_certs/chef360-root-ca.crt
```

Check for CIDR conflicts on nodes

```
# On a sample node - must return empty (no conflicts with Chef 360 pod CIDRs)
ip route | grep -E "10\.244\.|10\.96\."
```

9. Step 4 Export Data from Chef Infra Server

“Customer performs a data backup” Slide 51

Follow the “Export data from Chef Infra Server” section of the official DSM migration guide for the complete, authoritative steps: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/#export-data-from-chef-infra-server

This step covers three sub-tasks follow the official guide for each:

9.1 Extract the Pivotal Private Key

Execution location: Chef Infra Server

The pivotal key is required to authenticate knife-ec-backup . The extraction method uses an embedded binary on the Chef Infra Server. Always extract it from the Chef secrets vault never use the /etc/opscode/pivotal.pem file on disk as it may be out of date.

Follow exactly: [Get the pivotal private key](#)

9.2 Clean Up Stale Data Before Backup

Execution location: Chef Infra Server

Prerequisite: The knife-ec-backup Ruby gem must be installed on the Chef Infra Server before running this step. Reference: <https://github.com/chef/knife-ec-backup>

For 8,000-node environments, cleaning stale nodes (not checked in for 60+ days) before backup significantly reduces backup size and import time.

Follow: [Clean up stale data](#)

```
# Chef Infra Server — generate stale node report
knife tidy server report \
  --node-threshold 60 \
  -s <CHEF_INFRA_SERVER_URL> \
  -u pivotal \
  -k <PATH_TO_PIVOTAL_KEY>
```

9.3 Create the Chef Infra Server Backup

Execution location: Chef Infra Server

Prerequisite: The knife-ec-backup Ruby gem must be installed on the Chef Infra Server before running this step. Reference: <https://github.com/chef/knife-ec-backup>

Follow: [Create a Chef Infra Server backup](#)

Run the following command directly on the Chef Infra Server to create the backup. The pivotal key extracted in Step 9.1 is required:

```
# Chef Infra Server — create the backup using knife-ec-backup gem
knife ec backup backup_$(date '+%Y%m%d%H%M%s') \
--webui-key /etc/opscode/webui_priv.pem \
--with-key-sql \
-s <CHEF_INFRA_SERVER_URL> \
-u pivotal \
-k <PATH_TO_PIVOTAL_KEY>
```

Flag	Purpose
--with-key-sql	Exports ALL public keys for users/clients (not just default key) required for multi-key environments
--webui-key	Path to the webui private key on the Chef Infra Server
-s	Chef Infra Server URL
-u	Admin user (pivotal)
-k	Path to the pivotal key extracted in Step 9.1

After the backup directory is created, compress it and transfer it to the Migration Workstation to be used in the next steps for cleaning and importing data into Chef 360:

```
# Chef Infra Server — compress and transfer backup to Migration Workstation
tar -czf /tmp/chef-server-backup.tar.gz -C /tmp <BACKUP_DIR_NAME>/

# Transfer to Migration Workstation
scp /tmp/chef-server-backup.tar.gz <MW_USER>@<MIGRATION_WORKSTATION_IP>:/migration/
```

9.5 Generate SHA256 Checksum for Import

Execution location: Migration Workstation

With the backup transferred to the Migration Workstation, clean up unused cookbooks and stale objects from the backup directory using `knife tidy backup clean`. This reduces import time and avoids importing stale data into Chef 360.

Reference: `knife tidy backup clean` — https://docs.chef.io/360/1.7/reference/cli/knife/knife_tidy/

```
# Migration Workstation — clean up backup directory
hab pkg exec chef/knife-ec-backup knife tidy backup clean \
--backup-path /migration/chef-server-backup
```

9.5 Generate SHA256 Checksum for Import

Execution location: Migration Workstation

```
# Migration Workstation
chef-import-cli importservice generate-sha256 \
--file /migration/chef-server-backup.tar.gz

# Save the output required in Step 10
```

10. Step 5: Import Data into Chef 360

“Customer restores backup to a Chef 360 organization” Slide 51

Execution location: Migration Workstation

Follow the “Import data into Chef DSM” section of the official DSM migration guide for the complete, authoritative steps: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/#import-data-into-chef-dsm

10.1 Organization Import Behaviour Read Before Proceeding

Scenario	Behaviour
Org in backup, not in Chef 360	Automatically created
Org in backup, already in Chef 360	Overwritten with backup data
Org in Chef 360, not in backup	Left untouched

If a conflict exists that you do not want overwritten, rename or delete the Chef 360 org before importing.

Reference: [Review organization names](#)

10.2 Upload and Initiate the Import

Follow the official guide to:

1. Create an import process and obtain the signed upload URL
2. Upload the backup file (use curl for < 5 GB; use AWS CLI for 5 GB)
3. Obtain and save the returned IMPORT_ID

Follow: [Upload the backup file](#)

10.3 Monitor Import Progress

After upload, Chef 360 DSM automatically begins processing. The import progresses through these stages:

Stage	Description
waiting-for-upload	Awaiting file upload
acquiring-backup	Downloads and validates the backup
extracting	Extracts the archive
restoring-orgs	Creates organizations in Chef 360
restoring-users	Imports users and organization membership
restoring-data	Restores cookbooks, roles, environments, data bags, nodes
complete	Import finished

```
# Migration Workstation poll status periodically
chef-import-cli importservice get-import-process-status \
  --importId "<IMPORT_ID>" \
  --profile chef-server-import \
  --format json
```

Status	Meaning	Action
complete + errors: []	Full success	Proceed to Step 11
partial-success	Non-critical errors	Review errors; manually correct failed items
failure	Import failed	Check error field and email notification; see Section 18

Follow: [Monitor the import progress and results](#)

Expected duration for large datasets (8K nodes): 25+ hours depending on backup size and network bandwidth. The official guide notes that large numbers of nodes (>1,000) increase import time.

11. Step 6: Verify the Chef 360 Import

Execution location: Chef 360 UI + Migration Workstation

Follow the “Verify the import” section of the official DSM migration guide:

https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/#verify-the-import

11.1 Verify in Chef 360 UI

Reference: [Verify the import official steps](#)

Follow these steps from the official DSM migration guide to verify the import via the Chef 360 Platform UI:

1. Go to the DSM page in the [Chef 360 Platform UI](#)
2. Check that your organizations appear in the organization list
3. Select an organization to view:
 - > Nodes
 - > Cookbooks
 - > Roles
 - > Environments (including zone1 , zone2 , zone3)
 - > Data bags
4. Verify that users are members of the correct organizations
5. Delete the backup tar file and any private keys used during export, or move them to secure storage (see Section 11.3)

11.2 Verify via Knife Profiles

Execution location: Migration Workstation

Configure dual knife profiles to compare Chef Infra Server and Chef 360 DSM:

```

# Migration Workstation add to ~/.chef/credentials
[chef360]
chef_server_url = "https://<CHEF_360_FQDN>/organizations/<ORG_NAME>"
node_name = "<CHEF_360_ADMIN_USER>"
client_key = "~/.chef/chef360-admin.pem"

[legacy]
chef_server_url = "https://<CHEF_INFRA_SERVER_FQDN>/organizations/<ORG_NAME>"
node_name = "<INFRA_SERVER_ADMIN_USER>"
client_key = "~/.chef/infra-admin.pem"
# Cross-check cookbooks, roles, and environments
knife cookbook list --profile chef360
knife cookbook list --profile legacy

knife environment list --profile chef360 # Must include zone1, zone2, zone3
knife environment list --profile legacy

# Verify node count per zone in Chef 360 DSM
for ZONE in zone1 zone2 zone3; do
    echo "DSM ${ZONE}: $(knife search node "chef_environment:${ZONE}" -i --profile chef360 | wc -l) nodes"
done

```

Reference: [Connecting workstation to Chef 360 Platform](#)

11.3 Secure Post-Import Cleanup

```

# Chef Infra Server delete the pivotal key
rm -f /tmp/pivotal.pem

# Migration Workstation move backup to secure cold storage or delete
rm -f /migration/chef-server-backup.tar.gz
# Or: mv /migration/chef-server-backup.tar.gz /secure-cold-storage/

```

12. Step 7: Re-Point Zones to Chef 360

“To point nodes to Chef 360 Option A: Proxy | Option B: DNS | Option C: client.rb” Slides 5254

This is the core cutover step. Zone traffic is redirected from the Chef Infra Server to Chef 360 DSM. Execute **one zone at a time** complete Zone 1 before starting Zone 2.

Reference: [Redirect nodes to check in with Chef DSM](#)

Choose **one option** per zone:

Option A Add a Proxy Server to Redirect Traffic

“Add Proxy server to redirect traffic to Chef 360” Slide 52

Execution location: Proxy Server

Configure a reverse proxy to intercept HTTPS traffic for the Chef Infra Server FQDN and forward it to Chef 360 DSM. **Nodes require no changes.**

When to use: Nodes cannot be individually accessed; instant rollback needed; zone already uses a proxy.

```
# Proxy Server nginx example
upstream chef360_dsm {
    server <CHEF_360_FQDN>:443;
}

server {
    listen 443 ssl;
    server_name <CHEF_INFRA_SERVER_FQDN>; # Same FQDN nodes currently use

    ssl_certificate /etc/ssl/chef/server.crt;
    ssl_certificate_key /etc/ssl/chef/server.key;

    location / {
        proxy_pass https://chef360_dsm;
        proxy_ssl_verify off;
        proxy_set_header Host $host;
        proxy_set_header X-Real-IP $remote_addr;
        proxy_read_timeout 300s;
    }
}

sudo nginx -t && sudo systemctl reload nginx
```

Rollback: Update upstream chef360_dsm back to <CHEF_INFRA_SERVER_IP>:443 and reload nginx. Takes effect immediately.

Option B Change Existing DNS Record to Point to Chef 360

“Change existing DNS record to point to Chef 360” Slide 53

Execution location: DNS Provider / Internal DNS Server

Update the DNS A or CNAME record for the Chef Infra Server FQDN to resolve to the Chef 360 IP. **Nodes require no changes.**

When to use: DNS is centrally controlled; all zone nodes share the same FQDN in client.rb.

Pre-steps (2448 hours before migration day):

```
# Reduce DNS TTL for the Chef Infra Server FQDN to 60 seconds in your DNS provider
# Verify current TTL:
dig <CHEF_INFRA_SERVER_FQDN> | grep -i ttl
```

On migration day in your DNS provider:

```
Change A record: <CHEF_INFRA_SERVER_FQDN> <CHEF_360_IP>
# Migration Workstation verify propagation
dig +short <CHEF_INFRA_SERVER_FQDN>
# Must return: <CHEF_360_IP>

# On Node flush DNS cache and trigger chef-client run
sudo systemd-resolve --flush-caches # Linux (systemd)
# or: sudo nscd -i hosts # Linux (nscd)
# or: ipconfig /flushdns # Windows

sudo chef-client --once # Validate connection to Chef 360 DSM
```

Rollback: Revert DNS A record to <CHEF_INFRA_SERVER_IP>. Takes effect within TTL (60 seconds if pre-lowered).

Option C Change Node client.rb Configuration Directly

“Change node client config to point to Chef 360” Slide 54

Execution location: Migration Workstation pushed to Nodes via knife ssh/winrm

Update chef_server_url in each node's /etc/chef/client.rb (Linux) or C:\chef\client.rb (Windows).

When to use: DNS/proxy changes are not feasible; per-node control required.

```
# Migration Workstation Linux nodes (Zone 1 example)
knife ssh "chef_environment:zone1" \
  "sudo sed -i 's|chef_server_url.*|chef_server_url \"https://<CHEF_360_FQDN>/\
    organizations/<ORG_NAME>\"|g' /etc/chef/client.rb" \
  -x <SSH_USER> -i <SSH_KEY_PATH>

# Verify
knife ssh "chef_environment:zone1" "grep chef_server_url /etc/chef/client.rb" \
  -x <SSH_USER> -i <SSH_KEY_PATH>

# Trigger chef-client to validate new connection to Chef 360 DSM
knife ssh "chef_environment:zone1" "sudo chef-client --once" \
  -x <SSH_USER> -i <SSH_KEY_PATH>

# Migration Workstation Windows nodes
knife winrm "chef_environment:zone1" \
  "(Get-Content C:\chef\client.rb) -replace 'chef_server_url.*', 'chef_server_url\
    \"https://<CHEF_360_FQDN>/organizations/<ORG_NAME>\"' | Set-Content C:\
    \chef\client.rb" \
  -m -x <WINDOWS_USER> -P <WINDOWS_PASSWORD>
```

Rollback: Re-run the same sed/Set-Content command, replacing the Chef 360 FQDN with the original Chef Infra Server FQDN.

Re-Pointing Options Comparison

Status	Option A Proxy	Option B DNS	Option C client.rb
Node changes needed	None	None	Yes per node
Execution location	Proxy server	DNS provider	Workstation nodes
Rollback speed	Instant	~60 sec (pre-lowered TTL)	Minutes
Granularity	Per proxy rule	Per DNS record	Per node
Best for	Large fleets, proxied zones	Clean zone-wide switch	Per-node control
Risk	Low	Medium (DNS propagation lag)	Medium (config drift)

13. Step 8: Enroll Nodes into Chef 360 Node Management (Optional)

This step is OPTIONAL. It is required only if you plan to run Chef Courier jobs (remote execution, scheduled automation, InSpec compliance scans) on the migrated nodes. If you only need nodes to continue running chef-client runs via Chef 360 DSM (cookbooks, roles, environments), you may skip this step.

Complete Step 7 (Re-pointing) first, then return here when you are ready to enable Courier and full Node Management capabilities.

Execution location: Migration Workstation pushed to Nodes

Important distinction: Re-pointing (Step 7) redirects chef-client to Chef 360 DSM for cookbook runs. Node Management enrollment is an additional step that installs the Node Management Agent, Courier Runner, and other skills on nodes for full Chef 360 orchestration capabilities.

Follow the official Node Enrollment documentation for the complete steps: https://docs.chef.io/360/1.7/node_management/enroll_nodes/

Since all nodes already have Chef Infra Client installed, the recommended method is cookbook-based enrollment:

Reference: [Enroll nodes with a Chef Infra cookbook](#)

High-level steps:

1. **Generate a Self-Enrollment Key** Chef 360 UI **Node Management Self-Enrollment Keys Generate New Key Use as Cookbook**. Select the target zone cohort and copy the Signed Config File Contents.
2. **Create the enrollment cookbook** The Chef 360 UI provides step-by-step commands for generating the cookbook template, creating the enrollment script, and updating the default recipe. Follow those exactly.
3. **Upload the cookbook to Chef 360 DSM** after nodes are re-pointed (Step 7), they pull cookbooks from Chef 360 DSM:

```
# Migration Workstation
knife cookbook upload chef-360-enrollment --profile chef360
```

4. **Apply the cookbook per zone** trigger chef-client runs so nodes self-enroll:

```
# Migration Workstation Nodes
knife ssh "chef_environment:zone1" "sudo chef-client --once" \
-x <SSH_USER> -i <SSH_KEY_PATH>
```

5. **Verify enrollment status:**

```
# Migration Workstation
chef-node-management-cli management node find-one-node \
--nodeId <NODE_ID> \
--profile node-manager-profile
# Expect: "enrollmentLevel": "enrolled"
```

Reference: [Node Cohorts](#) for creating zone-specific cohorts **Reference:** [Skill Assemblies](#) for defining skills per zone

14. Step 9: Zone-by-Zone Parallel Migration Execution

"All 3 zones now pointing to Chef 360" Slide 55

Migrate one complete zone before starting the next. The Chef Infra Server remains live for unmigrated zones throughout.

14.1 Migration Sequence

Week 12 ZONE 1 (~2,0003,000 nodes)

1. Re-point to Chef 360 DSM (Option A, B, or C)
2. Verify chef-client check-ins in Chef 360 DSM
3. [Optional] Enroll nodes into Node Management (Step 8)
4. Observe for 72 hours no issues

Zone 1 COMPLETE proceed to Zone 2

Week 34 ZONE 2 (~2,0003,000 nodes)

(same steps Chef Infra Server still serves Zone 3)

Zone 2 COMPLETE proceed to Zone 3

Week 56 ZONE 3 (~2,0003,000 nodes most critical, migrate last)

(same steps Chef Infra Server now idle)

Zone 3 COMPLETE all zones on Chef 360 DSM

Week 7 PARALLEL VALIDATION WINDOW

Chef Infra Server running but idle (standby for rollback)

All nodes checking in with Chef 360 DSM 72h observation

Week 8 Chef Infra Server RETIRED (Step 10)

14.2 Per-Zone Validation

```
# Migration Workstation
# Node count in Chef 360 DSM must match pre-migration inventory
for ZONE in zone1 zone2 zone3; do
    LEGACY=$(knife search node "chef_environment:${ZONE}" -i --profile legacy | wc -l)
    DSM=$(knife search node "chef_environment:${ZONE}" -i --profile chef360 | wc -l)
    echo "Zone ${ZONE}: Legacy=${LEGACY}, Chef360 DSM=${DSM}"
done

# Confirm nodes converging (check ohai_time is recent)
knife search node "chef_environment:zone1" \
    -a ohai_time --profile chef360
```

14.3 During Parallel Operations

Activity	Action
New node provisioning	Enroll directly into Chef 360 DSM; do NOT add to Chef Infra Server
Cookbook updates	Deploy only via Chef 360 DSM after zone is re-pointed
Emergency rollback	See Section 17 revert re-pointing for affected zone
Chef Infra Server	Keep running but make no changes affecting already-migrated zones

15. Step 10: Retire the Chef Infra Server

“Retire Chef Infra Server” Slide 56

Execution location: Chef Infra Server + Migration Workstation

Execute only after all zones are validated and the 72-hour parallel observation window has passed.

15.1 Decommission

Follow your organization’s decommission process/policy to retire the Chef Infra Server. The steps below are provided as a technical reference your organization may have additional approval, change management, or compliance steps that must be followed before and after decommission.

Execution location: Chef Infra Server

```
# 1. Stop all Chef Infra Server services
sudo chef-server-ctl stop

# 2. Verify services stopped
sudo chef-server-ctl status

# 3. Revoke admin API access
sudo chef-server-ctl user-delete <ADMIN_USER>

# 4. Shutdown the server
sudo shutdown -h now
```

Reference: [chef-server-ctl commands](#)

Post-shutdown (follow your organization’s policy for each): – Update DNS (remove or permanently redirect old FQDN to Chef 360) – Update CMDB and internal documentation – Archive VM snapshot per your organization’s retention policy – Revoke any firewall rules or load balancer entries for the Chef Infra Server

TODO (Optional): The decommission steps above can optionally be automated via a cookbook or runbook. Consider codifying your organization’s decommission process as a cookbook for repeatability.

16. Post-Migration Validation Checklist

DSM Data Integrity

- All organizations from backup present in Chef 360 DSM
- Cookbook count and versions match pre-migration inventory
- Roles, environments (zone1 , zone2 , zone3), data bags present
- All 8,000 nodes visible in Chef 360 DSM
- Users are members of correct organizations

Node Check-ins

- All 8,000 nodes successfully checking in to Chef 360 DSM
- ohai_time updated within the last chef-client interval on all nodes
- No 401/403 authentication errors in chef-client logs

Node Management

The checks below apply only if you completed the optional Step 8 (Section 13 Enroll Nodes into Chef 360 Node Management).

- All nodes enrolled: "enrollmentLevel": "enrolled"
- Node Management Agent running; all skills installed per zone skill assemblies
- Node attributes visible in Chef 360 Admin Console

Courier Jobs

The checks below apply only if you completed the optional Step 8 (Section 13 Enroll Nodes into Chef 360 Node Management).

- Courier jobs executing on schedule without errors across all zones
- Job results and logs visible in Admin Console
- InSpec compliance scans running successfully

Security and Cleanup

- Chef Infra Server fully decommissioned
- Pivotal key deleted from all locations
- Backup tar file in secure cold storage or deleted
- Proxy decommissioned (if Option A)
- DNS TTL restored to normal values (if Option B)
- legacy knife profile retired from ~/.chef/credentials
- RBAC roles reviewed in Chef 360; excess accounts removed

17. Rollback Plan

17.1 Rollback Triggers

Trigger	Threshold
Node enrollment failures	> 10% of zone nodes fail
chef-client convergence failures	> 5% error rate post re-pointing
DSM import status	failure or partial-success with critical data missing
Production impact	Any outage directly attributable to migration

17.2 Rollback by Option

Option A Proxy Server:

```
# Revert upstream to Chef Infra Server instant, no node changes needed
# Update nginx upstream to <CHEF_INFRA_SERVER_IP>:443
sudo nginx -t && sudo systemctl reload nginx
```

Option B DNS Provider:

```
# Revert A record to <CHEF_INFRA_SERVER_IP>
# Propagates within ~60 seconds (if TTL was pre-lowered)
# On Node flush DNS cache
sudo systemd-resolve --flush-caches
```

Option C Migration Workstation Nodes:

```
knife ssh "chef_environment:<ZONE>" \
  "sudo sed -i 's|chef_server_url.*|chef_server_url \"https://<CHEF_INFRA_SERVER_FQDN>
organizations/<ORG>\"|g' /etc/chef/client.rb" \
  -x <SSH_USER> -i <SSH_KEY_PATH>
```

Note: Node enrollment in Chef 360 is not reversible (re-enrollment not supported). However, the Node Management Agent can be uninstalled from nodes without affecting Chef Infra Client, and Chef Infra Server can resume managing nodes independently.

Reference: [Chef 360 Node Enrollment](#)

18. Troubleshooting

Primary reference for DSM import issues: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/#troubleshooting

Issue	Cause	Fix
Import fails during extracting	Backup not in correct tar.gz format	Re-create using knife ec backup ; verify with file backup.tar.gz
Checksum mismatch on upload	SHA256 mismatch	Regenerate with chefimport-cli importservice generate-sha256
Organization already exists	Org name conflict in Chef 360	Rename or delete the existing Chef 360 org before re-importing
Import timeout at restoringdata	Large dataset; default timeout insufficient	Contact Progress Chef Support with Import ID and Tenant ID
401 Unauthorized after re-pointing	Node client key not in Chef 360 DSM	Re-register node or rebootstrap; delete and re-create client entry in DSM
DNS not propagating (Option B)	TTL not pre-lowered	Flush node DNS cache: sudo systemd-resolve --flush-caches
Proxy SSL error (Option A)	Certificate CN mismatch	Re-issue TLS cert for Chef 360 with old Chef Server FQDN as SAN
partial-success import status	Some objects failed to import	Review error report in notification email; manually re-import failed items

For support, collect the following and contact Progress Chef Support: - Import ID - Tenant ID - Import status JSON output - Error messages from logs - Backup file size

Reference: [Get support](#)

19. Reference: Key CLI Commands

chef-import-cli (Migration Workstation)

Full reference: https://docs.chef.io/360/1.7/reference/cli/chef_import_cli/

Command	Purpose
register-device	Register migration workstation with Chef 360
set-default-profile	Set active profile
importservice generate-sha256	Generate SHA256 checksum
importservice create-import-process	Any outage directly attributable to migration
importservice get-import-process-status	Monitor import progress

chef-node-management-cli (Migration Workstation)

Full reference: https://docs.chef.io/360/1.7/reference/cli/chef_node_management_cli/

Command	Purpose
management assembly create-assembly	Create skill assembly per zone
management setting create-setting	Create override settings
management cohort create-cohort	Create node cohort per zone
enrollment enroll-node	Enroll single node
importservice get-import-process-status	Enroll nodes from CSV
management node find-one-node	Check node enrollment status
status get-enrollmentId-status	Track bulk enrollment progress

knife (Migration Workstation)

Full reference: <https://docs.chef.io/workstation/latest/tools/knife/>

Command	Profile flag	Purpose
knife node list	--profile chef360 / --profile legacy	Compare node counts
knife cookbook list	--profile chef360	Verify cookbooks in DSM
knife environment list	--profile chef360	Verify zone environments
knife search node "chef_environment:zoneX"	both	Zone node counts
knife ssh "chef_environment:zoneX" ...		Push commands to zone nodes

chef-server-ctl (Chef Infra Server)

Full reference: https://docs.chef.io/server/ctl_chef_server/

Command	Purpose
chef-server-ctl stop	Stop all Chef Infra Server services
chef-server-ctl status	Check service status
chef-server-ctl user-delete <USER>	Revoke user access

20. Appendix

Appendix B: Zone Definition Summary

Attribute	Zone 1	Zone 2	Zone 3
Physical scope	DC/AZ/Region 1	DC/AZ/Region 2	DC/AZ/Region 3
Chef environment	zone1	zone2	zone3
Node count	~2,0003,000	~2,0003,000	~2,0003,000
Chef 360 cohort	zone1-cohort	zone2-cohort	zone3-cohort
Migration order	1st (lowest risk)	2nd	3rd (most critical)

Appendix C: Chef Infra Server Chef DSM Object Mapping

Reference: [Object mapping during migration](#)

Chef Infra Server	Chef 360 DSM Equivalent
Organizations	Tenant organizations (auto-created on import)
Users + membership	Users + organization membership
Cookbooks	Cookbooks in DSM
Roles	Roles in DSM
Environments	Environments in DSM
Data bags	Data bags in DSM
Nodes	Nodes in Node Management + DSM inventory
chef_server_url (client.rb)	https://<CHEF_360_FQDN>/organizations/<ORG>

Appendix D: All Scenarios Overview

Scenario	Description	Cluster Type
#1	Infra Server Feeding Automate	Parallel
#2	Infra Server Feeding Automate	Evolving
#3	Automate HA (11+ Node)	Parallel
#4	Automate HA (11+ Node)	Evolving
#5	Automate + Infra Server (Same Host)	Parallel
#6 This guide	Infra Server Only	Parallel
#7	Infra Server Only	Evolving
#8	5 Node Automate HA	Parallel
#9	5 Node Automate HA	Evolving

Document prepared based on:

Chef 360 Platform 1.7 documentation: <https://docs.chef.io/360/1.7/>

Chef DSM documentation: <https://docs.chef.io/360/1.7/dsm/>

DSM migration guide: https://docs.chef.io/360/1.7/dsm/migration_onprem_to_chef360/

Chef Infra Server documentation: <https://docs.chef.io/server/>

Chef 360 Upgrade Patterns presentation

Scenario #6, Slides 4956



Know More

About Progress Software

Progress Software (Nasdaq: PRGS) empowers organizations to achieve transformational success in the face of disruptive change. Our software enables our customers to develop, deploy and manage responsible AI-powered applications and personalized digital experiences with agility and ease. Businesses of all sizes get a trusted provider in Progress, with the products, expertise and vision they need to turn AI disruption into a competitive advantage. Millions of developers and technologists at hundreds of thousands of organizations depend on Progress every day. Learn more at www.progress.com

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates.
All rights reserved. Rev 2026/07 | RITM0377547



Worldwide Headquarters

Progress Software Corporation
15 Wayside Rd, Suite 400, Burlington, MA 01803, USA
Tel: +1-800-477-6473

facebook.com/progresssw
 x.com/progresssw
 youtube.com/progresssw
 linkedin.com/company/progress-software
 [progress_sw_](https://instagram.com/progress_sw_)