PH7X SYSTEMS ®

0.1.0 · CHAPTER 37

End hypercare on a criterion written before cutover, not on a date that arrives

SharePoint Migration Field Guide · Chapter 37

pH7x Systems

2026-11-10

37. End hypercare on a criterion written before cutover, not on a date that arrives

> The cutover proved one hour. Hypercare is the first full working week, the first scheduled crawl and the first month end, and it ends when somebody can show that they went well.

Chapter 35 moved write authority at one recorded instant and closed its cutover record. Chapter 36 wrote the way back and the date on which it closes. Chapter 38 switches the old farm off. This chapter owns the days between them: who is reachable, what is read on the new farm and how often, how a person's report becomes a finding with an owner, how the lines left open before cutover are closed, and what the old farm costs while it waits.

> **After cutover, who watches what, for how long, and what has to be true before the migration team steps back?**

The decision this chapter leads to is whether hypercare ends. The decider Chapter 1 named signs the exit against criteria written before cutover, and with it the date from which the old farm is Chapter 38's; or extends hypercare in writing, for a stated period, naming the lines still open.

Which periods this chapter is for, and what must already be true

First, and in most depth, the period after a SharePoint Server 2016 or 2019 farm was attached to SharePoint Server Subscription Edition (Chapter 23) and its address moved (Chapter 35). After a copy into SharePoint Online, Chapter 9 already holds hypercare, the rollback states and the gate for decommissioning; a short section near the end says what differs.

The chapter takes six things from the chapters before it, and repeats none of them:

- from Chapter 30, the baseline: the old farm's measures, taken over a stated window before its databases went read-only;
- from Chapter 32, every accepted finding and every line of the test plan still open, each with an owner and a date;
- from Chapter 33, the validators' confirmed hours, which already include hypercare, and the date of the next wave;
- from Chapter 34, the signed acceptance record, with each line left not established and its date inside hypercare;
- from Chapter 35, the cutover record: the instant, the runbook with its times, the readings before and after, both messages and the previous values;
- from Chapter 36, the trigger list and the expiry date of the way back.

It hands the exit record and the released old farm to Chapter 38, the readings and their new owners to Chapter 39, the date hypercare ended and the finding log to Chapter 40, the same log to Chapter 41, and the farm's build to Chapter 42.

What Microsoft documents, and what it leaves to you

We found no Microsoft page that defines a period of heightened support after an upgrade to Subscription Edition, or that says how long one lasts. Microsoft does not document one. The duration, the roster, the classes and the exit criteria in this chapter are our recommendation.

What Microsoft documents is what an administrator can read on the farm, and three points about when:

- The Subscription Edition upgrade pages end at validation. Their last timed instruction is for search: run queries, and "Twenty-four hours later, view the query reports and look for issues" (verify database upgrades, checked on 10 November 2026).
- The monitoring pages list Subscription Edition among the versions they apply to, although their text still names SharePoint Server 2016: the overview of monitoring, timer job status, search diagnostics, usage and health data collection, health reports and diagnostic logs, all checked on 10 November 2026.
- The one page that speaks of a farm's first days is written for SharePoint Server 2013 and lists no later version. It gives settings "that you can change to monitor your environment in its early stages", and adds: "Once the environment is stable and this detailed monitoring is no longer required, you may want to reverse the settings below to their default settings" (monitoring and maintaining SharePoint Server 2013, checked on 10 November 2026). We use it for method and for the names of counters, never as a Subscription Edition fact.

One limit applies to every script below. The reference pages for `Get-SPContentDatabase`, `Get-SPFarm`, `Get-SPTimerJob`, `Get-SPLogEvent`, `Merge-SPLogFile`, `Get-SPUsageDefinition`, `Get-SPEnterpriseSearchServiceApplication` and `Get-SPEnterpriseSearchStatus` list their parameters as applicable to SharePoint Server 2010 or 2013 up to 2019, and stop before Subscription Edition (each is linked where it is used; all checked on 10 November 2026). Microsoft's Subscription Edition upgrade procedure itself runs `Get-SPContentDatabase`, as Chapter 23 records. For the others, Chapter 31's rehearsal farm is where you confirm that they answer. No farm was available to us, so every script here is illustrative and was not run.

Before cutover: the hypercare plan

Hypercare starts at Chapter 35's instant, so its plan is signed before Chapter 35's decision. Chapter 35's second message already tells people "who to call, and until when the migration team is on hand": this plan is where that date comes from. One page per cutover. Print it; the last column is for a signature.

LINE	WHAT IT HOLDS	OWNER	SIGNED
Period and exit	The first day, the planned last day, a ceiling date, and the exit criteria below	Decider	
Roster	Who is reachable in which hours, with a deputy for each role, and how each is reached	Hypercare lead	
Report route	The one place a person reports to, the questions the helpdesk asks, and the classes a report can take	Helpdesk lead	
Finding log	Where it lives, its fields, and who may close a finding	Hypercare lead	
Watch list	Each reading below, with its expected result, how often it is read and by whom	Farm administrator	
Open lines	Every line Chapters 32 and 34 left not established, and every accepted finding with a fix date	Each line's owner	
Way back	Chapter 36's triggers and expiry date, who reads the old farm's state, and what holding it costs	Rollback plan's owner	
Monitoring settings	Every setting raised for hypercare, beside the value it returns to	Farm administrator	
Backups	The new farm's backup schedule from the instant, and the date of its first restore test	Database administrator	
Hand-over	What operations receives at the exit, and who accepts it	Operations owner	

Three things have to be staffed, not only written. The validators Chapter 34 named keep hours for hypercare, which Chapter 33 already counted. The owner of each integration's other end (Chapter 22) is reachable for its first scheduled run. And the people who can carry out Chapter 36's return stay reachable until its expiry date: a way back with nobody to walk it is not held.

How long it lasts, and what ends it

Chapter 9 gives a copy wave "days for a small wave and two weeks for a large one". An attach moves every site behind an address at once (Chapter 33), so we would plan no less than two working weeks for it, and we would write three dates: the first day, the planned last day, and a ceiling at which the decider either signs the exit or extends in writing. A period with no ceiling becomes the permanent arrangement, paid for by a team that has another wave to run.

None of the three dates ends hypercare. The criteria do. We recommend these, written before cutover with a figure of your own wherever a row asks for one:

EXIT CRITERION	MET WHEN	EVIDENCE	SIGNED BY
Open lines	Every line Chapters 32 and 34 left not established is established, a finding, or handed over with an owner and a date	The two records, updated	Each line's owner
Findings that stop work	None is open	The finding log	Hypercare lead
Other findings	Each is fixed and confirmed by its reporter, accepted in writing by the owner of what it affects, or handed over with an owner and a date	The finding log	Each owner
New reports	Reports per working day, counted per class and never in total, are at or below the level the helpdesk lead wrote, for the number of consecutive working days the plan wrote	The log's daily counts	Helpdesk lead
The watch list	Every reading gave its expected result, or its exception is a finding, on that same run of days	The readings, kept	Farm administrator
Performance	Each of Chapter 30's measures, read at a real peak, is within the limit written beside the baseline	Both sets of figures	Infrastructure owner
The business cycle	Each scheduled event that fell inside the period ran: a month end, a weekly report, a workflow with a long wait. One that fell outside is handed over by name	Its owner's confirmation	Process owner
Backups	The new farm's databases were backed up from the instant and one backup was restored elsewhere	The restore's output	Database administrator
The way back	Chapter 36's expiry date has passed, unused	The rollback plan	Rollback plan's owner
Hand-over	Operations accepted the farm, its readings and its open findings in writing	The hand-over record	Operations owner

The fourth row is not a score. It is a count of one kind of thing per day, read beside the class it belongs to, as Chapter 40 will read tickets.

We recommend that the expiry of the way back falls no later than the planned last day. If hypercare ended first, the people who could run a return would have dispersed while the plan still promised one.

Who is on call

ROLE	WHAT THEY DO IN HYPERCARE	REACHABLE
Hypercare lead	Owns the finding log, runs the daily review, calls for the decider when a trigger is met	The whole period, with a deputy
Helpdesk	Receives every report, asks the plan's questions, opens the finding	Its normal hours, briefed before cutover
Farm administrator	Reads the watch list, investigates platform findings	The whole period, with a deputy
Database administrator	Backups of the new farm, the state of the old farm's databases	The whole period
Search owner	The crawl log, the query reports, Chapter 11's tests when a result is disputed	The first week, then on request
Validators	Confirm or refute a report about their own sites, and confirm each fix	The hours Chapter 33 recorded
Owner of each integration's other end	Confirms the first scheduled run, and each run after a fix	At each scheduled run
Decider, and deputy	Takes Chapter 36's decision if a trigger is met, and signs the exit	Until the expiry date, then at the exit

The migration lead does not close findings about content. A validator does, for the reason Chapter 34 gives: the migration team signs the verification, never the acceptance.

From a report to a finding

Give people one place to report, and make it the place they already use. A report that reaches a team member in a corridor reaches the log the same day or it did not happen.

Each report becomes one finding with these fields: a number; the time it arrived; who reported it; the site and the address; what the person did, expected and saw, in their words; any identifier the error showed; its class; whether it was read against Chapter 36's triggers, by whom; its owner; its state; the evidence; and who confirmed the close.

We recommend Chapter 40's four classes, so that its ticket counts continue from yours without translation: missing content, access, how-to and platform. Add a flag, not a fifth class, for a dependency found late: a job, a report or a system that nobody listed and that failed or appeared after cutover. Chapters 38 and 41 both read that flag.

Triage in this order:

1. **Is it a trigger?** Read the finding against Chapter 36's trigger list before anything else. One that meets a trigger goes to the decider; the rest are fixed forward.

2. **Which farm did the person reach?** A read-only statement at the top of the page means the old farm, through a cached name (Chapter 35). It is a how-to finding, and nothing is lost.
3. **Is it already known?** The accepted findings of Chapters 32 and 34 answer many reports. Reply from the record, with the date of the fix if there is one.
4. **What do the records say?** For missing content, Chapter 8's read of the source says whether the item existed. The old farm is still there, read-only, and shows what a page, a view or a permission was before. This is the first thing the old farm is held for.
5. **What do the logs say?** See below.

A finding has five states: new; owned; fixed and confirmed by the person who reported it; accepted in writing by the owner of what it affects; or handed over with an owner and a date. A fix repeats the acceptance line it touches and every line that depends on it, as Chapter 34 requires.

When a report carries a correlation ID, `Merge-SPLogFile` "Combines trace log entries from all farm computers into a single log file on the local computer", and its `-Correlation` parameter "Specifies the correlation ID to filter on" (Merge-SPLogFile, checked on 10 November 2026). Microsoft recommends filtering by time as well; without `-StartTime` the default is "one hour prior to the current time on the local computer".

```
# Illustrative, not run against a farm. New farm (Subscription Edition), SharePoint Management She
# Reads the trace logs of every farm server and writes one new file. It changes nothing on the far
$id   = "<correlation ID from the finding>"
$from = "<start of the window, in the server's date format>"
$to   = "<end of the window>"
$out  = "<folder>\finding-<number>.log"
Write-Host "Merging trace log entries for $id between $from and $to into $out"
Merge-SPLogFile -Path $out -Correlation $id -StartTime $from -EndTime $to
```

Expected result. A file with the entries of that one request, from every server that handled it; if nothing matches, "a warning is written to the Windows PowerShell console window". **Validated by** attaching the file to the finding. For one server and one window, `Get-SPLogEvent -MinimumLevel "Warning" -StartTime <start> -EndTime <end>` returns the entries at or above that level, the valid values being "Error or Warning" (Get-SPLogEvent, checked on 10 November 2026).

What to watch on a Subscription Edition farm

READING	WHERE	EXPECTED	HOW OFTEN	VALIDATED BY
Health Analyzer	Central Administration, Monitoring, Review problems and solutions	No alert that was not on the list read before cutover	Daily	Each new alert is a finding, or is explained in the log
Database and site upgrade	Check upgrade status; `NeedsUpgradeIncludeChildren`; the upgrade logs	Unchanged since Chapter 23's readings	First hour, then after any change	Compared with Chapter 35's record
Timer jobs	Monitoring, Check job status; `Get-SPTimerJob`	No failed job without a finding; every timer service instance online	Daily	The job history, read for each kept customisation's job
Search components	`Get-SPEnterpriseSearchStatus`	Every component Active	First hour, then daily	Read twice before recording Degraded
Crawl	The crawl log	A crawl ran on schedule; no top-level error; an item created after the instant is found	Daily	The search owner, with the test item
Queries	Query health reports; usage reports	Latency within the search owner's limit; no known query newly without results	A day after the instant, then weekly	Compared with Chapter 11's exports
Usage and health data	Configure usage and health data collection; View health reports	Collection is on; the reports hold data	First day, then weekly	The slowest pages, read by the infrastructure owner
Performance	Chapter 30's measures, and the servers' counters	Within the limit beside the baseline, at a real peak	At each peak	Both sets of figures, side by side
Diagnostic logs	The trace log folder; its settings	Logs kept for longer than hypercare; the drive not filling	First day, then weekly	The setting and the free space, recorded

Sources for the table are the pages cited in the sections that follow.

The first reads

```
# Illustrative, not run against a farm. New farm (Subscription Edition), SharePoint Management She
# Reads only: it changes nothing.
Write-Host "Reading database upgrade state, timer service instances that are not online, and searc
Get-SPContentDatabase | ft Name, NeedsUpgradeIncludeChildren
$farm = Get-SPFarm
$farm.TimerService.Instances | where {$_.Status -ne "Online"}
Get-SPEnterpriseSearchServiceApplication | Get-SPEnterpriseSearchStatus -Text
```

Expected result. Every database false. Nothing from the third command. Every search component Active.

The first line is Microsoft's own, from the Subscription Edition procedure: it "returns a table-style list of databases in your farm and indicates whether the database needs an upgrade" (upgrade content databases; reference: Get-SPContentDatabase). The two lines about the timer service come from a Microsoft troubleshooting article: "typically after an upgrade", the timer service can be running on a server while its instance object is not online, and then administrative operations that depend on one-time timer jobs, "such as creating or deleting a Search Service Application, or deploying farm solutions", do not succeed (administrative timer jobs don't run after an upgrade; reference: Get-SPFarm). That article names no version and carries a repair script; we use only its read, and a result here is a finding to investigate before anybody runs the repair. The last line is the first example on the reference for Get-SPEnterpriseSearchStatus, which gives each component one of four states, Active, Degraded, Failed or Unknown, and says of Degraded that the reason "is typically a transient situation related to a restart or network issues" (reference for the first cmdlet in that line: Get-SPEnterpriseSearchServiceApplication). All six pages checked on 10 November 2026.

Validated by comparing the output with the same reads in Chapter 35's cutover record.

Health Analyzer

Health Analyzer "runs predefined health rules against servers in the farm". "When any rule fails, SharePoint Health Analyzer creates an alert on the Review problems and solutions page and writes status to the Windows Event log", and each rule belongs to one of four categories: "Security, Performance, Configuration, or Availability" (overview of monitoring, above). A rule runs "hourly, daily, weekly, monthly, or on demand only", and a rule set to repair automatically "will repair errors as soon as they are found" (configure Health Analyzer rules, checked on 10 November 2026).

Three consequences for hypercare, all our recommendation:

- Read the page before cutover and keep the list. A new farm may show alerts before anybody uses it; the ones hypercare cares about are the ones that arrived with real load.
- Two weeks do not contain every monthly rule's run. Before the exit, we would run the monthly rules on demand, which the overview says a farm administrator can do, and read the page once more.

- Do not let hypercare change a rule's repair setting without recording it. A repair that runs by itself changes the farm between two of your readings.

A page with no alerts shows that no rule failed. It does not show that people can work, and it is no substitute for the finding log.

Upgrade status and logs

The Upgrade Status page "lists all farm, service, or content database upgrades and their statuses", with "a count of errors or warnings", and the logs sit under `%COMMONPROGRAMFILES%\Microsoft Shared\Web server extensions\16\LOGS` (verify database upgrades, above). Chapter 23 read all of this before anybody used the farm, so in hypercare it changes only if something was left to happen later.

One thing can be. If a database was attached with the `SkipSiteUpgrade` parameter, each site collection "is upgraded when first browsed": "the site will be put in a queue and a timer job will pick it up for upgrade" (upgrade a site collection, checked on 10 November 2026). Those upgrades then happen during hypercare, one first visit at a time, and each site's administrator can read the result on the Upgrade Status page in Site Settings. Where Chapter 23's attach did not use the parameter, the site collections were upgraded with their database and nothing is pending.

Timer jobs

"Many features in SharePoint Server 2016 rely on timer jobs to run services according to a schedule", and Central Administration shows "currently scheduled jobs, failed jobs, currently running jobs, and a complete timer job history" (overview of monitoring, above). Read the failed jobs there. For one job's history, Microsoft gives the PowerShell:

```
# Illustrative, not run against a farm. New farm (Subscription Edition). Reads only.
$job = "<name of the timer job>"
Write-Host "Reading the last run of every timer job, then the history of $job. Times are GMT."
Get-SPTimerJob | Format-Table DisplayName,Id,LastRunTime,Status
(Get-SPTimerJob -Identity $job).HistoryEntries | Format-Table -Property Status,StartTime,EndTime,E
```

Both commands are on Microsoft's page, which adds that without `-Identity` "all timer jobs are returned", and warns that the cmdlet "will show you logs in GMT time zone whereas SharePoint Central Administration will show all events in local time" (timer job status, above; reference: Get-SPTimerJob, checked on 10 November 2026).

Expected result. Every job a kept customisation installed (Chapter 15) has run since the instant, on its schedule, with no error message. **Validated by** the owner of each such job: a timer job leaves no trace in the sites that depend on it, as Chapter 15 found, so nobody reports its silence until the thing it produced is missed.

Search: the crawl, then the queries

Chapter 34 waited for the full crawl. Hypercare watches what comes after it. An incremental crawl "only runs if it is scheduled", and the crawl log shows, per content source, "successes, warnings, errors, top-level errors, and deleted items", with the average duration of a crawl "for the last crawl, for the last 24 hours, for the last 7 days and for the last 30 days". A top-level error is one "in top-level documents, including start addresses, virtual servers, and content databases" (search diagnostics, above). So the daily read is: did a crawl run, did it finish, is there a top-level error, and is its duration growing.

The proof is an item. The search owner creates a test document after the instant, lists it in the record, and searches for it after the next crawl, as a person who may open it and as one who may not.

For queries, the Query Latency Trend report "shows the query latency (in milliseconds) by percentile" with the query rate over it, and the usage reports include "No Result Queries", the "popular search queries that returned no results" (search diagnostics, above). Read them a day after the instant, as Microsoft's checklist asks, and compare the queries without results with Chapter 11's exports from the old farm. A query that found something there and finds nothing here is a finding, as Chapter 40 says, whatever the cause turns out to be.

Usage and health data collection

An attach brings content, not configuration (Chapter 23), so this is set on the new farm or it is off. The settings "are farm-wide and cannot be set for individual servers in the farm", and Microsoft warns: "Logging uses system resources and can affect performance and disk usage. Only log those events for which you want regular reports" (usage and health data collection, above). Health reports are "pre-built reports" drawn from that usage data, "slowest pages and top active users" among them, and they exist only in Central Administration: "You cannot use Microsoft PowerShell to view health report data" (health reports, above).

```
# Illustrative, not run against a farm. New farm (Subscription Edition). Reads only.
Write-Host "Reading the usage definitions of the farm, then the one for page requests"
Get-SPUsageDefinition
Get-SPUsageDefinition -Identity "Page Requests"
```

Without `-Identity` the cmdlet "returns the collection of usage definitions in the farm", and the second line is the example on its reference page (Get-SPUsageDefinition, checked on 10 November 2026). Microsoft's page does not list the properties of what it returns, so read the output on the rehearsal farm before you rely on a column.

Expected result. The events you chose are collected, page requests among them, and the health reports show data after the first import. **Validated by** opening the slowest pages report. Usage data is kept for at most 31 days, the range of `DaysRetained` being "0 to 31 days" (usage and health data collection, above): export what Chapter 40 will want before it ages out.

Performance against the baseline

Chapter 34 left this line open because only real load answers it, and Chapter 32 handed over the measures. Read Chapter 30's measures on the new farm the same way and over the same kind of window as on the old one, at a real peak, and compare measure by measure with the limit the infrastructure owner wrote beside each. That comparison is the test. The counters below only help explain a measure that fails.

```
# Illustrative, not run against a farm. Any machine that can reach the web servers. Reads only.
$servers  = "<WebServer1>", "<WebServer2>"
$counters = "\Processor(_Total)\% Processor Time", "\Memory\Available MBytes", "\ASP.NET\Requests
Write-Host "Sampling $counters on $servers every 15 seconds, 20 times"
Get-Counter -Counter $counters -ComputerName $servers -SampleInterval 15 -MaxSamples 20
```

`Get-Counter` "gets performance counter data directly from the performance monitoring instrumentation in the Windows family of operating systems", and its reference notes that counter names differ on a system in another language: `(Get-Counter -ListSet Memory).Paths` lists the names your servers use (Get-Counter, checked on 10 November 2026). The three counters are from the 2013 page's list, which says of processor use that it should not "remain consistently high (over 80 percent)", and treats any rejected request as a problem. The same page says where to look first: "Often, the reason for the bottleneck will be inefficient custom code or third-party solutions", and "there is usually a gradual performance degradation that you can track if you monitor regularly" (monitoring and maintaining SharePoint Server 2013, above). We do not know of a Microsoft page that restates those thresholds for Subscription Edition. The limit that decides this line is the one written beside your own baseline.

Validated by the infrastructure owner, who signs the open line as established, or records a finding.

Diagnostic logs and backups

The trace logs are what every investigation above reads, and they delete themselves. Logs are kept for a number of days you set, from 1 to 366; "By default, the amount of disk space that diagnostic logging can use is unlimited"; and verbose logging "can quickly use drive space and affect drive and server performance" (configure diagnostic logging, checked on 10 November 2026). We would keep logs for longer than the ceiling date, restrict their disk space, and raise the level only for a named finding and for a stated time.

From the instant, the backups taken before the upgrade no longer hold what people write. The new farm needs its own, from the first hour. Microsoft's guidance, on a page that lists Subscription Edition, is: "Routinely test backups and validate their consistency", and it notes that SharePoint's backup "doesn't back up the Unified Logging Service (ULS) trace logs" (backup best practices, checked on 10 November 2026). Chapter 32's restore script applies unchanged to one of the new farm's backups.

The routine

WHEN	WHAT YOU DO	WHO	EXPECTED RESULT
First hour	Receive Chapter 35's cutover record; open the finding log with the instant's time	Hypercare lead	The log exists, and the record is attached
First hour	Call the roster; confirm the helpdesk has the questions and the classes	Hypercare lead	Every role answers, or its deputy
First hour	Run the first reads; read Health Analyzer against the list kept before cutover	Farm administrator	As above; every new alert recorded
First hour	Confirm the first backup of the new farm's databases is scheduled or running	Database administrator	A time, written
First day	Review new findings at fixed times, each read against the triggers first	Hypercare lead	No finding older than the plan's time without an owner
First day	Read failed timer jobs and the history of each kept customisation's job	Farm administrator	Each ran, or is a finding
First day	Read the crawl log; search for the test item	Search owner	A crawl ran; the item is found, trimmed
First day	Confirm each integration's first scheduled run with the owner of its other end	Each owner	Chapter 22's four results, on real traffic
First day	Read Chapter 30's measures at the day's peak	Infrastructure owner	Within the limit, or a finding
First day	Count the requests still reaching the old farm, by Chapter 22's method	Farm administrator, old farm	Falling; each caller that is not a person is flagged
A day after the instant	Read the query reports	Search owner	No known query newly without results
First week, daily	The watch list; the review; the daily count of reports per class	Roster	Each reading recorded, each exception a finding
First week	Read the measures on the first full working day and at the week's peak	Infrastructure owner	The open line closed or a finding
First week	Close the open lines whose dates fall in the week	Each line's owner	Signed, in Chapter 32's and Chapter 34's records
First week	Read the old farm: its content databases still read-only, with Chapter 35's query; its backups in reserve	Database administrator	`is_read_only` is 1; nothing changed
First week	Review the exit criteria, row by row	Decider, hypercare lead	Which rows are met, and what the others need
Exit	Run the monthly health rules on demand; read the whole watch list once more	Farm administrator	Clean, or findings

WHEN	WHAT YOU DO	WHO	EXPECTED RESULT
Exit	Return every raised monitoring setting to the value in the plan	Farm administrator	The settings match the plan's second column
Exit	Read the exit criteria line by line; sign, or extend in writing	Decider	The exit record
Exit	Hand over to operations; tell people where reports go from now on	Hypercare lead	Accepted in writing; message sent

The times of the daily review, the age at which an unowned finding is escalated and the number of consecutive days are yours. Write them in the plan before cutover. We give none, because none would be reproducible on your estate.

Close the lines Chapters 32 and 34 left open

Both chapters allowed one kind of line to cross the cutover: a line that could not be read before people arrived, with an owner and a date inside hypercare. Performance under real load is the usual one; the query reports of real traffic and a full cycle of use are the others Chapter 32 names.

Each ends, on or before its date, in one of Chapter 30's three states. Established: the reading, its window and its evidence, signed by the line's owner in the record that opened it. A finding: fixed and retested, or accepted in writing by the owner of what it affects. Or still not established, which hypercare may not leave as it is: the line gets a new owner in operations and a new date, and Chapter 40 reads it. A line that needs a month end when hypercare lasts two weeks is the plain case, and the plan says so before cutover.

The accepted findings travel the same way. One that Chapter 34 accepted "with a date for its index" has that date in the finding log from the first hour, with its owner, so that the acceptance does not become the permanent state by default.

Holding the way back, and what it costs

Until Chapter 36's expiry date, the SharePoint Server 2016 or 2019 farm stays as Chapter 35 left it: running, read-only and untouched. Microsoft's Subscription Edition procedure says the backups taken before the upgrade guarantee "that you have the data in a safe state if you must enable the original farm again", and asks you to "keep a copy of your original backups in reserve" (copy databases to the new farm, checked on 10 November 2026). It gives no period for which to keep either. The date is yours, and Chapter 36 wrote it.

Holding it has a cost, and we recommend writing it into the plan so that the expiry date is a decision and not a habit:

- **An unsupported farm stays on the network.** SharePoint Server 2016 and 2019 reached the end of support on 14 July 2026, after which there are "no new security updates, non-security updates,

free or paid assisted support options or online technical content updates" (ending support in 2026, checked on 10 November 2026). Every day of hold is a day of that exposure.

- **Its servers, its SQL Server and its backups keep running**, with whatever they cost you. Prices and licences are outside this book.
- **Its people stay reachable.** The farm administrator, the database administrator and the network owner who rehearsed the return cannot be released to the next wave until the date passes.
- **It can be changed by accident.** Its own Health Analyzer may raise the rule "Databases within this farm are set to read only and will fail to upgrade unless it's set to a read-write state", whose resolution reads "Set the databases to read-write using SQL Server" (the rule's page, checked on 10 November 2026). On a held farm that alert is expected, and following its resolution ends the hold. Tell whoever still administers the old farm, in writing.
- **It is worth less every day.** After the first business write, a return re-enters or loses everything written since the instant (Chapter 36). By the second week the way back mostly protects against one case: a fault in the upgrade itself that nobody can fix forward.

The old farm does a second job while it waits: it is the witness for step 4 of the triage. That job does not need the rollback option, so it does not justify extending the expiry. It ends when Chapter 9's evidence gate is met.

When the date passes, the rollback plan's owner records that it passed unused, and the old farm becomes a candidate for Chapter 38. Extending the date is Chapter 36's decision, taken by its owner with the cost above in front of them, never by letting the day go by.

Common failures

WHAT YOU SEE	LIKELY CAUSE	READ	THEN
"My document is gone"	It never existed, the person lacks access, or search has not caught up	Chapter 8's read; the old farm, read-only; the crawl log	Answer from the record. A real difference from the old farm is a trigger candidate
New documents are not found	No crawl has run since the full one	The crawl log's history	Schedule the crawl; an incremental one "only runs if it is scheduled"
Alerts, expiry or a scheduled report stopped	A timer job failed, or a timer service instance is not online	Failed jobs; the first reads	A finding with the job's owner. Investigate before running Microsoft's repair script
A page is slow for everybody	Custom code, a large list, or load the hardware does not carry	The slowest pages report; the counters; the view's size	Chapter 34 lists throttled views as a known kind of finding; the rest is the performance line
A site looks different on first visit	Its upgrade was deferred with `SkipSiteUpgrade`	The site's Upgrade Status page	Wait for the timer job; a failure is a finding with the log
A person sees the read-only statement	Their machine still resolves the old farm	Chapter 35's checks	Nothing is lost. Count it as how-to
A feature the old farm ran does nothing	One of the switches Microsoft's 2026 updates turned off	Chapter 34's list	Its owner decides, as Chapter 34 sets out
The health reports are empty	Usage data collection is off on the new farm, or the import has not run	The collection settings; the import job's history	Enable what you chose; read again after the next import
The log drive is filling	Verbose logging left on after an investigation	The diagnostic logging settings	Return the level to the plan's value
The old farm shows a health alert about read-only databases	Expected on a held farm	The rule's page	Do not follow its resolution before the expiry
A job nobody listed fails at month end	A dependency found late	The old farm's request logs; the other system's log	Flag it. Chapter 22 owns the fix, Chapters 38 and 41 read the flag

Sources as cited in the sections above, all checked on 10 November 2026.

Risks and limits

- **The period is ours.** Microsoft documents readings, not a hypercare. Two weeks is a recommendation, and the criteria are what end it.
- **The guidance is older than the product.** The monitoring pages name 2016, the counters page is for 2013 only, and the cmdlet references stop at 2019. Where your rehearsal farm disagrees with

a page, the farm is right about your estate and the difference is a finding.

- **Watching costs the thing watched.** Usage logging and verbose trace logs use processor and disk. Raise them for a reason, write the normal value beside each, and put them back.
- **A quiet log proves little.** People stop reporting when a report gets no answer. Read the log's silence beside the helpdesk's own count and the validators' word.
- **Some cycles outlast any hypercare.** A year end cannot be inside a two-week period. Hand it over by name.
- **Validators have day jobs.** If their hours were spent on acceptance, hypercare has none. Chapter 33 counted them, and the plan confirms them again.
- **We do not know how your customisations behave under real load.** No test farm shows it (Chapter 32). That is what the first week reads.

A copy into SharePoint Online: what differs

Chapter 9 holds this period: its length, triage against the verification record, the five rollback states and the gate for decommissioning. The plan, the finding log, the classes and the exit criteria of this chapter apply unchanged. Six things differ:

- **There is no farm to watch.** Microsoft operates the service, so the watch list above has no equivalent. What you read is the Service health page of the Microsoft 365 admin centre, where incidents and advisories appear and where "People who are assigned the Service Support admin and Helpdesk admin role can view service health". It keeps "the history of any service advisories and incidents that have affected your tenant in the past 30 days", and planned maintenance is not shown there (service health, checked on 10 November 2026). Record it during hypercare, so that a report raised during a Microsoft incident is classed as platform.
- **The unit is a site.** One site's finding holds that site, and each wave has its own hypercare, taken out of the same validators' hours as the next wave's acceptance (Chapter 33).
- **Missing content has one more place to look.** After Chapter 8's record, the recycle bin: a deleted item in SharePoint in Microsoft 365 is kept for 93 days from its deletion (restore items in the recycle bin, checked on 10 November 2026). An item deleted after cutover is not a migration finding.
- **Performance and search are read as Chapter 34 read them**: timed tasks from the same desks, and content that appears on the tenant's crawl schedule. There is no baseline of server measures to repeat.
- **The way back is the locked source sites**, still on a SharePoint Server 2016 or 2019 farm past its end of support, and the cost of holding them is the same list as above.
- **Reports that arrive late matter more.** Migration Manager's file share reports expire, as Chapter 9 records, so the archive is filled during hypercare, not after it.

The exit, and what it hands on

The decider reads the exit criteria line by line with the people who signed them, never a count of rows met, and signs one of two things. The exit: every row met, with the evidence named. Or an extension: a new planned last day, a new ceiling, and the rows still open with their owners. A third outcome does not exist. A hypercare that nobody ended has become operations without a hand-over.

At the exit, these hold, and the chapters after this one begin from them:

- **Chapter 38** receives the exit record, the expiry date passed unused, an old farm nobody changed, and the flagged dependencies found late, each of which is a caller its silent window will otherwise meet again.
- **Chapter 39** receives every reading of the watch list as a control with an owner in operations, and the monitoring settings back at their normal values.
- **Chapter 40** receives the date hypercare ended, from which its windows start, the finding log counted per class, and each line still not established with its owner and date.
- **Chapter 41** receives the finding log whole: every finding is a decision that the evidence later tested.
- **Chapter 42** receives the build the farm ran through hypercare. We would install no public update on the new farm during the period, unless it repairs a finding or its notes make it urgent, so that a change in behaviour has one possible cause; Chapter 42's calendar starts at the exit.

Hypercare does not end because two weeks passed. It ends when somebody signs, line by line, that the farm no longer needs the people who built it.

Out of scope

The cutover is Chapter 35 and the return Chapter 36; this chapter only reads incidents against Chapter 36's triggers. Switching off the old farm is Chapter 38, the controls that stay Chapter 39, measuring what changed Chapter 40, lessons Chapter 41 and public updates Chapter 42. The tests themselves stay with Chapters 8, 11, 22, 32 and 34. Monitoring products, System Center management packs, alerting design, service desk tooling, service level agreements, capacity planning for growth, costs and licences are not covered.

Before you continue

- [ ] The hypercare plan, with its exit criteria, roster, report route and ceiling date, was signed before the cutover decision
- [ ] Every report became a finding with a class, an owner and a state, read against Chapter 36's triggers before anything else
- [ ] The first hour, the first day and the first week were each read against expected results written beforehand, and the readings were kept
- [ ] Health Analyzer, upgrade status, timer jobs, the crawl log and the query reports were read on the new farm, and every exception is a finding
- [ ] Chapter 30's measures were read at a real peak, and every line Chapters 32 and 34 left open is established, a finding, or handed over with an owner and a date
- [ ] The old farm stayed read-only and untouched until the expiry date, and what holding it cost was written in the plan
- [ ] Every raised monitoring setting is back at its normal value, and operations accepted the farm, its readings and its open findings in writing
- [ ] The decider signed the exit on the criteria, line by line, and Chapter 38 holds the record