



Fabric Capacity Focus: Surge Protection, Notifications, and Capacity Overage

 **SYSM-358** - Jira project doesn't exist or you don't have permission to view it.

 This page focuses on three Fabric capacity options that can improve operational resilience:

- **Surge Protection**
- **Notifications / Alerts**
- **Capacity Overage**

For our platform, these options should be treated as **operational controls**, not as substitutes for proper capacity isolation and sizing. Since throttling is applied at the **capacity level**, the first protection for the **Data Platform Core workspace (prod Capacity only here)** remains capacity separation from Domain workloads

Version	Date	Description	Contributor
V0.1	15 Apr 2026	Initial document	COLOMBANI Théo
V0.1	17 Apr 2026	Update document	COLOMBANI Théo

- [Key message](#)
 - [What we should implement](#)
 - [Decision guide](#)
 - [Quick checklist](#)
- [1. Surge Protection](#)
- [2. Notifications / Alerts](#)
- [3. Capacity Overage](#)

Key message

These three options should be used as **operational safeguards**, not as substitutes for proper capacity design.


For our platform, the priority remains:

- protect the **Data Platform Core** workspace through **capacity isolation**
- use **Surge Protection** mainly to control variable Domain workloads
- make **Notifications** mandatory for production operations
- use **Capacity Overage** only as a controlled safety net for rare peaks. (learn.microsoft.com)

What we should implement

Feature	Data Platform Core capacity	Domain production capacity	Clear recommendation
Surge Protection	Secondary control only	Yes	Enable mainly on Domain capacity
Notifications / Alerts	Mandatory	Mandatory	Define owner, recipients, thresholds, escalation
Capacity Overage	Optional, capped	Optional, tightly capped	Use only for rare peaks and with budget approval

Recommended setup



For Data Platform Core

- keep it on a **dedicated capacity**
- make **alerts mandatory**
- use **Capacity Overage** only if uptime is critical and cost is approved
- do **not** rely on Surge Protection as the main protection layer. (learn.microsoft.com)

For Domain production

- allow shared capacity if needed
- enable **Surge Protection**
- enable **workspace-level controls**
- make **alerts mandatory**
- use **Capacity Overage** only with strict limits. (learn.microsoft.com)

Decision guide

Enable Surge Protection when	Make Notifications mandatory when	Enable Capacity Overage when
<ul style="list-style-type: none">• the capacity is shared• Domain workloads are variable• noisy workspaces need runtime control. (learn.microsoft.com)	<ul style="list-style-type: none">• the capacity is production• the platform team is expected to operate it properly <p>That means: always for production capacities. (learn.microsoft.com)</p>	<ul style="list-style-type: none">• the capacity is critical• peaks are occasional, not structural• the financial model is accepted• monitoring is already in place. (learn.microsoft.com)

Quick checklist

- Is **Surge Protection** enabled on Domain production capacity?
- Is **workspace-level Surge Protection** enabled for Domain workspaces?
- Are **Mission Critical** workspaces explicitly limited and documented?
- Does each production capacity have an **owner**?
- Are **alerts** configured for each production capacity?
- Is **Real-Time Hub alerting** planned or implemented?
- Is **Capacity Overage** enabled only where justified?
- Is every overage limit **capped and approved**?
- Do repeated alerts or overage events trigger a review of sizing or isolation?

1. Surge Protection


What to understand

Surge Protection helps reduce the impact of heavy **background activity** on a capacity. It is especially useful when interactive or user-facing workloads share capacity with more variable background jobs. Microsoft recommends it for shared capacities, but also states that critical solutions should still be isolated on a dedicated capacity. (learn.microsoft.com)

What it is good for	What it is not
<ul style="list-style-type: none">• protecting shared capacities from bursty workloads• reducing the impact of noisy workspaces• limiting background pressure on user-facing workloads. (learn.microsoft.com)	<ul style="list-style-type: none">• not a replacement for dedicated capacity• not a guarantee that all interactive requests will always succeed• not a fix for structural under-sizing. (learn.microsoft.com)

What we should put in place

- enable it mainly on **Domain production capacity**
- add **workspace-level Surge Protection**
- define a rule for handling repeated noisy workspaces
- keep **Mission Critical** status limited to a very small number of justified cases
- tune thresholds using the **Capacity Metrics App**, not guesswork. (learn.microsoft.com)


 **Key message**

Use **Surge Protection** to control shared Domain workloads, not to protect the Data Platform Core workspace instead of isolating it.

Example

🔍 Surge Protection

Workspace Consumption (Preview)

Set a maximum percentage of capacity that any single workspace can consume across both interactive and background operations. Lower limits help prevent any one workspace from overusing capacity. Limits apply over a rolling 24-hour window. High-priority capacities can be excluded. [Learn more](#) 

On

Rejection threshold: % Block: Hours

Background Operations

When total CU consumption reaches the rejection threshold, reject new background operation requests. When total CU consumption drops below the recovery threshold, accept new background operation requests.

On

Rejection threshold: % Recovery threshold: %

2. Notifications / Alerts

What to understand

Notifications are the minimum control that turns capacity health into an operational process. Fabric supports both **capacity notification emails** and **Real-Time Hub / Capacity Overview Events** for monitoring and alerting. (learn.microsoft.com, learn.microsoft.com)

What we should put in place

- For every production capacity:
 - one named **operational owner**
 - one shared **distribution list or team channel**
 - clear **alert thresholds**
 - one documented **escalation path**. (learn.microsoft.com)

Suggested implementation path

Minimum setup

- enable capacity notification emails

Recommended target

- implement alerts using **Fabric Capacity Overview Events** in **Real-Time Hub**. (learn.microsoft.com)



Key message

Alerts should be mandatory on all production capacities. A production capacity without an owner and alerting is not operationally governed.

Example



⏏ Notifications

Send notifications when

- You're using % of your available capacity
- You've exceeded your available capacity and might experience slowdowns

Send notifications to

- Capacity admins
- These contacts:

 Theo COLOMBANI ×
 Margaux COUDUN ×

3. Capacity Overage

What to understand

Capacity Overage allows Fabric to use extra compute beyond the purchased limit to avoid throttling. Microsoft positions it as a way to absorb **rare unexpected spikes** or **small regular peaks**, not as a substitute for proper sizing. It is available only on **F SKUs**. (learn.microsoft.com, learn.microsoft.com)

What it is good for	What it is not
<ul style="list-style-type: none"> reducing the risk of disruption during occasional overload protecting critical production continuity avoiding throttling for short, unexpected peaks. (learn.microsoft.com) 	<ul style="list-style-type: none"> not a performance booster not a strategy for permanent under-sizing not something to leave effectively unlimited. (learn.microsoft.com)

What we should put in place

For **Data Platform Core**:

- enable only if uptime is critical
- cap the limit
- require platform and budget owner approval
- review every overage event

For **Domain production**:

- use only if the business accepts the cost model
- keep tighter limits
- do not use it to hide repeated saturation. (learn.microsoft.com)



Key message

Use Capacity Overage as a controlled safety net, not as a normal operating model.

Example

Capacity Overage

Set a limit on CU hours billed at the overage rate in a rolling 24-hour period, thereby avoiding throttling. [Learn more](#)



On

Rolling 24-hour overage limit



48

CUhr

100% daily additional capacity.

This setting uses 1,172% of your subscription's CU quota (6 of 512).

Need more? [Learn how to increase your quota](#)