~~Loca~~tions: Practical Tips for Procurement Location Management in Oracle Fusion

Author: Michael Gibby

mgibby@hcg.com | www.mikegibby.com

2025 Oracle ACE Associate

Contents

Overview: Locations and Location Sets

Locations are one of those Oracle Fusion setup areas that seem simple—until they aren't. At go live, clients usually want to migrate their legacy data, and don't place much value on a clean up or re-structuring effort. Most clients don't think about them much until something breaks: requisition LOVs are mysteriously limited, change orders fail validation, or users swear a location "used to be there yesterday."

Almost every one of those issues traces back to how Locations and Location Sets were structured (or not structured) early on, often exacerbated by data quality during migration and go-live.

This article focuses on what actually matters: how locations are used by end users, why Common Set causes so many downstream issues, and how to structure Location Sets so you don't paint yourself into a corner later.

How Locations Are Used in Procurement (UX Comes First)

From an end-user perspective, locations matter for one reason: **they determine where stuff shows up** (just like when they order from Amazon at home!)

All users must set a preferred **Deliver-To** location, and they can change it during requisition checkout. When you order something, you want to make sure you actually get it — so getting this right is critically important.

If users can't easily find the right Deliver-To, adoption suffers immediately. In RSSP especially, **location naming is part of your UX design**, whether you intended it or not.

How Users Actually Search

Search behavior in Fusion is simple and unforgiving:

- Users type what they *think* the location is called
- They rarely scroll or refine search results
- They often don't know the "official" address — especially after mergers, rebrands, or remodels

(And yes, it doesn't help that advanced search is only available in Preferences, not in the cart. If that seems wrong to you, please upvote: Enable advanced search for Deliver to location in create requisition in redwood — Cloud Customer Connect)

If your naming doesn’t align with how humans think, adoption drops — no matter how good your data or other configurations are.

In fact, you need to be *better* than humans think, because addresses love to be difficult.
Portland, Oregon is across the *country* from Portland, Maine.
Georgia the state is across the *globe* from Georgia the country.

A Proven Location Naming Pattern That Works

What I’ve consistently seen work best is a **human-first, left-to-right structure**:

Building / Common Name – Street Address – Floor – Area / Office

Example:
Warbler Building – 155 W Van Buren – 2nd Floor – Office 123

Why this works:

- Users remember the **building name first**
- The address **confirms** they picked the right site
- Floor and office help with **last-mile** delivery
- Searching on *any* part of the name usually succeeds

In RSSP, users might search:

- “Warbler”
- “Van Buren”
- “2nd Floor”
- “Office 123”

All of those should work. Just keep in mind, when naming buildings or picking addresses, that Redwood requires 3 characters to search. This means if you have an "IT Building," users searching for just "IT" will be frustrated.

What Not to Do (Common Anti-Patterns)

Some naming patterns consistently cause pain:

- **Address-only names** — users don’t remember street numbers
- **Internal codes first** — “LOC-FL-0047” is as helpful as “IDK-LOL-HBU?”
- **Inconsistent ordering** — floor first sometimes, building first other times
- **Overloaded or inconsistent abbreviations** — BLDG, STE, RM
- **Too much punctuation** — you-get-it

Be Intentional About Prefixes

Prefixes can be extremely useful — if used consistently.

Good uses

- ZZZ – DO NOT USE – Old Dock
- TEMP – Mobile Clinic – Orlando
- DC – Main Warehouse – Atlanta

Bad uses

- Random numbers
- BU codes users don’t understand
- Prefixes that mean different things in different BUs – especially critical during mergers and acquisitions, or in global enterprises.

If a location should not be selected, make that obvious in the **first three characters**.

Deliver-To vs Ship-To (Simple Mental Model)

Here's the simplest way to think about it:

- **End users** care about **Deliver-To**
- **Suppliers** care about **Ship-To**

Ship-To is what appears on the PO as the destination for the truck — usually a dock. Deliver-To is where the goods actually need to end up — like a second-floor office or lab.

At the Location level, Oracle enforces a simple rule:

- A location can be its own Ship-To
- If it isn't, it must reference **one** Ship-To location

Where This Becomes Limiting

Many Deliver-To locations have **multiple delivery routes**, for example:

- Small packages → Front Desk → Research Lab
- Compressed gas → Hubbard Receiving Dock → Research Lab
- Pallets → Main Dock → Research Lab

End users only see Deliver-To in the requisition. They can't easily see or change Ship-To — it's hardcoded in the location relationship.

Until Oracle delivers a more flexible, derived Ship-To model, the leading practice is to create **multiple Deliver-To locations**, each tied to a different Ship-To. In this example, rather than create a single Research Lab deliver to, we create it 3 times, each time linked to a unique Ship To. Each Ship To just needs to be created once.

Example:

- Front Desk | Research Lab
- Hubbard Dock | Research Lab
- Main Dock | Research Lab

This makes routing explicit without exposing complexity to users.

Location Sets 101 (The Part People Miss)

A location does not need to exist globally in Fusion. It exists **inside a Reference Data Set**, and that data set is explicitly assigned to a BU via **Manage Business Functions**.

Key rules:

- **Common Set** locations are visible to all BUs
- **BU-specific Location Sets** are visible only to assigned BUs
- Oracle is doing exactly what you told it to do — even when it feels wrong

This explains the two most common complaints:

- *"I can't find locations I need"*
  → They're in a different Location Set or filtered by inventory org
- *"I see way too many irrelevant locations"*
  → They're defined in Common Set incorrectly

The Common Set Trap

Many rapid implementations put locations in Common Set because it's easy and everything shows up everywhere.

That works… until it really doesn't.

The moment you introduce:

- Multiple BUs
- Multiple inventory organizations
- Redwood Self Service Procurement
- Tighter data governance

Common Set becomes a liability.

Locations drive behavior — inventory org filtering, Deliver-To LOVs, change order validation. *Common Set removes the guardrails*. Suddenly people are ordering items across the globe and generating **massive** cross-entity accounting.

It will drive you mad!

The Right Fix (And the Safe Way to Do It)

The correct long-term solution is to move operational locations into **BU-specific Location Sets**.

But this must be done carefully.

What works in practice:

1. Create new BU-specific Location Sets
2. Re-create required locations in those sets
3. Leave Common Set locations temporarily
4. Rename old locations (e.g., ZZZ – DO NOT USE)
5. Deactivate only after confirming no active POs reference them

Oracle generally allows operational transactions (like receiving) to continue even if locations are no longer selectable. But once you initiate a change order, stricter validation kicks in.

That's why this needs to be phased — not rushed.

What Actually Works vs What Breaks (Test Results)

Receiving

- Existing PO references old Location Set
- Location no longer selectable
- **Result:** Receipt succeeds

Change Order – Not Invoiced or Receipted

- Deliver-To editable
- LOV restricted by inventory org
- **Result:** Works if replacements are valid

Change Order – Already Invoiced or Receipted

- Deliver-To locked
- **Result:** Change not allowed!!

Change Order Submission

- Deliver-To updated, Ship-To/Bill-To not updated
- Mixed Location Sets
- **Result:** Submission fails

Change Orders & Header vs Line Ship-To (A Subtle but Dangerous Quirk)

POs have both header-level and line-level Ship-To values. Changing one does not automatically update the other.

When all lines ship to the same place, Fusion **infers** the header display — but that does not mean the stored header value changed.

Once you edit the PO, Fusion stops inferring and shows the real value.

The Four Scenarios

1. Edit header and lines → everything matches
2. Edit header only → lines wrong
3. Edit lines only → UI looks right, PO prints wrong
4. Edit nothing → no issue

Training rule:
Edit everything or edit nothing.

The Inventory Org “Gotcha” in Redwood

In requisitions:

- Deliver-To is tied to an Inventory Org
- Preferences determine filtering without explicitly telling you that they are doing this filtering
- Checkout LOV is restricted accordingly to only show deliver to from the Inventory Org tied to your Preferred deliver to

If the correct location is in another inventory org:

1. You won’t see it
2. *Searching won’t help*
3. You must change Preferences first

This is especially confusing in industries that don’t *think* in inventory orgs — like healthcare, higher education and financial services. The very thing that users thought would help them, searching, betrays them.

Healthcare Feature: B2B Account Numbers on Ship-To

Oracle now supports assigning **B2B account numbers to Ship-To locations** and sending them on outbound POs — critical for healthcare and integrations like GHX.

This used to require custom DVMs or reporting hacks. Now it’s seeded, supported, and much cleaner.

For healthcare organizations with many facilities, this significantly reduces fulfillment errors and ambiguity for suppliers while reducing technical debt. Anyone still using an outdated method should move to the new supported method as soon as possible.

Final Takeaway

Location Sets aren't just setup — they shape UX, LOV behavior, and document lifecycle rules.

If you remember nothing else:

- Avoid Common Set for operational locations
- Align Location Sets intentionally to BUs
- Train your users about inventory org filtering in Redwood
- Phase migrations carefully with open PO's
- Test change orders — not just requisitions

Getting this right early saves months of confusion and prevents "Redwood bugs" that are really setup decisions catching up to you.

About the Author

Michael Gibby is a Director at Huron Consulting Group, where he specializes in Oracle Fusion Supply Chain in the Higher Education, Healthcare and Commercial industries. His diverse background in international manufacturing using EBS helps him drive practical but innovation solutions, rooted in solid business process design. When he isn't implementing Oracle P2P solutions, he is busy raising two future Oracle Cloud Consultants, and occasionally he finds time to rebreather dive in underwater caves in North Florida.