UPDATE MANAGEMENT

Get your development together.

Get your [stuff] together. Get it all together and put it in a back pack. All your [stuff]. So it's together.
—_Morty_

UPDATE SETS are used to track, maintain, and deploy development changes – whether it's configuration, code, layouts, or anything else that's tracked in Update Sets. Once your changes are ready for deployment to another instance, Update Sets are also how you move your changes between instances. You can export them to XML (once the State is set to Complete) and import that XML into the target instance, or you can retrieve them from the target instance before previewing and deploying them. Whichever method you use, it's important to _sanitize_ your Update Sets, keep them clean, and keep them functional. That's "functional" as in "functional programming".

In this chapter, we're going to learn about **batching**, **merging**, what is and **isn't** tracked in Update Sets, common pitfalls when dealing with scoped records, and how to correctly handle **preview** errors.

UPDATE SET BATCHING

ServiceNow released **Update Set batching** in the Jakarta release of the Now platform, allowing you to **group** Update Sets together in a logical hierarchical grouping. This enables you to preview and commit them in bulk, and in order.

Update Set batching takes a lot of the hassle and annoyance out of managing large groups of Update Sets for a major release, and allows you to extract a specific set of functionality from a release without impacting the rest. For example, consider a scenario where you need to do **A,** which enables you to do **B**, which enables you to then do **C**. If there's a problem with **B**, whoever does the deployment will have to be aware that **C** relies upon **B**, or the deployment of **C** will fail spectacularly!

"Dependencies" like this are a major issue, but also consider which of the records in your update sets might **reference** other records. For example, if you have a table which extends another table, you cannot push the **child**
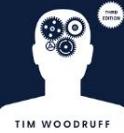
table before pushing the **parent** table, because the child references the parent table in order to inherit its fields and attributes.

Hierarchical Update Set batching makes this process easier, but is by no means a catch-all solution. You still have to **build**, **maintain**, and **control** the Update Set hierarchy and think carefully about where each Update Set belongs in the chain. Typically, the best-practice is to have the chain build chronologically, in order to avoid having a newer version of a record in a parent Update Set, and an older version of the record, in a child Update Set (which would result in an error).

If a child update set contains a record that also exists in one of the update sets that is "up-stream" from it, the version of that record in the child update set **must be newer** than the version in the parent set.

> **Pro-tip**: Deploying an update set will cause a system cache flush, which can impact performance. It is often a good idea to deploy during lunch, or before or after business hours.

MASTER UPDATE SETS

It's not a bad idea to begin a release or large project by creating a "Master" Update Set in your dev environment. This is the Update Set under which all related "child" Update Sets will live. It's also good to have a naming convention that makes sense, but that's something your team will have to figure out internally to determine what works best. I'll use the name **Release April 2018-Master** as my example "master" Update Set.

No updates will actually go into the master Update Set. Instead, this will just be used to "contain" the Update Sets that contain any development that should be moved along as part of the "April 2018" release. Whenever a developer creates a new Update Set for some development, they'll set the **Parent** field on their Update Set to this master set.

As an example, imagine I've created two Update Sets for some development that I've been assigned stories for: **TW-Automate Inc Assignment-v1.0** and **TW-Build Catalog Item-v1.0**, and give them both helpful descriptions. I set these Update Sets' **Parent** field to **Release April 2018-Master**, since they're both scheduled for the April 2018 release, but if they needed to be pushed back at some point, I could give them a different parent. The master Update Set would not have a parent.

After identifying something additional that needs to be added, I create **TW-Build Catalog Item-v1.1**. This "1.1" Update Set will be a child of **TW-Build Catalog Item-v1.0** *rather than* **Release April 2018-Master**, since it is dependent directly on v1.0. This way, if the story to which they both relate needs to be moved back to the next release, I can just reassign v1.0, and v1.1 will come along with it. By making 1.1 the child of 1.0 rather than the child of the master set, I also ensure that the order of deployment is maintained, which is important because 1.1 will surely reference records contained within 1.0.

Let's have a look at what these relationships look like visually. You can click the **Visualize Update Set Batch** from any Update Set form, to see a visual representation of the relationships between your Update Sets and their hierarchy.
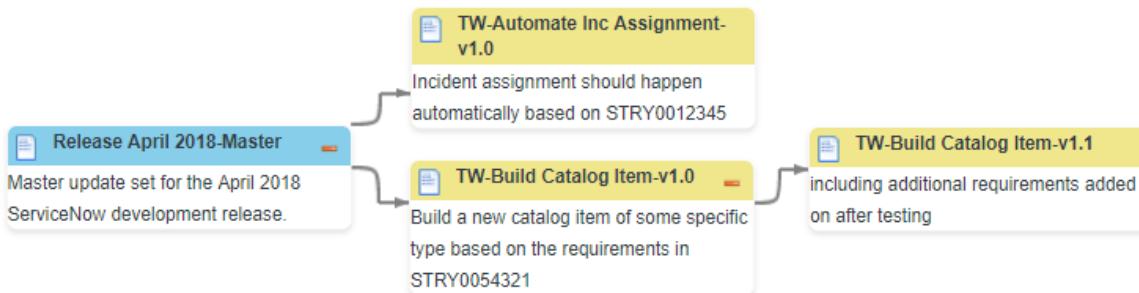
Note: Right-clicking an Update Set in this view, and clicking **Prune**, will only prune it from that **view**. This does not actually prune the Update Set from the batch hierarchy. If you refresh the batch visualizer, the pruned Update Set will reappear.

As you can see, from the master upset set, we split into two branches; but you can have as many child Update Sets as you like. If we then right-click on the master Update Set, and click **Edit**, we can see the **Child Update Sets** related list. This is a list of all of the Update Sets which are **immediate** children of the master set. It does not include the *children of the children* of the master set.

To see **all** of the Update Sets in the entire batch, go to the maser Update Set, and open the **Update Sets In Batch** related list.

Note: *Closing the master Update Set will close **all child Update Sets in the batch**. However, closing any child Update Set – even if it has children of its own – will not close any others in the batch, including its children. Weird, right?*

MERGING UPDATE SETS

Merging your Update Sets is advisable over batching in most scenarios where you're the only developer working on a given subset of functionality, because it provides a "safer" and more consistent deployment. That said, there's nothing wrong with batching to keep your Update Sets connected for tracking purposes, and then merging prior to deployment.

Batching is still a perfectly acceptable solution, especially when there are multiple developers working on a single project. You just have to be mindful of the "hierarchy" of the update sets in the batch.

WHAT IS AND ISN'T TRACKED

In ServiceNow, it's important to understand the difference between **data** and **configuration/customization**. Data is stuff that fills up tables.

Only what ServiceNow considers **configuration/customization** is captured in Update Sets, but it can sometimes be a little bit difficult to pin down what constitutes configuration without simply knowing in advance. For this reason, if you're not 100% certain that a particular record type is indeed captured in Update Sets, it's a good idea to open your Update Set in another tab, sort your updates by the *sys_updated* column, and make sure that your changes were captured.

Records which are not "captured" in Update Sets can still be forced into them. For example, if you want to force a specific record in the **Location** [`cmn_location`] table (which isn't tracked) into your Update Set, you can do so in a Background Script using the GlideUpdateManager2 API. In some cases, such as with a large number of records, it makes more sense just to export them all to XML and manually import them into the target instance after deploying your Update Set.

Pro-tip: *To make it a little easier to include data records (since nearly every developer needs to do this quite often), you can download a tool I wrote, which puts a "Include in Update Set" button in the Related Links of every record and list context menu for every table that isn't tracked in Update Sets if you're an admin. You can find this free tool at https://updatetracker.snc.guru/!*

While **not exhaustive**, here is a list of some of the records that are/aren't included in update sets. You might be surprised by some of the examples.

The following **are** tracked in Update Sets:

- Roles
- Form sections and layouts
- Reports
- Form and list views

The following are **not** tracked in Update Sets:

- Scheduled jobs & scheduled script executions
- Schedules
- Configuration Items
- Users
- Groups (and group *types*)
- Group memberships
- Role assignment (to users *or* groups)
- Homepages/Dashboards

You'll also find that there are certain records that are dependent on other records for their functionality. For example, if you create a role in the **Role** [`sys_user_role`] table, that will be tracked. If you then create a **group** (which isn't tracked), add the **role** to that group, and then force that group into your update set, you might be surprised to find that the roles you gave it **didn't come along with it**. This is because the associations between roles and groups is stored in the `sys_group_has_role` table which is also not tracked in update sets. Similarly, the association between **users** and roles is in the `sys_user_has_role` table, which is also not tracked in update sets.

There is a similar situation that happens with **dashboards**. Dashboards contain several linked *pieces* that are each not captured in update sets. Each new tab in the `pa_tabs` table is linked to two tables (`sys_portal_page` and `sys_grid_canvas`). To get dashboards into your update set, you have to grab each of those types of records.

PRIVATE SYSTEM PROPERTIES

One special case to be aware of, is **private System Properties** (that is, System Properties with the "**Private**" checkbox checked). These records *are* tracked in your Update Set, but when you go to deploy that Update Set in a target instance, you'll notice that the System Property is not actually created (if it doesn't already exist), and will not update the System Property if it does already exist in the target environment. This functionality is actually rather intentional, and is more "feature" than "bug", insofar as the whole purpose of the "**Private**" checkbox, is to prevent accidentally promoting System Properties containing data that is **instance-specific**. For example, imagine you've got a system property that controls which MID server to use for a given set of REST Methods, when invoked via script. You might very likely have different MID Servers for your development environment, than you do for your production environment. Therefore, you would want this property to be **private** so that when you set the property in dev, it doesn't accidentally get carried over to prod and overwrite the setting pointing to your production mid server.
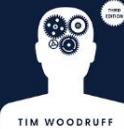
Normally, this is a good thing. However, it's something that's important to be mindful of when – not updating but **creating** a System Property that's private. This is because on the target instance, the INSERT within your update set for a private System Property is completely ignored, just like an UPDATE to that System Property would be. This means that the *sys_properties* record would not even be created! Because of this, you must be mindful to do a few things:

1. Always make a "pre-deployment" and "post-deployment checklist", and include that checklist in your Update Set, to ensure that – among any other necessary steps – any private System Properties are created.
2. As mentioned in the System Properties section of the Configurability chapter in this book, always specify a "default value" (the second argument) for System Properties that aren't found, when calling `gs.getProperty()`.
3. Always practice **defensive programming** – *even against yourself*! In part, this means thinking about all of the ways that anything could possibly go wrong and checking for those issues in your code. For example, perhaps you can't reasonably specify a "default value" for a System Property, because you'd have no way to know an acceptable value if the property isn't set. In that case, specify a default value that is just a blank string (`''`), and then have another couple lines of code that check for that like so:

```
var myProp = gs.getProperty('some_prop', '');
if (!myProp) { throw new Error('System property "some_prop" not set.'); }[12]
```

TRACKING SCOPED RECORDS

To be honest, working with Update Sets inside of scoped applications in ServiceNow can be a real bother. In this section, we'll go over a few specific rules to be aware of, to hopefully save you and your team some headaches.

For starters, it's important to understand that a single Update Set should never contain records from multiple application scopes. The platform tries to prevent this, but for some reason, it happens anyway. If at all possible, you want to avoid it at all costs. Update Sets have a specific scope, and you cannot deploy updates in one scope using an Update Set in another scope.

There are plenty of scenarios where you may need to work within multiple scopes. For example, the **HR** application consists of multiple separate application scopes including **Human Resources: Core** and **Human Resources: Service Portal**. One development effort may require that you work with records in both HR scopes, as well as records in the Global scope. This typically requires manually creating **three separate Update Sets**: one for

---

[1] *Please note that unfortunately, ServiceNow's `Error` constructor does not properly set the `lineNumber` property of the `Error` object. For this reason (and obviously, others), it's a good idea to be very descriptive in your error messages if your use a `throw` statement.*

Exceptions caught by ServiceNow's script processor will correctly set the `.lineNumber` property, so if you put your code in a `try/catch` block, and there is an error that you didn't manually `throw`, you'll be able to use that property in your error messaging to determine where exactly the error was.

[2] *The practice of detecting errors early and throwing or returning them, is called "failing early". This best-practice makes your code more clear and succinct, and also makes failure more apparent and performance-friendly. It also minimizes the potential negative externalities that may be caused by erroring code. For example, if your code updates a record or does a REST call or something before the error is detected, but the error prevents your code from continuing, you may be left with artefacts of all the work it did between the error being detectable, and being detected.*

Always aim to detect errors early, and – if they're show-stoppers – `throw`, halt, `continue`, or `return` early!

each scope. The Update Set you have selected persists when you switch back to a specific scope, but not between scopes.

For example, if you have **Update Set 1** selected in the **Human Resources: Core** scope, then you switch to the **Global** scope and select **Update Set 2**, and finally switch *back* to the **Human Resources: Core** scope, you'll find that you still have **Update Set 1** selected. This will persist until you change your Update Set *while in that scope*.

There are some additional headaches to contend with, too. For example: typically, when you visit a record that's in a different scope than your session is in, you cannot edit the record and you get a warning like this:

This record is in the Human Resources: Core application, but Global is the current application. To edit this record click here.

You can click the link at the end of that message to switch (for the duration of that one update) to the scope in which the record resides, but you won't be able to see what Update Set is selected in that scope unless you switch your whole session to the relevant scope.

Another reason to be very careful when dealing with scoped updates, is that some **tables** are in a certain scope, but the **records** in them are in a different scope. Again, ServiceNow is working on this, but when this happens, there is no good way to tell that you're effectively poisoning your Update Set with updates in multiple scopes, making it so it won't be able to be deployed.

For an example of this, consider the following steps:

1. Select the **Human Resources: Core** application
2. Create an Update Set in the **Human Resources: Core** application. Set it as your current Update Set.
3. Navigate to the **sn_hr_core_service** table.
4. Open any record in that table, and make any update to it, so the record is captured in your Update Set.

You'll notice that the record is tracked in your Update Set. However, in some instances, you'll also notice that even though the **sn_hr_core_service** table (like your Update Set) is in the **Human Resources: Core** scope, the update itself is in the **Global** scope. This Update Set cannot be deployed until all updates it contains are in the same scope that it's in.

Unfortunately, the only way around this right now when dealing with multiple scopes, is to double-check that any new updates you make are tracked in your Update Set, and in the correct scope.
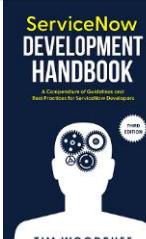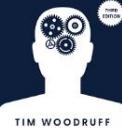
Note: *ServiceNow has been making an effort to patch this issue wherever they find it by updating the plugin, but they do not usually offer retroactive patches if you've already deployed a given scoped application or plugin. You might have some luck if you notice this issue in your instance, by opening a HI ticket (http://hi.servicenow.com/). The key is just to keep an eye on it to make sure your Update Set and update scopes match.*

PROMOTION

Code promotion is obviously one of the most important processes you can have as a development team. I mean, you can write the best code in the world, but if it borks production because of a bad merge, nobody's going to be patting you on the back!

A robust code promotion process consists of a few key factors:

1. Is your **source instance** clean and (aside from the code you're promoting), relatively in-sync with the target instance?
   a. This question often comes down to asking yourself "when was the last time this instance was cloned over?" and "how much work has been started but not finished or abandoned in this environment, since the last clone?"
   b. See the upcoming Cloning section of this chapter for more info.
2. Is everything you intended to deploy, captured in your Update Sets?
   a. Review the section of this chapter on things that *are* and *are not* captured in Update Sets, and ensure everything you want to move is tracked.
   b. Do you have any private System Properties that need to be manually created in the target instance? As mentioned in the Private System Properties section of What is and Isn't Tracked, even if private System Properties appear in your Update Set, they may not be created in the target instance!
   c. Does your functionality rely on any **data** (records which aren't in a tracked table)? If so, make sure that these records are manually pushed into your Update Set, using a tool such as my "**Include in Update Set**" tool (available for free, at https://include.snc.guru). Don't forget to make sure that you've got the **latest** version of every data record you intend to move over![3]
3. Are there any steps that need to be taken manually, **before** or **after** deployment in the target environment? If so, be sure to mention those steps in your Update Set description.
   a. It's important to have a standardized process of deployment which includes either reading Update Set descriptions, or otherwise ensuring that any necessary manual steps are taken on the target environment. This should be **firmly** a part of **any company's deployment strategy**!
4. If issues are identified in testing, **fix those issues in the dev environment** in a **new Update Set**, and push them up via an Update Set. You can batch or re-merge the Update Sets in Test once you're ready to promote, but never make changes directly in Test, except to facilitate testing; for example, by setting a System Property or following a Guided Setup to configure an application for testing – but be sure to do so in the Default Update Set so those changes aren't captured in what's promoted to production. And don't forget to make sure that the necessary set-up steps are captured somewhere!

---

[3] *Don't forget that if you manually push a record into your Update Set using a tool such as my **Include in Update Set** tool (https://include.snc.guru), and then modify that record, it's the **old** version that will still be captured in your Update Set! You must manually re-push the updated record into your Update Set. This applies for data records, and things like scheduled script executions and anything else not captured in Update Sets.*