

TWiki > DICE Web > ResearchAndTeachingUnit > PrometheusTheon (29 Sep 2017, TimColles)

Theon/Prometheus API

This is a formal definition of what Prometheus should receive in the feed of data from Theon.

Views

Theon provides three views that can be queried by Prometheus. Access to these views on [PostgreSQL](#) is constrained to the "prometheus" user. Prometheus provides a kerberos identity per-hostname that needs access and this will be appropriately mapped.

A note on view naming. Many Theon database objects are currently suffixed with `_3g` while legacy aspects remain in the system. When the Inventory structures have been migrated out then the `_3g` will be removed. Ensure that all interfaces to Theon can be modified at that time.

prometheus_user[_3g]

The "session register" in Theon (and accessible to computing staff via the "Register" desktop in [TheonUI](#)) holds the set of user registrations that are considered to be "Current" (live) within the current academic session (which runs from the 1st August to the 31st July). A person may have multiple registrations in this register based on different underlying status values, for example a student may be taking a degree programme with the School while simultaneously applying for a subsequent degree programme with the school or an on-programme student may also be a contracted staff member (through providing Teaching Support, or as a Research Associate).

The `prometheus_user[_3g]` view contains one row for each of a specifically defined and limited set of user session registrations (not necessarily one row per person as explained above). This view represents the master list of all users that we would expect to be entitled to have a DICE account within Informatics, either now or at some defined point in the future. Conversely a person who does not have a user registration listed in this view should not be entitled to a DICE account, and if they have currently got a DICE account it should be placed into expiry as per the lifecycle process defined in Prometheus. A user registration listed on this view is not necessarily given a DICE account however. The "entitlement" to a DICE account is actually defined by the particular roles that are assigned to that user registration in `prometheus_role[_3g]` and whether any of those roles (or roles they inherit) are DICE account granting. If a user is in this view they are guaranteed to always have at least one role record in `prometheus_role[_3g]`.

Each user registration has a unique row identifier which is the UUID. This is guaranteed to always be unique. However, it may move between user registrations in very rare circumstances. An example would be duplicate user registration entries under the same categorisation (see below), in which case a "merge" would be carried out.

As stated above do not assume a one-to-one correspondence between a person and a user registration in the `prometheus_user[_3g]` view. It is possible for the same person to exist with multiple user registrations in this view, each with a different UUN. There are three UUN categories represented in Theon. A UUN will always be unique across all categories. The three categories are: Student, Staff and Visitor (although note that visitors can be Staff, do not assume a v1 UUN is actually a visitor). One person may have up to three separate rows in the `prometheus_user[_3g]` view, although generally two at most while traversing from one category to another. A student that is also employed as research staff for example could permanently have multiple user registrations represented. Theon makes no attempt to conjoin user registrations for the same person, they are all different people as far as Theon is concerned.

The columns in the `prometheus_user[_3g]` view are:

- "person" - the UUID value and the unique identifier for the user registration
- "surname" - the surname of the person associated with the user registration
- "firstname" - the informal firstname (falling back to formal firstname) of the person associated with the user registration
- "enrolment" - the matriculation number of the person associated with the user registration, or blank if this is not a student user registration
- "username" - the UUN of the person associated with the user registration, or blank if the UUN is not defined (see below)

It should not be the case that when "enrolment" has a value that the corresponding "username" value is not "s" concatenated with the enrolment. Conversely when the "username" value is not a student UUN then it should always be the case that the "enrolment" value will be blank. When either of these do happen it means that multiple categories of identity within Theon have been merged into one category record (e.g. when a student record has incorrectly been given the UUN from a staff record) and this should be corrected in Theon.

A person will appear in the view as soon as they have at least one primary role from the Theon set of roles (see below). This is not necessarily a DICE account granting entitlement role, so while they may appear in this view as being entitled to a DICE account there may not be a correlation in Prometheus with the DICE account granting entitlement and role to give them a DICE account. A person will disappear from this view as soon as they have no primary roles attributable from the Theon set of roles.

A blank username will occur when the same value of UUN is applied to a different record, irrespective of whether its the same underlying person or not. So for example a person with a staff record has username "abc" and a row in this view. If the username "abc" is assigned to another person record (staff, student or visitor) correctly or in error and irrespective of whether the record is associated with the same person or not then this does not result in a duplicate error (a UUN must be distinct across all person records in Theon) but instead the original instance of that username "abc" will be blanked. So while the record may still result in a valid user registration with roles and will appear in this view, the username will be empty. This can represent simple and correct traversal of a UUN (for example staff member leaving and becoming visitor and then the UUN being reconciled) but can also be considered as an error condition, temporary or otherwise. At the moment Prometheus will ignore user registrations in the feed with an empty username (treat them as not entitled to a DICE account and to enter grace period under lifecycle rules if they previously had a DICE account), it may be better to flag them either within Theon or in Prometheus.

prometheus_role[_3g]

The prometheus_role[_3g] view contains one row per role per user registration. It can be tied back to the prometheus_user[_3g] view by the unique person UUID value. The columns in the prometheus_role[_3g] view are:

- "person" - the UUID value and the unique identifier for the corresponding user registration in prometheus_user[_3g]
- "role" - the name of a primary role (see below)

Each user registration can have any number of roles. There is no delay in assignment or removal of roles within Theon, although some roles are used to represent specific timed periods of status.

prometheus_contact[_3g]

The prometheus_contact[_3g] view contains at most one row per user registration. It can be tied back to the prometheus_user[_3g] view by the unique person UUID value. This view contains the contact details the associated person has recorded for that user registration. The columns in the prometheus_contact[_3g] view are:

- "person" - the UUID value and the unique identifier for the corresponding user registration in prometheus_user[_3g]
- "email" - the preferred public email address of the person associated with the user registration. While a particular user may have multiple email addresses associated with them, only one is ever returned in this field based on rules reflecting the preferred address (usually set via the self-service interface) followed by some fallbacks based on UUN and user category. Note that the email address held here may not always correlate with the email address published on the generated personal web page due to the different rule sets in force (that choose between various possibilities of email address) on those pages which also use some legacy data.
- "extension" - the telephone extension number for where the person is located. This will contain multiple numbers aggregated and separated by a single semi-colon on the rare cases where a person has more than one extension recorded for them.
- "room" - the room number where the person is located. This will contain multiple rooms aggregated and separated by a single semi-colon on the rare cases where a person has more than one room recorded for them. The format of the room will be all uppercase site acronym dash room designation. For example, IF-4.02A, IF-G.07 or AT-3.12. The naming notation may vary slightly however for remote offices in non-Informatics buildings.

Roles

Roles (and entitlements) are used to manage authorisation policies within Informatics. For every primary role allocated by Theon, there should be a corresponding rfe map. The management and use of roles and entitlements is described extensively in the Prometheus documentation.

Theon defines the following primary roles.

Staff Roles

Staff will have one and only one of the following three staff specific roles.

- The "staff" role is given to any record in Theon that has a status of 'Academic', 'Administrative', 'Computing', 'Research', 'Technical' and a currency of "Existing" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. They will also be given an additional "STATUS-staff" role where status is one of the specific named categories above.
- The "future-staff" role is given to any record in Theon that has a status of 'Academic', 'Administrative', 'Computing', 'Research', 'Technical' and a currency of "Upcoming" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. If they have a start date against their record it must be greater than the current date plus one month (i.e. they must be more than one month from their start date). They must not have been marked as deleted by the feed processes. They will not be given an additional "STATUS-staff" role while they have the "future-staff" role.
- The "new-staff" role is given to any record in Theon that has a status of 'Academic', 'Administrative', 'Computing', 'Research', 'Technical' and a currency of "Upcoming" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. If they have a start date against their record the current date must be within one month of their start date. They must not have been marked as deleted by the feed processes. They will not be given an additional "STATUS-staff" role while they have the "new-staff" role.

Very rarely a staff member may have more than one status and in these circumstances will have each corresponding "STATUS-staff" role, and potentially more than one primary staff role (if currently "Existing" on one status and transitioning to the other status and so "Upcoming" for that).

It is beyond the scope of this document to define how staff are mapped from upstream HR/IDM feed data into the five local staff status categories and two session temporal categories shown here. Note however that it is not an entirely determined process as we do not have a formal definition of the upstream feed itself so cannot entirely know what to always expect in it. So sometimes users which should receive a staff role (and hence a user registration) will not do so simply because either their status or currency have not been appropriately mapped.

Visitor Roles

Theon has five distinct visitor categories that all result in a status of "Visitor":

- "VisitorStaff" - these staff alike visitors are all taken from VRS and correspond with this type in VRS.
- "VisitorStudent" - these student alike visitors are all taken from VRS and correspond with this type in VRS. Note this is distinct from students locally registered in our University on visiting degree programmes which are handled alongside students on all other degree programmes, they are categorised as students **not** visitors.
- "LocalVisitorStaff" - this is a Theon registered only visitor (not in VRS) and represents a member of staff at another school within the University visiting us (requiring a local DICE account), formally "current" staff and students should never be entered into the VRS (although it happens occasionally by accident).
- "LocalVisitorStudent" - this is a Theon registered only visitor (not in VRS) and represents a student at another school within the University visiting us (requiring a local DICE account), formally "current" staff and students should never be entered into the VRS (although it happens occasionally by accident). Note this is not a student on a "visiting degree programme" with us as they would then be a student **not** a visitor. This is often used to give a DICE account to students in other schools that have a Teaching Support contract with us.
- "LocalVisitorVisitor" - this is a Theon registered visitor that is also in VRS with a visitor type of "VisitorStaff" or "VisitorStudent" but originally/primarily registered as a visitor at another school in the University - the VRS can only represent one affiliation concurrently so this is required when a visitor to us has a current VRS record for elsewhere.

Note that the VRS does not hold the "sponsor person" so an entry in VRS will not result in a role (and hence user registration) until this has been locally added in Theon (see rules below).

Visiting staff (as defined by VRS and within Theon above) will have one and only one of the following three visiting staff specific roles.

- The "tempvisitor" role is given to any record in Theon that has a status of "Visitor" and a currency of "Existing" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. They must have a sponsor person defined for the currently applicable visit record and that sponsor person must be a current member of staff (must have the "staff" role exactly as it is defined above). Additionally the visitor category must be only one of "VisitorStaff", "LocalVisitorStaff", "LocalVisitorVisitor".
- The "future-tempvisitor" role is given to any record in Theon that has a status of "Visitor" and a currency of "Upcoming" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. If they have a start date against their record it must be greater than the current date plus one month (i.e. they must be more than one month from their start date). They must have a sponsor person defined for the currently applicable visit record and that sponsor person must be a current member of staff (must have the "staff" role exactly as it is defined above). They must not have been marked as deleted by the feed processes. Additionally the visitor category must be only one of "VisitorStaff", "LocalVisitorStaff", "LocalVisitorVisitor".
- The "new-tempvisitor" role is given to any record in Theon that has a status of "Visitor" and a currency of "Upcoming" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. If they have a start date against their record the current date must be within one month of their start date. They must have a sponsor person defined for the currently applicable visit record and that sponsor person must be a current member of staff (must have the "staff" role exactly as it is defined above). They must not have been marked as deleted by the feed processes. Additionally the visitor category must be only one of "VisitorStaff", "LocalVisitorStaff", "LocalVisitorVisitor".

Visiting students (as defined by VRS and within Theon above) will have one and only one of the following three visiting student specific roles.

- The "visitingstudent" role is given to any record in Theon that has a status of "Visitor" and a currency of "Existing" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. They must have a sponsor person defined for the currently applicable visit record and that sponsor person must be a current member of staff (must have the "staff" role exactly as it is defined above). Additionally the visitor category must be only one of "VisitorStudent", "LocalVisitorStudent".
- The "future-visitingstudent" role is given to any record in Theon that has a status of "Visitor" and a currency of "Upcoming" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. If they have a start date against their record it must be greater than the current date plus one month (i.e. they must be more than one month from their start date). They must have a sponsor person defined for the currently applicable visit record and that sponsor person must be a current member of staff (must have the "staff" role exactly as it is defined above). They must not have been marked as deleted by the feed processes. Additionally the visitor category must be only one of "VisitorStudent", "LocalVisitorStudent".
- The "new-visitingstudent" role is given to any record in Theon that has a status of "Visitor" and a currency of "Upcoming" within the current academic session. If they have an end date against their record it must be greater than or equal to the current date. If they have a start date against their record the current date must be within one month of their start date. They must have a sponsor person defined for the currently applicable visit record and that sponsor person must be a current member of staff (must have the "staff" role exactly as it is defined above). They must not have been marked as deleted by the feed processes. Additionally the visitor category must be only one of "VisitorStudent", "LocalVisitorStudent".

The "sponsor person" must remain a current member of staff (must have the "staff" role exactly as it is defined above), as soon as they lose that role then any visitors they are sponsoring will lose their associated visitor roles and potentially they will lose their user registration as a result.

Note however that the VRS feed content is not robust, in particular with respect to UUN reconciliation, nor is it an entirely determined process as we do not have a formal definition of the upstream feed itself so cannot entirely know what to always expect in it. So sometimes users which should receive a visitor role (and hence a user registration) will not do so simply because either their status or currency have not been appropriately mapped.

Student Roles

Theon has five distinct student categories that all result in a status of "Student":

- "UG[12345]" - these are students on an undergraduate degree programme, usually qualified by the current year of study on that programme (this is different from the number of years they may have been on-programme which could be

substantially longer due to interruptions). The year of study is not always set in the feed (nor in Theon) resulting in an unqualified status value.

- "VUG" - these are students specifically on a visiting undergraduate degree programme, these students are only here for one semester or one year.
- "PGT" - these are students on a taught postgraduate degree programme, either full-time (one year) or part-time (two year).
- "PT[12]" - these are students on a two-year taught postgraduate degree programme, usually qualified by the current year of study on that programme.
- "PGR" - these are students on a research postgraduate degree programme

The same categories are used irrespective of whether the students are applying (currency will be "Applying"), have been given a formal offer (currency will be "Upcoming"), are on-programme including during interruption of studies for whatever reason (currency will be "Existing"), or have completed/exited successfully or otherwise (currency will be "Previous"). A students start date will always be set and is the programme start date (not necessarily the application start date as this can be deferred). A students end date for a research postgraduate student is set locally in Theon based on their programmes defined "maximum" duration and interruptions. A students end date for a taught student is open and is only set when they have either been awarded, exited for some other reason or a period of time beyond failure to confirm attendance. Note however that while a research postgraduate student will also be automatically assigned a record in the upcoming session by Theon, taught students will not (hence giving them an effective end date of the end of the current session). This is done to manage effectively with transitions in the feed that occur over the upstream session roll-over period.

Four types of role are produced specifically for students. These are "cohort", "degree", "year" and "module". Students will usually have one cohort role, one degree role, one year role and many module roles. However in various circumstances it is possible for students to have more than one cohort role and/or more than one degree role. For example on transition from a UG programme to a PGT programme or from a PGT programme to a PGR programme there will often be programme overlap where Theon will interpret their registration status as placing them on multiple concurrent programmes. Also for some cohorts students can apply for multiple degree programmes, be offered a place on multiple programmes and do not need to finalise their choice until some time into the first session of study.

All students will generally have one and only one of the following student specific "cohort" roles.

- The "cohort-ug" role is given to any record in Theon that has a status of "UG[12345]" and either a currency of "Existing" or a currency of "Upcoming" and a programme start date falling within one month from the current date within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. The "current" academic session is defined in this context as running from 01/08/YYYY to 30/09/YYYY+1, so note that the existing and upcoming sessions overlap.
- The "cohort-vug" role is given to any record in Theon that has a status of "VUG" and either a currency of "Existing" or a currency of "Upcoming" and a programme start date falling within one month from the current date within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. The "current" academic session is defined in this context as normal, running from 01/08/YYYY to 31/07/YYYY+1.
- The "cohort-pgt" role is given to any record in Theon that has a status of "PGT" and either a currency of "Existing" or a currency of "Upcoming" and a programme start date falling within one month from the current date within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. The "current" academic session is defined in this context as running from 01/08/YYYY to 30/09/YYYY+1, so note that the existing and upcoming sessions overlap.
- The "cohort-pt" role is given to any record in Theon that has a status of "PT[12]" and either a currency of "Existing" or a currency of "Upcoming" and a programme start date falling within one month from the current date within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. The "current" academic session is defined in this context as running from 01/08/YYYY to 30/09/YYYY+1, so note that the existing and upcoming sessions overlap.
- The "cohort-pgr" role is given to any record in Theon that has a status of "PGR" and either a currency of "Existing" or a currency of "Upcoming" and a programme start date falling within one month from the current date within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. The "current" academic session is defined in this context as normal, running from 01/08/YYYY to 31/07/YYYY+1.

All students will generally have one and only one of the following student specific family of "degree-PROGRAMME" roles. For students on a degree programme run by our school the PROGRAMME part of the role is set according to the students corresponding programme tag value which is managed by administrative staff and outside our control, see below. For students taking our courses but not on a specific degree programme run by our school the PROGRAMME part of the role is always set to "ext".

In all cases below the "current" academic session is defined for taught students (excluding students on visiting degree programmes) as running from 01/08/YYYY to 30/09/YYYY+1, so note that the existing and upcoming sessions overlap. The "current" academic session is defined for research students (and students on visiting degree programme) as normal, running from 01/08/YYYY to 31/07/YYYY+1.

- The "degree-PROGRAMME" role is given to any record in Theon that has a status of "Student" and a currency of "Existing" within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes.
- The "future-degree-PROGRAMME" role is given to any record in Theon that has a status of "Student" and a currency of "Upcoming" within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. The programme start date must be greater than the current date plus one month (i.e. they must be more than one month from their start date). They must not have been marked as deleted by the feed processes.
- The "new-degree-PROGRAMME" role is given to any record in Theon that has a status of "Student" and a currency of "Upcoming" within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. The current date must be within one month away of their programme start date. They must not have been marked as deleted by the feed processes.

The setting of PROGRAMME is entirely at the discretion of our administrative staff. Usually it will be a locally understood abbreviation for the programme, for example "Artificial Intelligence and Computer Science" would be "aics". New

programmes do not necessarily get this configured straight away in which case the default will be the lower case version of the EUCLID code, for example something like "ptmscdatsc1f", until changed. It is likely that over time we will start to see more EUCLID encodings and fewer local abbreviations because the local abbreviations will only be maintained by ITO administrative staff while they have a local requirement and are engaged to do so - as programme assessment is now done upstream in EUCLID and students/staff themselves will become more familiar with the EUCLID code rather than our own abbreviations the local requirement to have our own naming scheme will reduce/go and and there will be little point in us maintaining local programme names ourselves which will have no meaning other than to us.

Taught students only will have one and only one of the following student specific "year-YEAR" roles. The YEAR part of the role is set according to the students cohort and current year of study and will only have one of the following values:

- "ug1" - an undergraduate student in year 1 of study, or a visiting undergraduate student (these are only actually here for one semester usually)
- "ug2", "ug3", "ug4", "ug5" - an undergraduate student in year 2, 3, 4 or 5 of study (the number of additional years on a programme varies)
- "msc" - a taught postgraduate student (this is always a one year programme, distinct years "on-programme" in part-time variations are not represented)
- "pt1", "pt2" - a taught postgraduate student in year 1 or 2 of study (this is always a two year programme, distinct years "on-programme" in part-time variations are not represented)
- The "year-YEAR" role is given to any record in Theon that has a status of "Student" and either a currency of "Existing" within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. The "current" academic session is defined for one year taught postgraduate students as running from 01/08/YYYY to 30/09/YYYY+1, so note that the existing and upcoming sessions overlap. The "current" academic session is defined for all other taught students as normal, running from 01/08/YYYY to 31/07/YYYY+1. The student does not need to be on a degree programme run by our school, so students taking our courses but not on our programmes will also have a "year-YEAR" role.

All students (although usually just taught students) will have any number of the following student specific "module-COURSE" roles, normally up to about twelve. The COURSE part of the role is set according to the corresponding course's tag value which is managed by administrative staff and outside our control, see below. Note that "module" is an entirely deprecated term administratively and in the central records, these roles should ideally be re-named "course-COURSE".

- The "module-COURSE" role is given to any record in Theon that has a status of "Student" and a currency of "Existing" within the "current" academic session and is formally registered onto COURSE within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. The course registration status must not be set as "Withdrawn". The "current" academic session is defined in this context as normal, running from 01/08/YYYY to 31/07/YYYY+1. The student does not need to be on a degree programme run by our school, so students taking our courses but not on our programmes will also have a "module-COURSE" role. Only courses run by our School will result in a role, other courses a student may be registered on outside our School will not result in a role.

The setting of COURSE is entirely at the discretion of our administrative staff. Usually it will be a locally understood abbreviation for the course, for example "Accelerated Natural Language Processing" would be "anlp". New courses do not necessarily get this configured straight away in which case the default will be the lower case version of the EUCLID code, for example something like "infr11125", until changed. It is likely that over time we will start to see more EUCLID encodings and fewer local abbreviations because the local abbreviations will only be maintained by ITO administrative staff while they have a local requirement and are engaged to do so - as more courses are managed through Learn (including submission of work), more distance learning variations of courses become available, all assessment is done upstream in EUCLID and students/staff themselves are more familiar with the EUCLID code than our own abbreviations the local requirement to have our own naming scheme will reduce/go and and there will be little point in us maintaining local course names ourselves which will have no meaning other than to us.

Teaching Support Roles

Multiple types of role are produced specifically for teaching support duties. These are all of the form "DUTY-COURSE". The COURSE part of the role is set according to the corresponding course's tag value which is managed by administrative staff and outside our control, see above for the "module-COURSE" role. The DUTY part of the role is one of "demonstrator", "marker", "ta", "tutor". Any other duties are not specifically represented. Any person in any category, staff, visitor and student (although really only UG4, UG5, PGT and PGR) can have any number of these roles depending on their allocated teaching duties for courses.

- The "DUTY-COURSE" role is given to any record in Theon that has a currency of "Existing" or "Upcoming" within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. They must have been allocated one of the four defined duties on a course, their allocation must have been approved and must be "current" or "upcoming". The "current" academic session is defined in this context as normal, running from 01/08/YYYY to 31/07/YYYY+1. Upcoming records are included to handle duties being assigned to people not started yet, or where duties are being transferred between people mid-session.

Membership Roles

Multiple types of role are produced specifically for organisational unit membership. These are all of the form "UNIT-member". The UNIT part of the role is set according to the corresponding unit's tag value which is managed by administrative staff and outside our control, see below.

- The "UNIT-member" role is given to any record in Theon that has a currency of "Existing" or "Upcoming" within the "current" academic session. If they have an end date against their record it must be greater than or equal to the current date. They must not have been marked as deleted by the feed processes. They must have been assigned as a member of the UNIT, the current date must fall between the start date and end date of their allocation and the assignment must not have been marked as deleted by the feed processes (where applicable). In the case where the UNIT is an Institute then the membership type must be one of "Affiliate", "Member", "Visitor" otherwise any membership type would be valid. The "current" academic session is defined in this context as normal, running from 01/08/YYYY to 31/07/YYYY+1. Upcoming records are included to handle affiliations being assigned to people not started yet.

The setting of UNIT is entirely at the discretion of our administrative staff. Usually it will be a locally understood abbreviation for the unit, for example "Laboratory for Foundations of Computer Science" would be "lfcs" and "Informatics Teaching Organisation" would be "ito". We do not necessarily in advance when new units are created. Student research institute membership assignments are largely made automatically, based on the association of the degree programme, however they can also be made manually. Other membership assignments are made manually by administrative staff and outside our control.

New Roles

We want to look at the possibility of adding some more finer grained roles that Theon could easily generate so as to reduce the number of "secondary roles" that need to be manually applied through Prometheus. One issue here is role naming as we would be adding a lot of new roles that don't necessarily have rfe/unix/ldap friendly short names, and we either have to just accept that, use an identifier other than the name which will probably be obscure, or commit to maintaining short name mappings ourselves within Theon. For example, access to parts of the [TheonPortal](#) could be managed by entitlements granted through user roles if they were fine-grained enough. Access to areas of [TheonPortal](#) are currently constrained by access files generated directly from Theon. These are based on rules like the authenticated person must be allocated an administrative duty such as "Director of Graduate School" or the authenticated person must have an official staff position title such as "Institute Portfolio Manager" or the person must have an allocated teaching duty such as "Lecturer".

Upstream Feed Handling

As previously noted we have no formal definition of the feed data we receive from upstream. Consequently it is impossible to guarantee the limits in what to expect when that content transfers through Theon and into the views defined above. In particular the timing of key events at registry we know by first hand experience only, they could change (and have changed in the past) without our knowledge and we are highly unlikely to trap this until after it happens and has propagated through Theon into Prometheus. So we need formal procedures to deal with (clean up after) "anomalies" as we cannot predict nor prevent them.

Timing

The following are approximate upstream timings of "relevant" feeds as they roll into Theon.

- IDM dataset is fetched and processed at 0230
- EUGEX course data is fetched and processed at 0400
- EUGEX student data is fetched at 0400 but not explicitly processed
- BIS dataset updates (including VRS dataset) and processing occur between 0400 and 0530
- Internal "session" processing occurs between 0400 and 0430
- Internal "spawning" processing occurs between 0430 and 0500
- EUGEX student data is processed at 0745
- HR staff and TSP contract data is processed between 0830 and 0900
- Internal "spawning" re-processing occurs at 0900

The processing of any single feed is atomic, that is a failure will roll back the entire update including all data sync processes and triggered functional processing. However the whole span of overnight processing is not atomic and it is possible for Theon to have an inconsistent state which would be reflected in an erroneous "status" or "currency" on person records. This clearly then affects the content of the Prometheus views as they are constructed around these values. Within the window defined by the timing of feed updates above Prometheus should not therefore access (or process any changes at least) from the views. So from above, between 0230 and approximately 0915, many changes are happening and downstream data can be considered at-risk.

Issues

There are known issues with the VRS dataset, reconciled UUNs will not appear. This is worked around in Theon but not always successfully.

The HR/TSP datasets are run from the daily snapshot which reflects the exact contract status on that day. However, the daily snapshot is taken at the end of the day and only available to be retrieved the following day. Consequently new staff only appear in this feed 24 hours after their contract start date. Staff transfers from other parts of the University can also be problematic because, irrespective of the start date, the actual School/Unit affiliation transfers in the HR data are set on monthly boundaries (related to the payroll cycle), so these can appear even later. Administrative staff work around these issues locally in Theon by pre-registering staff which then appear downstream under the "LocalVisitorStaff" category (i.e. note **not** with staff status) until they finally appear in the feed, however we have in the past had subsequent reconciliation issues when they do appear in the feed - these seem to largely be resolved.

We only fetch student data once in the morning as that used to reflect the upstream update cycle which tipped EUCLID data into BIS. However this now happens hourly upstream so we could increase our own update frequency. We don't generally have any call to do this (although it is sometimes done manually at extra times during the day at certain key points of the year), but if we did it would have some consequence for how we stage Prometheus updating data in order to avoid inconsistencies (e.g. people dropping out of the view and then re-appearing on the next refresh from the view). A small buffer on "deletions" may be recommended prior to lifecycle kicking in so as to suppress the bulk of these brief transient events, however this would be difficult to implement on the Prometheus or Theon side, as it would mean maintaining state more than is currently done.

Big Churn Events

Most daily changes are small, but there are points in the year relating to the student data set that will result in a large number of simultaneous changes that it is worth knowing about.

- 1st August - formal academic session rollover internal to Theon
- Early August - upstream rollover, undergraduate year of study change, new session unconditional confirmed offer students appear
- 1st September - bulk research student programme start date (although smaller numbers of PGR students start throughout the session)

- Mid September - bulk taught student programme start date
- 1st September - internal Theon "extended" session period ends (covering previous session resit period)
- End September - last working day in September and registry changes will drop the previous session PGT students from our primary feed
- 1st October - internal Theon "overrun" session period ends (covering previous session PGT project and award assessment period)
- Shortly after a semester ends - visiting student end dates
- Early June - undergraduate awarded (4th/5th year) student end dates

UUN Change

At the discretion of administrative staff a UUN can be moved from a "staff" record for a person to a "visitor" record for the same person and vice-versa. In Theon and in the Prometheus view these are treated as different people (distinct UUID value). Similarly a "student" UUN may move to a teaching support "staff" record. A UUN can only exist once within Theon, so be aware of the following scenario example. User X has a staff record with UUID=1 and UUN="abc" but leaves and will get a visitor record with UUID=2 and UUN="v1abc". This will (depending on overlap timing) result in two different user registrations in the Prometheus view. If at any point the visitor record is reconciled back to "abc" then the administrative staff will update the visitor record and replace "v1abc" with "abc". This will change the username for the UUID=2 user registration in the Prometheus feed from "v1abc" to "abc" and the UUID=1 user registration will have an empty username in the Prometheus view (to avoid UUN conflict the "abc" is automatically removed from the staff record when the visitor record is updated).

-- [TimColles](#) - 14 Dec 2016

Topic revision: r8 - 29 Sep 2017 - 09:21:56 - [TimColles](#)

Copyright © by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
Ideas, requests, problems regarding TWiki? [Send feedback](#)
This Wiki uses [Cookies](#)

