



# Enterprise Architecture HUIT Technology Partner Services

## IT Resource Standard for Enterprise Tags

<b>Authors:</b> <ul style="list-style-type: none"><li>Sevier, Raoul</li></ul>	<b>Audience Level:</b> <ul style="list-style-type: none"><li>IT Director / Manager</li><li>Solution Architect and Project Manager</li><li>Application Developer and Designer</li><li>DevOps Staff</li><li>Senior IT Engineers</li></ul>
<b>Version:</b> 1.0 <b>Last Revised:</b> 11/30/2020 7:22:00 PM <b>Status:</b> Released <b>Document Type:</b> Single Topic Guidance	<b>Distribution Scope:</b> <ul style="list-style-type: none"><li>HUIT</li><li>Schools/Units needing HUIT resource management</li></ul>
<b>Workgroup Members:</b> <ul style="list-style-type: none"><li>Fanton, Joel – Automation</li><li>Hall, Vicki – Network</li><li>Hartman, Scott - Network</li><li>Kingsview, Sebeyon – AWS, Azure</li><li>McGowan, Patrick – VMWare</li><li>Phelan, Anna - CMDB</li></ul>	<b>Reviewers:</b> <ul style="list-style-type: none"><li>Burson, Jefferson – EA</li><li>Charest, Greg – EA</li><li>Eskridge , Myra Hope - CMDB</li><li>LaPorte, David – Network</li><li>Larsen, Tamara – DevOps</li><li>Lazri, Cody – IAM</li><li>Rota, Benjamin – Cloud</li><li>Vaverchak, Tim – Shared Services</li></ul>

## IT Resource Enterprise Tag Standard

### Document Change Log

<b>Rev</b>	<b>Date</b>	<b>Editor</b>	<b>Changes</b>
0.1	01-Jul-2020	Raoul Sevier	Initial Draft
0.2	02-Jul-2020	Greg Charest	Comments
0.3	02-Jul-2020	Greg Charest	Revised to 0.3
0.4	08-Jul-2020	Raoul Sevier	Edits to body
0.5	09-Jul-2020	Raoul Sevier	Edits to table
0.6	20-Jul-2020	Raoul Sevier	Incorporated RFC comments
0.61	23-Jul-2020	Raoul Sevier	Added Mandatory / Optional / n/a table
0.62	31-Jul-2020	Raoul Sevier	Clarified content and format of Exception tag
0.90	13-Aug-2020	Raoul Sevier	Finalized lists of Enterprise, Pending, and Other tags
0.91	18-Aug-2020	Raoul Sevier	Minor Edits
0.92	20-Aug-2020	Raoul Sevier	Edits to 'Created By' and 'Resource Name'
1.0	23-Nov-2020	Raoul Sevier	Updates to the initial tag list, creation of side-car spreadsheet

## Table of Contents

1. Problem Statement .....	4
2. Recommendations .....	4
3. Discussion .....	4
3.1. The role of Tags in HUIT .....	4
3.2. Standardization of Selected Tags for Enterprise use .....	5
3.3. Criteria for Enterprise Tags .....	5
3.4. Mapping of Enterprise Tags Within Resource Domains .....	5
3.5. Propagation of Tag Information to the Enterprise CMDB .....	6
3.6. Use of Automation for Effectiveness and Efficiency .....	6
3.7. Mandatory vs. Optional Enterprise Tags .....	6
3.8. Tagging Practices at Harvard .....	7
4. Updates, Exceptions, and Waivers .....	7
4.1. Updates to these Standards .....	7
4.2. Waivers and Exceptions from these Standards .....	7
5. Standard Enterprise Tag Names .....	8
5.1. Enterprise Tag List .....	8
5.2. Enterprise Tags Pending Implementation .....	9
5.3. Other Candidate Tags for Inclusion in the Standard .....	10
5.4. Deprecated tags .....	11
6. Usage Standard .....	12
6.1. Usage for Enterprise Tags .....	12
6.2. Usage for Enterprise Tags Pending Implementation .....	13
7. References .....	14

## 1. Problem Statement

Effective management of IT resources requires that various attributes describing the resources be available for cost allocation, cost optimization, reporting, compliance, and security purposes. These attributes must be consistently defined and available across organizational boundaries.

The general concept of tagging resources applies to cloud and on-premise infrastructure in the server, network, storage and application domains. Cloud based computing environments offer standard tagging mechanisms, but nomenclatures and availability vary by vendor.

The lack of a consistent set of enterprise level meta-data on IT resources across domains and organizational boundaries currently limits HUIT's ability to address the above requirements.

## 2. Recommendations

- Define a limited set of tags to be 'Enterprise' in scope and specify the naming, meaning and allowed content of the tags
- Identify a group with the responsibility to maintain and evolve the set of enterprise tags and approve exceptions as appropriate
- Implement automated mapping or translation processes if needed to convert tags from existing domain systems to the enterprise specification
- Require all HUIT teams to maintain accurate inventories of IT resources and their appropriate Enterprise tags, in the central Configuration Management Database (CMDB).
- Encourage all HUIT IT teams to provide the CMDB with additional information about resources which may be deemed important, but without the constraint of consistent naming or value formats.
- Inform Harvard School IT partners of HUIT's enterprise tagging policies and processes and use the Cloud Community of Practice to begin an effort to align tagging efforts across the University.

## 3. Discussion

### 3.1. The role of Tags in HUIT

HUIT, like many IT organizations, is largely organized by IT technical domains. Platform and server teams manage compute resources, network organizations own and operate switches and access points, backup teams manage the resources necessary to store and protect data.

Each organization's management of the technologies under their purview is augmented by the use of meta data or tags which hold attribute information about the resources they manage. Examples of tag information include the environment (development, stage, production), owner, cost center as well as technical data.

Many of the systems, although robust, expect users to conform to specific design and usage conventions limiting their cross-domain value. Differences in the implementation and design of these individual tagging systems currently limit their use in answering broad cross organization questions

## IT Resource Enterprise Tag Standard

around usage, cost, dependencies and security. For example, the AWS platform team may use a tag named 'Environment' with a value of 'Production', while the Azure team may use a tag named 'State' with a value of 'P-1'.

When IT management review IT resources across different organizations and technical domains, they are faced with inconsistent tag usage. This presents itself both in the names of the tags used by different tools and organizations, as well as the allowed formats and values in those tags. These inconsistencies make aggregate reporting and analysis very difficult.

### 3.2. Standardization of Selected Tags for Enterprise use

Requiring adherence to a single monolithic tagging nomenclature is inadvisable. Many of the existing individual systems have developed over time, provide specific domain functionality and cannot be easily changed. Much of the information about IT resources in a technical domain is important, if not essential, to the successful management of the resources by the operators of the domain. For example, AWS administrators need to know the configuration of elastic load-balancers, which style of RDS database is used by an application, and if the virtual server is used for production operation. These systems provide significant value and should continue to operate. Nevertheless, there is a clear need for a relatively small number of tags to be 'Enterprise' in scope, with consistent naming conventions, standard content and meaning.

The primary tool for managing IT service-related data in an ITSM environment is the Configuration Management Data Base (CMDB). The existing ServiceNow based CMDB is the logical container for managing the meta-data held in enterprise level tags as well additional information about resources which may be deemed important but does not require consistent naming or value formats.

### 3.3. Criteria for Enterprise Tags

From an enterprise management perspective, the non-technical and/or operational attributes of a resource are often the focus of interest.

Criteria for selecting Enterprise Tags may include:

- Data necessary to understand:
  - Usage
  - State
  - Ownership
  - Cost allocation
  - Service and Applications support by the resource
- Attribute information that supports operational requirements such as backups and patching
- Attribute information that describes and supports automated resource creation and management

These criteria, and others, should be considered when determining the need for an Enterprise Tag. As a practical matter, this process will start organically and stabilize over time.

### 3.4. Mapping of Enterprise Tags Within Resource Domains

## IT Resource Enterprise Tag Standard

The names and values of tags within each existing resource domain may, or may not, currently conform to the naming of the attribute in the Enterprise CMDB, nor may the values align to the standard created for enterprise level tags.

Whenever possible, existing tags nomenclatures should be aligned to the standard for enterprise level tags in order to minimize the number of variant names and values of the same attribute across different technical domains. However, as discussed above, this may not always be practical. In these cases, an automated translation or mapping process can be implemented in order to convert the existing local meta-data into the enterprise tag form and content.

### 3.5. Propagation of Tag Information to the Enterprise CMDB

Where enterprise management identifies attributes in local resource domains that must be propagated to the Enterprise CMDB, it is up to the organization owning the resource domain to craft a means of providing the required information.

The HUIT organization has procured several tools, such as CloudAware and LogicMonitor, that do resource discovery by looking into platforms, such as AWS and VMWare, and identifying IT resources. This information is aggregated into the tool's internal database. This is a convenient way of acquiring information about resources. Further these tools often have pre-built connectors that populate Configuration Management tools such as ServiceNow's CMDB.

When using tools such as these, it is up to the local resource domain team to ensure the Enterprise Tags and values conform to the Enterprise Tag standard as the information is moved to the CMDB.

In other situations, there may not be an existing tool that can provide discovery or propagation to a CMDB. In these cases, custom procedures or code may be needed to provision the CMDB with the requisite information.

### 3.6. Use of Automation for Effectiveness and Efficiency

A central goal of the Enterprise Tag Standard process is the use of automation to the greatest degree possible.

In principle, tagging of IT resources should take place at the time the resource is created. In the case of resources that are created by Ansible Tower or by SCCM, the scripts used to create the resources should contain the directives that create the tags in the appropriate platforms. For example, an Ansible script to create an AWS virtual machine should also create the tags that identify the server as a production machine using the correct Enterprise Tag name and value.

Where automated resource management is not available, in principle the tag names and values should be defined as early in the life-cycle as possible, and only once.

### 3.7. Mandatory vs. Optional Enterprise Tags

Given the great variety of resource domains, it is a given that not all IT resources will be able to support all Enterprise Tags.

**It is up to the organization managing the local resource domain and the team responsible for enterprise tag management to agree which tags will be required for each domain. Once agreed, then the resource domain must provide the agreed tag information in the standard name with standard values.**

### 3.8. Tagging Practices at Harvard

While initially scoped to the HUIT organizations, HUIT routinely manages IT resources belonging to Schools and other organizations. As a practical matter, any IT resources that fall under the management of HUIT organizations must conform to the tagging standard. As HUIT interacts with Schools and other organizations, they must inform them of this approach and take steps to ensure that School and other organization resource information makes its way to the Enterprise CMDB.

**The Cloud Community of Practice can be used as an initial vehicle for discussion of and alignment of enterprise tags with HUIT's University partners.**

## 4. Updates, Exceptions, and Waivers

IT environments undergo continuous change. As a practical matter, it is important to manage that changes with as much automation as possible to maximize both effectiveness and efficiency of IT operations. This means updating the standards as the mix of resources change. Just as important as knowing what a resource should be tagged, is a sense of where exceptions are important, and an inventory of waivers to the standards with the reasons the waivers were given.

### 4.1. Updates to these Standards

**The responsibility for tasks related to maintaining and updating enterprise tags should be clearly defined. Should the scope of change be large enough, an additional round of peer and management review may be required. This material and updates will be cross-published on the EA web site and in the HUIT TPS Confluence wiki site.**

### 4.2. Waivers and Exceptions from these Standards

There may be some circumstances where standards have not yet been defined for a class of resource. Under these circumstances a waiver for that class can be allowed, as long as there is management concurrence. This becomes the basis for updating the standards document. The responsible organization must keep a log of waivers.

In the event there is an applicable name or tag standard for a resource, but there are compelling reasons to deviate from them, exceptions may be granted. Under these circumstances, the responsible organization must grant the exception in writing. The responsible organization must keep a log of exceptions.

## 5. Standard Enterprise Tag Names

The list of Standard Enterprise Tags identified in this document represents that essential set of information that management needs to conduct cross-organizational and cross-technical domain management.

This list is simplified for presentation purposes. A complete table with additional descriptive material can be found in the side-car Excel spreadsheet for this document.

### 5.1. Enterprise Tag List

The following table describes tags that must be used in a consistent fashion across IT organizations and technical domains.

Enterprise Tag Name	Purpose	Content
<b>Automated Build Enabled</b> CMDB name: automation_bootstrapped	Indicates that server has been configured to receive automation from Ansible Tower or SCCM.	One of: <ul style="list-style-type: none"> <li>successful ssh connection</li> <li>ssh unreachable</li> <li>generic ssh failure</li> <li>incorrect ssh key</li> <li>unknown</li> <li>windows host</li> </ul>
<b>Backup Policy</b> CMDB name: backup_policy	Indicates the backup policy in effect for the resource. Applies primarily to resource type of server.	One of: <ul style="list-style-type: none"> <li>No policy</li> <li>Daily Incremental, monthly full</li> <li>Daily incremental only</li> <li>Monthly full only</li> </ul>
<b>Created By</b> CMDB name: created_by	Denotes the Ansible Tower ID used to create Stackbuilder resources, or the SCCM ID.	One of: <ul style="list-style-type: none"> <li>&lt; AnsibleTower username&gt;</li> <li>&lt; SCCM username&gt;</li> </ul>
<b>Data Classification</b> CMDB name: data_class	Denotes the level of sensitivity of data contained or managed by the resource. This usually applies to servers and databases.	One of: <ul style="list-style-type: none"> <li>Level4</li> <li>Non-level4</li> <li>Public</li> </ul>
<b>Environment</b> CMDB name: environment	Used to indicate the operating level of a resource. Production indicates the resource is configured to full standards and policies for operational use by the intended users. Non-production indicates a (generic) lower-level environment that is less exposed, and may use lower standards and policies for operation. Other values provide a finer-grained indication of operating level.	One of: <ul style="list-style-type: none"> <li>Production</li> <li>Non-production</li> <li>Stage</li> <li>Test</li> <li>Development</li> <li>Sandbox</li> </ul>
<b>Exceptions</b> CMDB name: exception	This tag indicates that there is a quality about the resource that is non-conformant. The content of the tag is the compliance date that was given when the exception was granted.	One of: <ul style="list-style-type: none"> <li>No tag on resources that are compliant</li> <li>For resources that are non-compliant, a value of the compliance date as text in the</li> </ul>

## IT Resource Enterprise Tag Standard

		format of YYMMDD to allow collation. In the case of multiple exceptions, the earliest compliance date must be used
<b>Hosted By</b> CMDB name: hosted_by	Used to indicate the support group that is responsible for providing tier-2 level support and beyond.	One of: <ul style="list-style-type: none"> <li>• DevOps-APT12</li> <li>• DevOps-APT3</li> <li>• DevOps</li> <li>• IAM</li> <li>• Library</li> <li>• (Other values to follow)</li> </ul>
<b>HUIT Asset ID</b> CMDB name: huit_assetid	Used to associate resources with administrative codes. This enables ICAPS to divide up a single account among multiple billing codes.	<Asset ID number assigned by ICAPS>
<b>Resource Name</b> CMDB name: name	Used for people to identify the resource. The tag value should be configured to present the essence of the resource's identity, perhaps by concatenating several attributes. When a friendly name is not possible, the value should reflect the resource unique identifier.	One of: <ul style="list-style-type: none"> <li>• &lt;A friendly name for a given resource that helps with readability&gt;</li> <li>• &lt;A technical identifier from the technical domain that operates the resource&gt;</li> </ul>
<b>Resource Type</b> CMDB name: resource_type	Used to indicate the type of the resource.  This is intended to simplify the selection of resources in the CMDB and allow efficient and effective reporting and analytics.	One of: <ul style="list-style-type: none"> <li>• Appliance</li> <li>• Database</li> <li>• Security_device</li> <li>• Load_Balancer</li> <li>• Storage_Device</li> <li>• Network Device</li> <li>• Virtual_Guest</li> <li>• VPC</li> </ul>
<b>Server Platform</b> CMDB name: platform	Used to distinguish between the major operating system types and enable better visibility within configuration management.  This is primarily used with server resources.	One of: <ul style="list-style-type: none"> <li>• linux</li> <li>• windows</li> </ul>

### 5.2. Enterprise Tags Pending Implementation

The following table describes tags that will become Enterprise Tags once an implementation plan is completed.

Enterprise Tag Name	Purpose	Content
<b>AD Domain</b> CMDB name: ad_domain	Used to relate resource to AD domain, and manage the use of Active Directory to control access and permissions.	<Conforms to the domain names in UNIVAD>

## IT Resource Enterprise Tag Standard

<b>Application CI ID</b> CMDB name: app_ci_id	Associates an IT resource with an application configurable item.  The process for creation application IDs is being developed by the CMDB Workgroup.	<Encoded by CMDB>
<b>Criticality</b> CMDB name: criticality	Used to identify the criticality of the resource, typically thought of in terms of availability, but can accommodate other criteria.	One of: <ul style="list-style-type: none"> <li>• Life Safety</li> <li>• High</li> <li>• Medium</li> <li>• Low</li> </ul>
<b>Operations Policy</b> CMDB name: op_policy	Used to define operational timing and criteria.	One of: <ul style="list-style-type: none"> <li>• (Values to follow)</li> </ul>
<b>Patching Policy</b> CMDB name: patching_policy	Used to define patching timeslots.	One of: <ul style="list-style-type: none"> <li>• Zero-day + Monthly</li> <li>• Zero-day + Weekly</li> <li>• On-demand</li> <li>• (Other values to follow)</li> </ul>
<b>Product</b> CMDB name: product	Used to associate a resource with an application. The primary use-case is where an application uses a whole resource (eg: a server or a database). Another use-case is where several applications share a resource (eg: VPC or VMWare Host).  In the cases where there are a large number of applications using a resource (eg: storage array), or the nature of a shared resource does not lend itself to association with applications (eg: network switch), this tag should not be used.	Multi-valued list separated by commas, using values <Encoded by CMDB>

### 5.3. Other Candidate Tags for Inclusion in the Standard

These tags are candidates for inclusion into the list of Enterprise Tags. A Tag Review Board or similar group will determine the appropriateness for inclusion, as well as the naming construct and allowed values.

Tag	Notes and Examples
Backed Up	
Building Root	
Lifecycle	
Network priority	
Patched	
Self-managed	
Server owner	

## IT Resource Enterprise Tag Standard

### 5.4. Deprecated tags

These tag names are published on the Cloud web site, they are not intended to be used by HUIT.

<b>Deprecated Tag</b>	<b>Notes</b>
<b>owner</b>	Deprecated in the DevOps domain
<b>department</b>	Deprecated in the DevOps domain
<b>project</b>	Deprecated in the DevOps domain

## 6. Usage Standard

The list of Standard Enterprise Tags identified in this document represents that essential set of information that management needs to conduct cross-organizational and cross-technical domain management. Since not all tags are relevant to all technology domains, this section sets out the mandatory and optional tag usage per technology domain.

This list is simplified for presentation purposes. A complete table with additional descriptive material can be found in the side-car Excel spreadsheet for this document.

### 6.1. Usage for Enterprise Tags

The following table describes tags that must be used in a consistent fashion across IT organizations and technical domains. Mandatory indicates the tag must be associated with every resource within the technology domain. Optional indicates the tag is helpful when associated with the resources in the technology domain. ‘n/a’ indicates the tag is not relevant to the technology domain.

ENTERPRISE TAG	TECHNOLOGY DOMAIN									
	Enterprise Tag Name	AWS	Azure	VMWare	Network	Storage	Backup	IAM	Log-ging	Auto-mation
<b>Automated Build Enabled</b> CMDB name: automation_bootstrapped	M	M	M	n/a	n/a	n/a	M	n/a	n/a	n/a
<b>Backup Policy</b> CMDB name: backup_policy	M	M	M	n/a	O	n/a	M	n/a	n/a	n/a
<b>Created By</b> CMDB name: created_by	O	O	M	O	O	O	O	O	O	O
<b>Data Classification</b> CMDB name: data_class	M	M	M	n/a	O	O	M	O	n/a	M
<b>Environment</b> CMDB name: environment	M	M	M	M	M	M	M	M	M	M
<b>Exceptions</b> CMDB name: exception	M	M	M	M	M	M	M	M	M	M
<b>Hosted By</b> CMDB name: hosted_by	M	M	M	n/a	n/a	n/a	M	n/a	n/a	n/a
<b>HUIT Asset ID</b> CMDB name: huit_assetid	M	M	M	O	O	O	M	n/a	n/a	O
<b>Resource Name</b> CMDB name: name	O	O	M	O	O	O	O	O	O	O
<b>Resource Type</b> CMDB name: resource_type	M	M	M	M	M	M	O	M	M	M
<b>Server Platform</b> CMDB name: platform	M	M	M	n/a	n/a	n/a	M	n/a	n/a	M
<b>Legend:</b> M = Mandatory O = Optional n/a = Not Applicable										

## IT Resource Enterprise Tag Standard

### 6.2. Usage for Enterprise Tags Pending Implementation

The following table describes tags that will become Enterprise Tags once an implementation plan is completed. Mandatory indicates the tag must be associated with every resource within the technology domain. Optional indicates the tag is helpful when associated with the resources in the technology domain. 'n/a' indicates the tag is not relevant to the technology domain.

ENTERPRISE TAG	TECHNOLOGY DOMAIN									
Enterprise Tag Name	AWS	Azure	VMWare	Network	Storage	Backup	IAM	Log- ging	Auto- mation	API
<b>AD Domain</b> CMDB name: ad_domain	n/a	M	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
<b>Application CI ID</b> CMDB name: app_ci_id	M	M	M	O	O	O	n/a	n/a	n/a	n/a
<b>Criticality</b> CMDB name: criticality	M	M	M	M	M	M	M	n/a	n/a	M
<b>Operations Policy</b> CMDB name: op_policy	M	M	M	O	O	O	M	O	O	M
<b>Patching Policy</b> CMDB name: patching_policy	M	M	M	M	M	M	M	M	M	M
<b>Product</b> CMDB name: product	M	M	M	O	O	O	M	O	O	M

## 7. References

- Enterprise Tag Standard – Wiki site with additional details  
<https://wiki.harvard.edu/confluence/display/huitenterprisearchitecture/TPS+Tagging+Standards>
- Harvard Internal Tagging Standards:
  - Standard Tags available for Cost Analysis  
<https://cloud.huit.harvard.edu/aws-financial-account-management#tags>
  - HUIT Deployed Application Documentation -  
<https://confluence.huit.harvard.edu/display/CLOPS/Deployed+Application+Documentation>
  - Original Cloud tagging standard:  
<https://confluence.huit.harvard.edu/display/Cloud/Tagging#space-menu-link-content>
  - Cloud Architecture tagging standard:  
<https://confluence.huit.harvard.edu/display/CLA/Cloud+Resource+Tagging>
  - LogicMonitor Tagging Standards  
<https://confluence.huit.harvard.edu/display/ITSDevOps/ITS+DevOps+LogicMonitor+Property+%28Tagging%29+Standards+and+Examples>
  - DMS tagging standard:  
<https://wiki.harvard.edu/confluence/display/huitdmsops/Cloud+Inventory+Tagging>
- AWS Documentation - <https://aws.amazon.com/answers/account-management/aws-tagging-strategies/>
- DLT.com - <https://opscenter.dlt.com/hc/en-us/articles/235230267-AWS-Tagging-with-a-DLT-AWS-Account>
- Cornell.edu - <https://confluence.cornell.edu/display/CLOUD/Standard+Tagging>