Université Nice Sophia Antipolis
Membre de UNIVERSITÉ CÔTE D'AZUR

Numéro d'anonymat
(Partie réservée à l'administration)

CONCOURS EXTERNE

ASSISTANT INGÉNIEUR

BAP E

Emploi type : « Assistant(e) en ingénierie logicielle»

Session 2018

Épreuve écrite d'admissibilité

Mardi 5 juin 2018 de 9h à 12h

(Durée : 3 heures – coefficient 4)

AMPHITHÉÂTRE PV 01

NOM DE NAISSANCE : ……

NOM D'USAGE : ……

PRÉNOM : ……

CONSIGNES À LIRE AVANT DE COMMENCER L'ÉPREUVE

- **Le sujet qui vous a été remis comporte une partie exercices (pages 1 à 7) et une partie annexes (pages 1 à 18)**.

- Veillez à ce que cet exemplaire soit complet sinon demandez-en un autre aux surveillants de salle.

- Le sujet peut comporter plusieurs types de questions : questions à choix multiples, questions à réponse courte, questions réponse à développer, exercices et cas pratiques

- **Les réponses sont à apporter sur les copies qui vous ont été distribuées pour les exercices 1 à 4 (questions 4.1 et 4.2) et 5.**

- **Les réponses sont à apporter directement sur le sujet qui vous a été remis pour la partie QCM (exercice 4, questions 4.3.1 à 4.3.5)**

- **Attention : le sujet ne doit pas être dégrafé et devra être remis dans son intégralité aux surveillants à l'issue de l'épreuve.**

- Il vous est demandé d'écrire soigneusement, en bleu ou en noir uniquement et de **ne pas utiliser de crayon à papier**.

- Les questions à choix multiples nécessitent de cocher la ou les bonne(s) réponse(s)

- L'usage de tout ouvrage de référence, de tout document autre que ceux qui vous seront remis lors de l'épreuve ainsi que l'usage de tout matériel électronique est interdit. L'usage de la calculatrice simple est autorisé.

- Les téléphones portables sont éteints pendant toute la durée de l'épreuve, et rangés dans le sac.

- Veillez à respecter l'anonymat dans vos réponses. Il vous est rappelé que **votre identité ne doit figurer que dans la partie supérieure de la bande entête de la première page des copies mises à disposition pour composer et dans la partie inférieure de la 1ère page du sujet.**
  Toute mention d'identité, quelle qu'elle soit, portée sur tout ou partie de la copie que vous remettrez en fin d'épreuve (signature, paraphe, initiales, mention d'identité…) est un signe distinctif pouvant identifier la provenance de la copie et mènera à l'annulation de votre épreuve.

EXERCICES

Exercice 1 : (5 points)

Requêtes SQL

Soient les trois tables suivantes :

Films

idFilm	Titre	Réalisateur
1	Oscar	Réalisateur d'Oscar
2	La soupe aux choux	Réalisateur de la soupe aux choux

Acteurs

idActeur	Nom	Prénom
1	Dupont	Scarlett
2	Durand	James
3	Martin	Georges

Role

idActeur	idFilm	Personnage
1	1	Dory
2	1	Bruce
3	2	Némo

A partir de ces tables, nous vous proposons quelques exercices. Proposer les requêtes pour les cas suivants :

1.1 Ecrire la requête qui permet de créer la table Role
Sachant que :
idFilm et idActeur sont des clés primaire dont on laisse le serveur de base de données donner une valeur (AUTO_INCREMENT)
Les champs texte sont de type VARCHAR (chaînes de caractères de taille variable) de taille maximale de 255 caractères.

1.2 Comment insérer les données dans cette table selon l'exemple ci-dessus.

1.3 Remplacer « Dory » par « Dora »

1.4 Compter le nombre de films.

1.5 Afficher l'ensemble des titres, nom de l'acteur, prénom de l'acteur et rôle joué.

Exercice 2 : (6 points)

Installation d'un WIKI

Une université qui compte 30 000 étudiants et 5 000 personnels souhaite disposer d'un wiki à destination de tous les usagers.
Pour cela, il est décidé d'installer l'application Confluence dont vous trouverez la documentation en annexe.

2.1 Vous avez en charge de :

2.1.1 Dimensionner le serveur qui hébergera l'application (justifiez votre réponse).
2.1.2 Choisir le système d'exploitation (justifiez votre réponse).

2.2 Vous décrirez les prérequis et les principales étapes à réaliser avant de procéder à l'installation de Confluence.

2.3 Vous décrirez la procédure d'installation que vous avez choisie ainsi que les éléments de configuration à finaliser après l'installation.

2.4 Vous détaillerez également les éléments qui vous semblent nécessaire à la sécurisation et la sauvegarde du serveur (détaillez et expliquez vos choix).

Exercice 3 : (5 points)

Une analyse du fichier /var/monappli/log_data est lancée chaque nuit à 22h30, sur le serveur linux de production.

3.1 Ecrire la ligne du crontab correspondant au lancement du job check_log_data.sh

3.2 Afin de sauvegarder les données, proposer un script shell permettant
(a) d'éteindre la base de données proprement
(b) de faire la copie physique des fichiers dans un dossier /nfs,
(c) précisant date:heure de début et fin de traitement

Le nom du répertoire de la sauvegarde est ainsi composé :
/nfs/Nomdelabase_dateetheuredelasauvegarde
Exemple : /nfs/mabase_20170210_04:02:30
Nom de la base : mabase
/datab/idx/mabase : contient tous les fichiers index
/datab/dbf/mabase : contient tous les fichiers de la bd
/datab/div/mabase : contient des fichiers temporaires à effacer

Proposer la ligne du crontab permettant de lancer votre script toutes les 4 heures du lundi au vendredi inclus et d'envoyer les logs de sortie dans un fichier « logCron »

Reportez-vous aux annexes "crontab", "Date", et "Arrêt de la base"

Exercice 4 Programmation : (7 points)

4.1 Soient deux variables : 'A' qui contient la valeur 5 et 'B' qui contient la valeur 3

Ecrire en javascript un programme qui inverse et affiche les valeurs des variables

4.2 Indiquez pour chacun des extraits de code ci-dessous à quels langages de programmation il appartient et ce qu'il fait.

a)

```
int i = 1;
for (i; i <= 10; i++)
{
    printf("%s", "Hello, world !\n") ;
}
for ($a = 0; $a < 11; $a++)
{
    echo 'Un message';
}
```

b)

```
<!DOCTYPE html>
<html>
<body>

<h2>What Can we Do?</h2>

<p id="demo">We can change content.</p>

<button type="button"
onclick='document.getElementById("demo").innerHTML = "Hello
world!"'>Click Me!</button>

</body>
</html>
```

4.3 QCM

Réponse juste : 1 Point
Non répondu : 0 point
Réponse fausse : -0,5 point

4.3.1 Un utilisateur vous signale que son programme échoue avec le message suivant :

```
$ ./prog.py
-bash: ./prog.py: Permission denied
$
```

Que peut-on dire sur les causes de l'échec :

- □ Rien, il faut aller regarder dans le code source ou interroger l'utilisateur.
- □ Le programme a tenté d'utiliser des ressources qui lui sont interdites.
- □ Le fichier `prog.py` n'est pas un fichier exécutable.

4.3.2 L'exécution d'un programme (créé sur la même machine) échoue avec le message:

```
$ ./prog
./a.out: error while loading shared libraries:
libboost_mpi.so.1.65.1: cannot open shared object file: No
such file or directory
$
```

Quelles sont les causes possibles du problème ?

- □ La librairie n'a pas été installée sur la machine.
- □ Le répertoire contenant la librairie n'a pas été indiqué lors de l'édition de liens.
- □ L'environnement d'exécution ne sait pas où trouver la librairie.

4.3.3 L'exécution d'un programme échoue avec le message :

```
$ ./prog1.py
Traceback (most recent call last):
  File "./prog1.py", line 3, in <module>
    import numpy as np
ImportError: No module named numpy
```

Quelles sont les causes possibles du problème ?

- □ On ne peut en savoir plus sans regarder le script
- □ Il y a un problème dans le module numpy, il faut poursuivre les recherches dans ce module.
- □ Le module numpy n'a pu être trouvé.

4.3.4 Sur les machine de type Unix/Linux à quoi peut servir la variable d'environnement LD_LIBRARY_PATH ?

- ☐ Elle indique où trouver les bibliothèques dynamiques/partagées au programme.
- ☐ Elle indique dans quels répertoires trouver les documents, notamment les pages de manuels.
- ☐ Elle indique le chemin du chargeur de librairies.

4.3.5 Sur les plateformes de type Unix/Linux, certains fichiers ont des premières lignes commençant par #! tels que :

```
#!/usr/bin/env bash
#!/usr/bin/python3
#!/usr/bin/env perl
#!/usr/bin/env sh
```

À quoi sert cette ligne ?

- ☐ Comme toutes les lignes commençant par #, il s'agît d'un commentaire destiné à renseigner le lecteur sur le type de fichier qu'il est en train de lire.
- ☐ Elle peut indiquer au système qu'il s'agît d'un script à interpréter.
- ☐ Elle peut indiquer au système quel programme appeler pour exécuter ce fichier.

Exercice 5 Questions diverses (9 points)

5.1 Décrire de façon précise ce qu'est un web service ?
5.2 Expliquez les principales différences entre SOAP et REST?
5.3 Qu'appelle-t-on un système scalable, résiliant?
5.4 Ci-dessous une liste de démons. Indiquer très brièvement leur rôle ?

atd :
acpid :
cupsd :
hcid :
smartd :
mountd:
lpd :

5.5 Quelles sont les principales différences entre un langage compilé et un langage interprété ?

5.6 Citez des exemples de chacun de ces deux types de langage.
5.7 Qu'est-ce que le RGPD ?
5.8 A qui le RGPD s'applique-t-il ?
5.9 Quelle est la différence entre une page web dynamique et une page web statique ?
5.10 Qu'est-ce que le CSS et à quoi cela sert ?
5.11 Qu'est-ce qu'un « logiciel de gestion de versions » (ou Version Control System) ?
5.12 Quels sont les fonctionnalités apportées par cet outil ?
5.13 Dans quel cadre, l'utiliseriez-vous ?
5.14 Citez le nom de deux de ces logiciels.
5.15 Qu'est-ce que l'intégration continue ?
5.16 Citez des outils permettant de la mettre en place au sein d'une équipe de développeurs.
5.17 En gestion de projet, deux types de modèles conceptuels sont le « cycle en V » et les méthodes agiles. Donnez rapidement les différences majeures entre ces deux modèles.
5.18 Qu'est-ce que Docker ?
5.19 Quel en est le principe de fonctionnement ?
5.20 Citez un exemple d'utilisation de cette technologie.
5.21 Qu'est que ITIL ?
5.22 Décrire succinctement la différence entre MOA et MOE
5.23 Quel est le rôle du fichier .htaccess pour les serveurs web apache ?
5.24 A quoi sert le cahier des charges ?
5.25 A quoi sert un cahier de tests ?
5.26 Qui sont les acteurs autour du cahier de test ?
5.27 Qu'est-ce que le DPO, quel est son rôle ?

ANNEXES

Installing Confluence on Windows

In this guide we'll run you through installing Confluence in a production environment, with an external database, using the Windows installer.

Before you begin

Before you install Confluence, there's a few questions you need to answer.

Are you using a supported operating system?	• We don't support installing Confluence on OSX. • The Confluence installer includes Java (JRE) and Tomcat, so you don't need to install these seperately.
Are ports 8090 and 8091 available?	Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port. Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete.
Is your database set up and ready to use?	To run Confluence in production you'll need an external database. • Set up your database before you begin. Step-by-step guides are available for PostgreSQL, Oracle, MySQL, and SQL Server. • If you're using Oracle or MySQL you'll need to download the driver for your database. • To use a datasource see Configuring a datasource connection as there are some steps you need to perform before running the setup wizard. • The embedded H2 database can be used for evaluating Confluence, but you'll need to migrate to another database before running in production. You may find it easier to use external database from the start.
Do you have a Confluence license?	You'll need a valid Confluence Server license to use Confluence. • If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup. • If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup. • If you're migrating from Confluence Cloud, you'll need a new license.

Install Confluence

1. Download Confluence

Download the installer for your operating system
- https://www.atlassian.com/software/confluence/download

2. Run the installer

a) Run the installer. We recommend using a Windows administrator account.
b) Follow the prompts to install Confluence. You'll be asked for the following info:

1. **Destination directory** – this is where Confluence will be installed.
2. **H ome directory** – this is where Confluence data like logs, search indexes and files will be stored.
3. **TCP ports** – these are the HTTP connector port and control port Confluence will run on. Stick with the default unless you're running another application on the same port.
4. **Install as service** – this option is only available if you ran the installer as administrator.

c) Confluence will start up in your browser once installation is complete.

Set up Confluence

3. Choose installation type

a) Choose **Production installation**.
b) Choose any **add-ons** you'd also like to install.

4. Enter your license

Follow the prompts to log in to my.atlassian.com to retrieve your license, or enter a license key.

5. Connect to your database

a) If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
b) Choose **My own database** then select your particular database from the **Database type**dropdown menu.
c) For MySQL and Oracle, follow the prompts to download and install the required driver.
d) Enter your database details. Use **test connection** to check your database is set up correctly.

6. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

a) You can only import sites from the **same** or **earlier** Confluence version.
b) The system administrator account and all other user data and content will be imported from your previous installation.**In the setup wizard:**
c) **Upload a backup file** – use this option if your site export file is small (25mb or less).
d) **Restore a backup file from the file system** – use this option if your backup file is large. Drop the file into your `<confluence-home>/restore` directory then follow the prompts to restore the backup.
e) **Build Index** – we'll need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later. Your content won't be searchable until the index is built.

7. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Desk

8. Create your administrator account

a) Enter details for the administrator account.
b) Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

9. Start using Confluence

That's it! Your Confluence site is accessible from a URL like this:
http://<computer_name_or_IP_address>:<port>

Installing Confluence on Linux from Archive File

In this guide we'll run you through installing Confluence in a production environment, with an external database, manually using a zip file.

This method gives you the most control over the installation process.

Before you begin

Before you install Confluence, there are a few questions you need to answer.

Are you using a supported operating system?	• We don't support installing Confluence on OS X or mac OS for production environments. • **Confluence can't run on OpenJDK**. You'll need to install Oracle Java. • You can use either the JDK (Java Development Kit) or JRE (Java Runtime Environment). • We only support the version of Apache Tomcat that is bundled with Confluence.
Are ports 8090 and 8091 available?	Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port. Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete.
Is your database set up and ready to use?	To run Confluence in production you'll need an external database. • Set up your database before you begin. Step-by-step guides are available for PostgreSQL, Oracle, MySQL, and SQL Server. • If you're using Oracle or MySQL you'll need to download the driver for your database. • To use a datasource see Configuring a datasource connection as there are some steps you need to perform before running the setup wizard. • The embedded H2 database can be used for evaluating Confluence, but you'll need to migrate to another database before running in production. You may find it easier to use external database from the start.

Is your JAVA_HOME variable set correctly?	Before you install Confluence, check that you're running a supported Java version and that the `JAVA_HOME` environment variable is set correctly.
	
	Confluence can only run with Oracle JDK or JRE.
	
	To check your Java version:
	`$ java -version`
	
	To check your JAVA_HOME variable is set correctly:
	`$ echo $JAVA_HOME`
	
	If you see a path to your Java installation directory, the `JAVA_Homeenvironment` variable has been set correctly. If a path is not returned you'll need to set your `JAVA_HOME` environment variable manually before installing Confluence.

Have you created a dedicated user to run Confluence?	We strongly recommend running Confluence as a dedicated user. You should create this user before you begin, so that when creating the installation and home directories, you can give this user appropriate read and write permissions. In this example, we'll create a user called `confluence`: `$ sudo /usr/sbin/useradd --create-home --comment "Account for running Confluence" --shell /bin/bash confluence`
Do you have a Confluence license?	You'll need a valid Confluence Server license to use Confluence. • If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup. • If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup. • If you're migrating from Confluence Cloud, you'll need a new license.

Install Confluence

1. Download Confluence

Download the `tar.gz` file for your operating system
- https://www.atlassian.com/software/confluence/download.

2. Create the installation directory

a) Create your installation directory – this is where Confluence will be installed. Avoid using spaces or special characters in the path. We'll refer to this directory as your `<installation-directory>`.
b) Extract the Confluence `tar.gz` file to your `<installation-directory>`. We recommend using a GNU version of the archive utility, especially on Solaris.
c) Give your dedicated Confluence user read, write and execute permission to your `<installation-directory>`.

3. Create the home directory

a) Create your home directory – this is where Confluence application data like logs, search indexes and files will be stored. This should be separate to your installation directory, with no spaces or special characters in the path. We'll refer to this directory as your `<home-directory>`.
b) Give your dedicated Confluence user read, write and execute permissions to the `<home-directory>`.

c) Edit `<installation-directory>\confluence\WEB-INF\classes\confluence-init.properties`.
d) At the bottom of the file, enter the absolute path to your `<home-directory>`. This tells Confluence where to find your `<home-directory>` when it starts up.

4. Check the ports

By default Confluence listens on port `8090`. If you have another application running on your server that uses the same ports, you'll need to tell Confluence to use a different port.

5. Start Confluence

a) Run `<installation-directory>/bin/start-confluence.sh` to start the setup process.
b) Go to `http://localhost:8090/` to launch Confluence in your browser (change the port if you've updated the Connector port).

Set up Confluence

6. Choose installation type

a) Choose **Production installation**.
b) Choose any **add-ons** you'd also like to install.

7. Enter your license

Follow the prompts to log in to my.atlassian.com to retrieve your license, or enter a license key.

8. Connect to your database

a) If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
b) Choose **My own database** then select your particular database from the **Database type**dropdown menu.
c) For MySQL and Oracle, follow the prompts to download and install the required driver.
d) Enter your database details. Use **test connection** to check your database is set up correctly.

9. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

a) You can only import sites from the **same** or **earlier** Confluence version.

b) The system administrator account and all other user data and content will be imported from your previous installation.
c) **Upload a backup file** – use this option if your site export file is small (25mb or less).
d) **Restore a backup file from the file system** – use this option if your backup file is large. Drop the file into your `<confluence-home>/restore` directory then follow the prompts to restore the backup.
e) **Build Index** – we'll need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later. Your content won't be searchable until the index is built.

10. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Desk

11. Create your administrator account

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

12. Start using Confluence

That's it! Your Confluence site is accessible from a URL like this:
http://<computer_name_or_IP_address>:<port>

Server Hardware Requirements Guide

Server administrators can use this guide in combination with the free Confluence trial period to evaluate their server hardware requirements. Because server load is difficult to predict, live testing is the best way to determine what hardware a Confluence instance will require in production.

Peak visitors are the maximum number of browsers simultaneously making requests to access or update pages in Confluence. Visitors are counted from their first page request until the connection is closed and if public access is enabled, this includes internet visitors as well as logged in users. Storage requirements will vary depending on how many pages and attachments you wish to store inside Confluence.

Minimum hardware requirements

The values below refer to the minimum available hardware required to run Confluence only; for example, the minimum heap size to allocate to Confluence is 1 GB and 1 GB for Synchrony (which is required for collaborative editing). You'll need additional physical hardware, of at least the minimum amount required by your Operating System and any other applications that run on the server.

On small instances, server load is primarily driven by peak visitors, so minimum system requirements are difficult to judge. We provide these figures as a guide to the absolute minimum required to run Confluence, and your configuration will likely require better hardware.

Here is our minimum hardware recommendation:

- **CPU:** Quad core 2GHz+ CPU
- **RAM:** 6GB
- **Minimum database space:** 10GB

Note: Please be aware that while some of our customers run Confluence on SPARC-based hardware, we only officially support Confluence running on x86 hardware and 64-bit derivatives of x86 hardware. Confluence typically will not perform well in a tightly constrained, shared environment - examples include an AWS micro.t1 instance. Please be careful to ensure that your choice of hosting platform is capable of supplying sustained processing and memory capacity for the server, particularly the processing-intense startup process.

Example hardware specifications

These are example hardware specifications for non-clustered Confluence instances. It is not recorded whether the amount of RAM refers to either the total server memory or memory allocated to the JVM, while blank settings indicate that the information was not provided.

Accounts	Spaces	Pages	CPUs	CPU (GHz)	RAM (MB)	Notes
150	30	1,000	1	2.6	1,024	
350	100	15,000	2	2.8	1,536	
5,000	500		4	3	2,048	
10,000	350	16,000	2	3.8	2,048	
10,000	60	3,500	2	3.6	4,096	
21,000	950		2	3.6	4,096	
85,000	100	12,500	4	2.6	4,096	3 machines total: application server, database server, Apache HTTPD + LDAP tunnel server.

Server load and scalability

When planning server hardware requirements for your Confluence deployment, you will need to estimate the server scalability based on peak visitors, the editor to viewer ratio and total content.

- The editor to viewer ratio is how many visitors are performing updates versus those only viewing content
- Total content is best estimated by a count of total spaces

Confluence scales best with a steady flow of visitors rather than defined peak visitor times, few editors and few spaces. Users should also take into account:

- Total pages is not a major consideration for performance. For example, instances hosting 80K of pages can consume under 512MB of memory
- Always use an external database, and check out the performance tuning guides.

Maximum reported usages

These values are largest customer instances reported to Atlassian or used for performance testing. Clustering, database tuning and other performance tuning is recommended for instances exceeding these values.

Most Spaces	1700
Most Internal Users	15K
Most LDAP Users	100K
Most Pages	80K

Hard disk requirements

All page content is stored in the database, while attachments are stored in the file system. The more attachments you have, the more disk space you will require.

Private and public comparison

Private instances manage their users either internally or through a user repository such as LDAP, while online instances have public signup enabled and must handle the additional load of anonymous internet visitors. Please keep in mind that these are examples only, not recommendations:

Use Case	Spaces	User Accounts	Editors	Editor To Viewer Ratio	Pages	Page Revisions	Attachments

Online Documentation	140	11,500	1,000	9%	8,800	65,000	7,300
Private Intranet	130	180	140	78%	8,000	84,000	3,800
Company-Wide Collaboration	100	85,000	1,000+	1%+	12,500	120,000	15,000

Professional assistance

For large instances, it may be worthwhile contacting an Atlassian Solution Partner for expertise on hardware sizing, testing and performance tuning.

Example site

Here is a breakdown of the disk usage and memory requirements a large documentation site as at April 2013:

Database size	2827 MB
Home directory size	116 GB
Average memory in use	1.9 GB

Size of selected database tables

Data	Relevant Table	Rows	Size
Attachment metadata	`attachments`	193903	60 MB
Content and user properties	`os_propertyentry (?)`	639737	255 MB
Content bodies (incl. all versions of blogs, pages and comments)	`bodycontent`	517520	1354 MB
Content metadata (incl. title, author)	`content`	623155	459 MB

Labels	`label (5982, 1264 kB), content_label (134151, 46 MB)`	140133	47.2 MB
Users	`users`	38766	6200 kB

Note: not all database tables or indexes are shown, and average row size may vary between instances.

Size of selected home directory components

Data	Files	Size
Attachments (incl. all versions)	207659	105 GB
Did-you-mean search index	10	14 MB
Office Connector cache	3506	456 MB
Plugin files	1851	669 MB
Search index	448	3.9 GB
Temporary files	14232	5 GB
Thumbnails	86516	1.7 GB
Usage index (now disabled)	239	2.6 GB

Note: not all files are shown, and average file size may vary between instances.

- Crontab -

```
# For details see man 4 crontabs

# Example of job definition:
# .---------------- minute (0 - 59)
# |  .------------- hour (0 - 23)
# |  |  .---------- day of month (1 - 31)
# |  |  |  .------- month (1 - 12) OR jan,feb,mar,apr ...
# |  |  |  |  .---- day of week (0 - 6) (Sunday=0 or 7) OR sun,mon,tue,wed,thu,fri,sat
# |  |  |  |  |
# *  *  *  *  * user-name command to be executed
```

- Date -

Display or change the date.

Syntax

date [*option*]... [+*Format*]

date [*option*] [MMDDhhmm[[CC]YY][.ss]]

'date' with no arguments prints the current time and date, in the format of the %c directive (described below).

If given an argument that starts with a +, date prints the current time and date (or the time and date specified by the --date option, see below) in the format defined by that argument, which is the same as in the strftime function.

Except for directives, which start with %, characters in the format string are printed unchanged. The directives are described below.

Options:

-d, --date=*String* Display time described by *String*, instead of 'now'
this can be in almost any common format.
It can contain month names, timezones, 'am' and 'pm',
'yesterday', 'ago', 'next', etc.

-f, --file=*DateFile* like --date once for each line of *DateFile*
If *DateFile* is '-', use standard input. This is
useful when you have many dates to process,
because the system overhead of starting up the
'date' executable many times can be considerable.

-I, --iso-8601[=*Timespec*] Output an ISO 8601 compliant date/time string., '%Y-%m-%d'.
Timespec='date' (or missing) for date only,
'hours', 'minutes', or 'seconds' for date and
time to the indicated precision.

If showing any time terms, then include the time zone
using the format '%z'. If '--utc' is also specified,
use '%Z' in place of '%z'.

-r, --reference=*File* Display the last modification time of *File*

-R, --rfc-822 Output RFC-822 compliant date string
Example: Mon, 19 Nov 2012 12:44:56 -0600

-s, --set=*String* Set time described by *String* (see -d above)

-u, --utc, --universal Print or set Coordinated Universal Time

--help Display this help and exit
--version output version information and exit

Format controls the output as follows. The only valid option for the second form (MMDDhhmm[[CC]YY][.ss])will specify Coordinated Universal Time.

Interpreted sequences are:

Date:
D Date in mm/dd/yy format (06/24/13)
x Date in standard format for locale (09/24/13 for English-US)

Year:
C Century (20 for 2015)
Y Year in 4-digit format (2015)
y Year in 2-digit format (14)
G Same as 'Y'
g Same as 'y'

Month:
b Month name - abbreviated (Jan)
B Month name - full (January)
h Same as 'b'
m Month number (09)

Week:
W Week of the year (00-52)

V Week of the year (01-53)
If the week containing January 1 has four or
more days in the new year, then it is considered week 1;
otherwise, it is week 53 of the previous year, and the next week
is week 1. Similar to ISO 8601 (but not 100% compliant.)

U Same as 'W'

Day:
a Day of the week - abbreviated name (Mon)
A Day of the week - full name (Monday)
u Day of the week - number (Monday = 1)
d Day of the month - 2 digits (05)
e Day of the month - digit preceded by a space (5)
j Day of the year - (1-366)

w Same as 'u'

Time:
p AM or PM
r Time in 12-hour format (09:15:36 AM)
R Time in 24-hour format - no seconds (17:45)
T Time in 24-hour format (17:45:52)
X Same as 'T'
Z Time offset from UTC (-07) This generally consists of Time Zone+DST

Hour:
H Hour in 24-hour format (17)
I Hour in 12 hour format (05)
k Same as 'H'
l Same as 'I' (Upper-case I = Lower-case L)

Minutes & Seconds:
M Minutes (35)
S Seconds (05)
s Seconds elapsed since January 1, 1970 00:00:00 GMT (Unix time)

Here are the same format codes in alphabetical order:

%% a literal %
%a locale's abbreviated weekday name (Sun..Sat)
%A locale's full weekday name, variable length (Sunday..Saturday)
%b locale's abbreviated month name (Jan..Dec)
%B locale's full month name, variable length (January..December)
%c locale's date and time (Sat Nov 04 12:02:33 EST 1989)
%d day of month (01..31)
%D date (mm/dd/yy)
%e day of month, blank padded (1..31)
%h same as %b, locale's abbreviated month name (Jan..Dec)
%H hour :24 hour(00..23)
%I hour :12 hour(01..12)
%j day of year (001..366)
%k hour :24 hour(00..23)
%l hour :12 hour(01..12)
%m month (01..12)
%M minute (00..59)
%n a newline
%p locale's AM or PM
%r Time, 12-hour (hh:mm:ss [AP]M)
%s Seconds since 1970-01-01 00:00:00, (a GNU extension)
Note that this value is defined by the localtime system
call. It isn't changed by the '--date' option.
%S second (00..60)
%t a horizontal tab
%T Time, 24-hour (hh:mm:ss)
%U Week number of year with Sunday as first day of week (00..53)
%V Week number of year with Monday as first day of week (01..53)
If the week containing January 1 has four or
more days in the new year, then it is considered week 1;
otherwise, it is week 53 of the previous year, and the next week
is week 1. Similar to ISO 8601 (but not 100% compliant.)

```
 %w   day of week (0..6);  0 represents Sunday
 %W   week number of year with Monday as first day of week (00..53)
 %x   locale's date representation (mm/dd/yy)
 %X   locale's time representation (%H:%M:%S)
 %y   last two digits of year (00..99)
 %Y   year (1970...)
 %z   RFC-822 style numeric timezone (-0500) (a nonstandard extension)
     This value reflects the current time zone.
     Is not changed by the --date option.
 %Z   Time offset from UTC (-07) This generally consists of Time Zone+DST
     Is not changed by the --date option.
```

By default, date pads numeric fields with zeroes. GNU date recognizes the following modifiers between % and a numeric directive.

- (hyphen) do not pad the field; useful if the output is intended for human consumption.
_ (underscore) pad the field with spaces; useful if you need a fixed number of characters in the output, but zeroes are too distracting.

The - and _ are GNU extensions. Here is an example illustrating the differences:

```
date +%d/%m -d "Feb 1"
=> 01/02
date +%-d/%-m -d "Feb 1"
=> 1/2
date +%_d/%_m -d "Feb 1"
=>  1/ 2
```

Setting the time

If given an argument that does not start with +, date sets the system clock to the time and date specified by that argument (as described below). You must have appropriate privileges to set the system clock. The --date and --set options can not be used with such an argument. The --universal option can be used with such an argument to indicate that the specified time and date are relative to Coordinated Universal Time rather than to the local time zone.

The argument must consist entirely of digits, which have the following meaning:

```
MM     month
DD     day within month
HH     hour
MM     minute
CC     first two digits of year (optional)
YY     last two digits of year (optional)
SS     second (optional)
```

The '--set' option also sets the system clock; see the examples below.

Examples

Print the date of the day before yesterday:

$ date --date='2 days ago'

Rename a file with the current date and time
$ STAMPME=$HOME/demo_file_$(date +%Y%m%d-%H%M).txt
$ mv $HOME/demo_file $STAMPME

Print the date of the day three months and one day hence:
$ date --date='3 months 1 day'

Print the day of year of Christmas in the current year:
$ date --date='25 Dec' +%j

Print the current full month name and the day of the month:
$ date '+%B %d'
Note that the '%d' expands to a zero-padded two-digit field, for example:
$ date -d 1may '+%B %d' will print 'May 01'.

Print a date without the leading zero for one-digit days of the month, you can use the (GNU extension) '-' modifier to suppress the padding altogether.
$ date -d=1may '+%B %-d'

Print the current date and time in the format required by many non-GNU versions of 'date' when setting the system clock:
$ date +%m%d%H%M%Y.%S

Set the system date and time
$ date --set="2012-6-29 11:59 AM"

Set the system clock forward by two minutes:
$ date --set='+2 minutes'

Print the date in the format specified by RFC-822 (day month year hh:mm:ss zzz), use
$ date --rfc

To convert a date string to the number of seconds since the epoch 1970-01-01 00:00:00 GMT (Unix time), use the '--date' option with the '%s' format. That can be useful in sorting and/or graphing /or comparing data by date. The following command outputs the number of the seconds since the epoch for the time one second later than the epoch, but in time zone five hours later (Cambridge,Massachusetts), thus a total of five hours and one second after the epoch:

$ date --date='2000-01-01 00:00:01 UTC +5 hours' +%s
946706400

Suppose you had **not** specified time zone information in the example above. Then, date would have used your computer's idea of the time zone (and DST) when interpreting the string. Here's what you would get if you were in Greenwich, England:

```
# local time zone used
$ date --date='2000-01-01 00:00:01' +%s
946684800
```

Seconds since the 1970 epoch can be useful when sorting or graphing dated data . But to convert a number of seconds back to a more readable date, use a command like:

```
$ date -d '1970-01-01 946684800 sec' +"%Y-%m-%d %T %z"
2000-01-01 00:00:00 +0000
```

- Arrêt de la base -

```
$ORACLE_HOME/bin/sqlplus /nolog << EOF
connect / as sysdba
shutdown immediate;
EOF
```