



Optimizing SQL Server on Amazon EC2 for Oracle JD Edwards EnterpriseOne

AWS Prescriptive Guidance



AWS Prescriptive Guidance: Optimizing SQL Server on Amazon EC2 for Oracle JD Edwards EnterpriseOne

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Introduction	1
Overview of JD Edwards EnterpriseOne behavior on SQL Server	3
Optimize layout and use appropriate resources	4
Placement of resources	4
EC2 instances and layout	4
Storage layout	5
SQL Server editions	6
Optimize SQL Server configuration	8
Adjust memory settings	8
Configure maximum and minimum memory	8
Lock pages in memory	10
Adjust CPU settings	10
Adjust MAXDOP	10
Adjust the cost threshold for parallelism	11
Enable instant file initialization	12
Configure data compression	12
Check disk space utilization before compression	13
Run the enumeration script	14
Run the compression script	15
Check disk space utilization after compression	16
Add and balance data files	17
Complete file size calculations	18
Create new files	19
Temporarily empty the MDF file	20
Resize the MDF file	20
Clean up	20
Validate results	21
Configure RCSI	22
Move tempdb to instance storage	23
Enable delayed durability	23
Next steps	26
Resources	28
Document history	29
Glossary	30

#	30
A	31
B	34
C	36
D	39
E	43
F	45
G	46
H	47
I	48
L	51
M	52
O	56
P	58
Q	61
R	61
S	64
T	68
U	69
V	70
W	70
Z	71

Optimizing SQL Server on Amazon EC2 for Oracle JD Edwards EnterpriseOne

Jeremy Shearer, Amazon Web Services (AWS)

December 2022 ([document history](#))

JD Edwards EnterpriseOne can be used with multiple database platforms, including Oracle Database, SQL Server, and IBM Db2. Many users find that SQL Server is a good database choice because of its balance of cost and features combined with their existing skills for managing a SQL Server database.

Each database platform supports multiple deployment options for EnterpriseOne on AWS, including Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Relational Database Service (Amazon RDS), as the following table shows.

EnterpriseOne platform	Deployment options on AWS		
	Amazon EC2	Amazon RDS	Other
Oracle Database	Yes	Yes	IBM Power Systems (i/AIX) and AWS Hybrid Architecture
SQL Server	Yes	Yes	
IBM Db2	Yes	No	IBM Power Systems (i/AIX) and AWS Hybrid Architecture

This guide focuses on deploying an EnterpriseOne database with SQL Server on Amazon EC2. For a detailed discussion of other SQL Server deployment options, see [Choosing between Amazon EC2 and Amazon RDS](#).

When you use Oracle JD Edwards EnterpriseOne with a SQL Server database on Amazon EC2, you can take advantage of specific optimization techniques to achieve a highly performant and cost-optimized system. This guide focuses on performance optimization of a SQL Server instance and

doesn't cover high availability, disaster recovery, backups, or other complementary configurations covered in other documents, including [Migrating Microsoft SQL Server databases to the AWS Cloud](#).

This guide builds upon the guide [Best practices for deploying SQL Server on Amazon EC2](#) and is intended for architects and DBAs who have a good understanding of SQL Server and JD Edwards EnterpriseOne.

Overview of JD Edwards EnterpriseOne behavior on SQL Server

EnterpriseOne business logic is primarily handled within applications. Only basic data manipulation language (DML) statements are passed to the database from the application. In standard processing, the record set is opened on the database but is managed by the application. The application then typically performs multiple DML operations for each record in the record set. This approach generates a substantial volume of *chatty* DML operations against the database. The latency of each DML operation is one of the key drivers of performance. Because of this architecture, CPU usage of the database that supports EnterpriseOne tends to be minimal, whereas network and disk I/O characteristics are the primary drivers of process performance. EnterpriseOne database tuning focuses heavily on the minimization of DML latency.

To mitigate the latency impact of disk read I/O, a large buffer cache is often used. This can be combined with SQL Server data compression to make the buffer cache substantially more effective. Although using data compression affects CPU, the overhead is minimal when you use this approach with EnterpriseOne. When the buffer cache is adequately sized, disk read I/O latency isn't typically an area of concern.

The SQL Server buffer cache doesn't address the latency of write I/O. When an EnterpriseOne process generates a large number of chatty write operations, performance may be constrained by the latency of each write operation that commits to the transaction log. To minimize this latency, you can use `io2` and/or `io2 Block Express` volumes for the LDF file. If `io2` or `io2 Block Express` alone is insufficient to deliver the required performance or is otherwise cost-prohibitive, you can use a delayed durability configuration to improve performance.

Because many EnterpriseOne processes create record sets that might overlap with other open record sets, you should enable read committed snapshot isolation (RCSI) on each EnterpriseOne database to minimize blocking. When this feature is enabled, it can create a substantial I/O requirement for `tempdb`. `tempdb` is by nature ephemeral and doesn't require the durability of standard block storage. In most cases, local instance non-volatile memory express (NVMe) storage is the best choice for `tempdb`.

The following sections of this guide explore these and other best practices for optimizing SQL Server for JD Edwards EnterpriseOne.

Optimize layout and use appropriate resources

Optimizing the resource layout and selecting appropriate resources affect the cost and performance of the system. When you use a SQL Server database with EnterpriseOne, consider the optimization patterns that are discussed in the following sections.

Topics

- [Placement of resources](#)
- [EC2 instances and layout](#)
- [Storage layout](#)
- [SQL Server editions](#)

Placement of resources

Because EnterpriseOne completes most business logic in the application tier, it tends to be very chatty across the network between the database and application tiers. As a result, processes that run on the application tier and access the database tier are often sensitive to network latency. To minimize network latency, we recommend that you place the EnterpriseOne database servers in the same placement group, within the same Availability Zone and Region, as the EnterpriseOne application servers.

If you are architecting a high availability configuration, you can use multiple techniques to ensure that the most sensitive processes run close to the database server. These techniques include using EnterpriseOne Object Configuration Manager (OCM) to map specific batch jobs (also known as UBEs) to specific servers, and using Virtual Batch Queues (VBQ) with remote nodes disabled.

For information about how to use a placement group on AWS, see [Placement groups](#) in the Amazon EC2 documentation.

EC2 instances and layout

SQL Server databases that support EnterpriseOne typically require:

- x86/x64 CPUs
- High-performance local instance storage for tempdb
- Large amount of memory for buffer cache

- High storage throughput and IOPS
- High network throughput
- Low vCPU count

Note

This section provides specific EC2 instance type and Amazon Elastic Block Store (Amazon EBS) storage recommendations, based on the information available at the time of this writing. As AWS adds support for new EC2 instances, Amazon EBS storage types, and Amazon FSx storage types, better options might become available. For the latest information, see the [Resources](#) section of this guide.

The Amazon EC2 [X2iedn](#) instance type is the preferred instance type for SQL Server databases that support EnterpriseOne. X2iedn provides high Amazon EBS throughput, high network throughput, and a large quantity of memory and quantity of instance storage per vCPU provisioned. It also supports [Provisioned IOPS SSD \(io2\) Block Express](#).

Some EnterpriseOne processes might require low-latency write I/O to support chatty commits. The volume type with the lowest latency write I/O is io2 Block Express, which is available only on a subset of x86/x64 instances that contain instance storage, including X2idn and X2iedn instances. When you use other x86/x64 instances that have instance storage, the lowest latency write I/O volume type will be io2.

Storage layout

When you use SQL Server database files with EnterpriseOne, they exhibit characteristics that support various disk types depending on their function.

- tempdb files should be placed on NVMe instance storage. When RCSI is enabled, a substantial workload is created in the tempdb database to store record set snapshots. These snapshots are ephemeral and do not require the durability of traditional elastic block storage. When you use NVMe instance storage, the database will receive very low latency I/O, high IOPS, and high throughput at a low price point.
- MDF and NDF data files should be placed on one or more [General Purpose SSD \(gp3\) volumes](#). These files tend to be read IOPS heavy but aren't very latency sensitive when they're used with

a large buffer cache. You can use multiple MDF and NDF file for each database to stripe your database across multiple disks to achieve the desired performance level.

- LDF files should be placed on a single gp3, or [Provisioned IOPS SSD](#) io2 or io2 Block Express volume based on requirements. Many JD Edwards processes perform operations that create chatty write I/O, which is latency sensitive. For many users, gp3 latency is sufficient to meet requirements. However, if you have a runtime-sensitive process, io2 or io2 Block Express might be required to meet your performance requirements for the workload. You might also consider enabling delayed durability in the SQL Server database to mitigate the performance impact of chatty write I/O. When delayed durability is enabled, you can use gp3 storage without being concerned about write I/O latency.
- Backup files should be placed on high-throughput, low-cost storage such as [Throughput Optimized HDD](#) (st1) or in an [Amazon Simple Storage Service \(Amazon S3\)](#) bucket. Additionally, because EnterpriseOne data tends to be repetitive and sparse, we recommend that you use SQL Server backup compression for backups you make through the database.
- Buffer pool extensions (BPEs) can provide value when you use an instance with substantial NVMe instance storage. However, when you use X2iedn instances, the benefit of BPE is substantially mitigated by the large amount of available memory, and it's better to use the available NVMe storage for tempdb.

SQL Server editions

Most users are able to leverage SQL Server Standard edition to meet the business requirements of their production systems and SQL Server Developer edition for their non-production environments. SQL Server Enterprise edition tends to be infrequently used for EnterpriseOne because of its high costs and because Microsoft moves features from Enterprise edition to Standard edition with each release. Many of the features that EnterpriseOne typically uses have been moved to SQL Server Standard edition, including the following:

- Maximum memory was increased to 128 GB in SQL Server 2012.
- Basic Always On availability groups for single databases were made available in SQL Server 2016.
- Database compression was made available in SQL Server 2016 SP1.
- BPEs became available in SQL Server 2017.
- Transparent data encryption became available in SQL Server 2019.

However, some features are available only in Enterprise edition. These include:

- Online index operations
- Using more than 128 GB of RAM per database instance
- Using more than 24 cores
- Resource Governor to manage workload and system resource consumption
- Read-ahead operations

Most EnterpriseOne users can take advantage of other solutions to meet their business requirements without using these Enterprise edition features.

Optimize SQL Server configuration

The default configuration of SQL Server is not optimized for JD Edwards EnterpriseOne. You must apply the appropriate configurations to ensure the optimal performance for EnterpriseOne running on a SQL Server database. The following sections describe these settings in detail.

Topics

- [Adjust memory settings](#)
- [Adjust CPU settings](#)
- [Enable instant file initialization](#)
- [Configure data compression](#)
- [Add and balance data files](#)
- [Configure RCSI](#)
- [Move tempdb to instance storage](#)
- [Enable delayed durability](#)

Adjust memory settings

We recommend that you configure the default memory values for a SQL Server database that is running JD Edwards workloads. These include:

- Configuring maximum and minimum memory settings
- Locking pages in memory

Configure maximum and minimum memory

Setting the maximum memory of the SQL Server database ensures that the operating systems and other processes have enough memory to perform their actions without paging to disk. Setting maximum and minimum memory can prevent multiple SQL Server instances that are installed on the same EC2 instance from starving one another for memory.

You can use the following script to automatically configure maximum and minimum settings with conservative values. This script reserves 1 GB for the operating system, and 25 percent of

the memory under 16 GB and 12.5 percent of the remaining memory as overhead. SQL Server minimum memory is set to half of maximum memory. The script assumes that you have a single SQL Server database installed on the EC2 instance.

```
DECLARE @OSMemoryTotalKB bigint;
DECLARE @OSMemoryUnder16GB bigint;
DECLARE @OSMemoryOver16GB bigint;
DECLARE @OSOverhead bigint;
DECLARE @MemoryOverheadLower bigint;
DECLARE @MemoryOverheadUpper bigint;
DECLARE @MemoryOverheadTotal bigint;
DECLARE @SQLMaxMemory int;
DECLARE @SQLMinMemory int;

-- Find how much memory is available on the OS
SELECT @OSMemoryTotalKB = total_physical_memory_kb from sys.dm_os_sys_memory;
SET @OSMemoryUnder16GB = IIF(@OSMemoryTotalKB>16777216, 16777216, @OSMemoryTotalKB);
SET @OSMemoryOver16GB = IIF(@OSMemoryTotalKB>16777216, @OSMemoryTotalKB-16777216, 0);

-- Calculate overhead for the OS
SET @OSOverhead= 1048576; -- static 1GB reservation

-- Calculate overhead for managing memory
SET @MemoryOverheadLower = @OSMemoryUnder16GB/4; --reserve 25% of memory under 16GB for
overhead
SET @MemoryOverheadUpper = @OSMemoryOver16GB/8; -- reserve 12.5% of memory over 16GB
for overhead
SET @MemoryOverheadTotal = @OSOverhead + @MemoryOverheadLower + @MemoryOverheadUpper;

-- Calculate remaining memory available for SQL
SET @SQLMaxMemory = (@OSMemoryTotalKB-@MemoryOverheadTotal)/1024;
SET @SQLMinMemory = @SQLMaxMemory/2; -- set minimum to half of maximum

Print N'Total Server memory (KB): ' + CAST(@OSMemoryTotalKB as NVARCHAR);
Print N'Memory Overhead for OS Overhead (KB): ' + CAST(@OSOverhead as NVARCHAR);
Print N'Memory Overhead for management of lower 16GB (KB): ' +
CAST(@MemoryOverheadLower as NVARCHAR);
Print N'Memory Overhead for management of over 16GB (KB): ' + CAST(@MemoryOverheadUpper
as NVARCHAR);
Print N'Memory Overhead Total: ' + CAST(@MemoryOverheadTotal as NVARCHAR);
Print N'SQL Minimum Memory (MB): ' + CAST(@SQLMinMemory as NVARCHAR);
Print N'SQL Maximum Memory (MB): ' + CAST(@SQLMaxMemory as NVARCHAR);
```

```
EXEC sp_configure 'show advanced options', 1;
RECONFIGURE;
EXEC sp_configure 'min server memory', @SQLMinMemory
RECONFIGURE;
EXEC sp_configure 'max server memory', @SQLMaxMemory;
RECONFIGURE;
```

Lock pages in memory

To ensure stability of the memory used for an EnterpriseOne SQL Server database, we recommend that you lock pages in memory. Follow the steps in the [Best practices for deploying SQL Server on Amazon EC2](#) guide to complete this configuration.

Adjust CPU settings

The default CPU settings on a SQL Server database allow processes to consume all available resources to complete their tasks. This configuration can starve EnterpriseOne processes on the CPU resources they require, causing performance issues and timeouts. To mitigate this issue, you can adjust the maximum degree of parallelism and cost threshold settings.

Adjust MAXDOP

By default, the maximum degree of parallelism (MAXDOP) is set to unlimited (0). Setting MAXDOP to a value of 1 disables parallelism and forces queries to run single-threaded. A value other than 0 or 1 sets the maximum number of parallel threads (vCPUs) that a single query can use.

To set the appropriate value for MAXDOP, consider the following:

- If you're running SQL Server Enterprise edition, you can use Resource Governor to control CPU allocation. However, because SQL Server Standard edition is typically more cost-effective, many EnterpriseOne installations cannot use Resource Governor.
- Most EnterpriseOne processes are short DML operations and do not use parallelism. However, many third-party applications benefit from parallelism and might experience performance degradation when parallelism is reduced or disabled.
- You can set a smaller MAXDOP value to limit the ability of any single process to saturate the system.

We recommend that you set the MAXDOP value, at most, to half the number of the vCPUs available in the instance. The minimum MAXDOP value would be 1, which disables parallelism entirely. The following query disables parallelism by setting MAXDOP to 1, but you can adjust the script to set it to any other MAXDOP value.

Note

The scripts in this guide use JDE_Prist920 as the EnterpriseOne database name. To use the scripts, update the database name to reflect your database.

```
USE JDE_Prist920;
GO
EXEC sp_configure 'show advanced options', 1;
GO
RECONFIGURE WITH OVERRIDE;
GO
EXEC sp_configure 'max degree of parallelism', 1;
GO
RECONFIGURE WITH OVERRIDE;
GO
```

Adjust the cost threshold for parallelism

If you enable parallelism by setting MAXDOP to a value greater than 1, set the cost threshold for parallelism to 50 or higher to limit the number of EnterpriseOne queries that are considered for parallelism. You can use the following script to set the value.

```
EXEC sp_configure 'show advanced options', 1;
GO
RECONFIGURE
GO
EXEC sp_configure 'cost threshold for parallelism', 50;
GO
RECONFIGURE
GO
```

Enable instant file initialization

When a database files grows, it fills the new disk space with zeros (0x0) by default. This creates significant system I/O and can degrade system performance. Instant file initialization prevents zeroing operations on the allocated disk space. To enable instant file initialization, follow the steps in the [Best practices for deploying SQL Server on Amazon EC2](#) guide.

Configure data compression

You can compress the tables and indexes in EnterpriseOne business data and control tables by using page or row compression. Most EnterpriseOne workloads on AWS exhibit the best performance with page compression, but extremely large workloads (multiples of uncompressed terabytes) might perform better with row compression. A detailed discussion of page versus row compression is beyond the scope of this guide. This section focuses primarily on page compression.

When you enable compression for normal EnterpriseOne workloads, there is a minimal increase in CPU usage but significant benefits to overall system performance, which can be measured in the following areas:

- Smaller database sizes and storage requirements, because the data is stored on disk in a compressed format.
- Higher buffer cache hit ratio, because the buffer cache can hold much more data when it's compressed.
- Lower required Amazon EBS IOPS and throughput, because each I/O operation returns much more data, and fewer operations are required, because the buffer cache is more effective.
- Faster backups, because data remains compressed throughout the backup process.

You can enable compression individually by table or by index alone. You can also choose the type of compression, either page or row, by table and index. It might be advantageous to not compress tables that are updated regularly, such as the F0002 (Next Number) and F0902 (Account Balances) tables. In many circumstances, enabling compression across all tables and indexes provides the easiest solution, because it provides most of the benefits without requiring an object-by-object analysis. The steps in this guide will compress all tables and indexes with page compression.

In some circumstances, compression might cause performance degradation, especially when third-party systems directly access the JD Edwards databases and perform table and index scan

operations. This degradation is typically driven by poorly performing queries. In these cases, review the slow queries and use common optimization techniques to improve their performance. For example, consider rewriting the queries to use existing indexes or build new indexes.

Enabling compression is a multi-step process. Many of these steps require exclusive access to the database objects, which means that you would have to take EnterpriseOne and other systems offline. Follow these high-level steps to enable page compression on all tables and indexes in the DTA and CTL schemas:

1. [Check disk space utilization before compression.](#)
2. [Run the enumeration script.](#)
3. [Run the compression script.](#)
4. [Check disk space utilization after compression.](#)

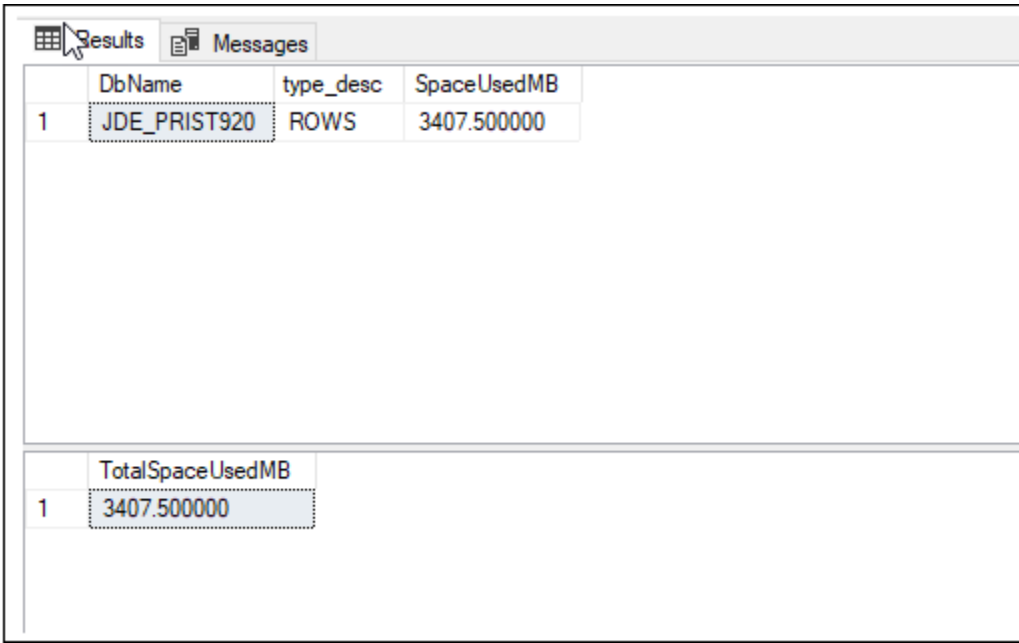
Check disk space utilization before compression

To check the current disk space utilization of the database, run the following scripts.

```
USE JDE_PRIST920
SELECT DB_NAME() AS DbName,
       type_desc,
       CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS SpaceUsedMB
FROM sys.database_files
WHERE type IN (0,1)
AND type_desc = 'ROWS';

SELECT SUM(CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0) AS TotalSpaceUsedMB
FROM sys.database_files
WHERE type IN (0,1)
AND type_desc = 'ROWS'
```

The output should be similar to the following:



	DbName	type_desc	SpaceUsedMB
1	JDE_PRIST920	ROWS	3407.500000

	TotalSpaceUsedMB
1	3407.500000

In this example, the table rows occupy 3,407 MB of disk space.

Run the enumeration script

Because of the large volume of tables and indexes in the EnterpriseOne database, you can use a script to enumerate the objects to be compressed. The output of the enumeration script is the compression script that is used in the next section. Before you run the following script, update the schema owner names to reflect the owners of the tables and indexes that you want to compress.

```
declare @tblname as varchar(100)
declare @idxname as varchar(100)
declare @schemaname as varchar(100)
declare @sqlstatement as varchar(512)
declare tblcurs CURSOR for
    select t.name as tblname, s.name as schemaname
    from sys.tables t
    inner join sys.schemas s on t.schema_id = s.schema_id
    inner join sys.indexes i on i.object_id = t.object_id
    inner join sys.partitions p on i.object_id = p.object_id AND i.index_id =
p.index_id
    where s.name in ('PS920DTA', 'PS920CTL')
    and i.type_desc='CLUSTERED'
    and p.data_compression_desc <> 'PAGE'
open tblcurs
FETCH next from tblcurs into @tblname, @schemaname
```

```
while @@FETCH_STATUS = 0
begin

    FETCH next from tblcurs into @tblname, @schemaname
    set @sqlstatement = 'alter table ' + @schemaname + '.' + @tblname + '
rebuild with (DATA_COMPRESSION = PAGE)'
    print @sqlstatement

end
close tblcurs
deallocate tblcurs
declare idxcurs CURSOR for
    select i.name as idxname, t.name as tblname, s.name as schemaname
    from sys.tables t
    inner join sys.schemas s on t.schema_id = s.schema_id
    inner join sys.indexes i on i.object_id = t.object_id
    inner JOIN sys.partitions p ON i.object_id = p.object_id AND i.index_id =
p.index_id
    where s.name in ('PS920DTA', 'PS920CTL')
    and p.data_compression_desc <> 'PAGE'
    and i.type_desc='NONCLUSTERED'
    and i.name is not null
open idxcurs
FETCH next from idxcurs into @idxname, @tblname, @schemaname
while @@FETCH_STATUS = 0
begin

    FETCH next from idxcurs into @idxname, @tblname, @schemaname
    set @sqlstatement = 'alter index ' + @idxname + ' on ' + @schemaname +
'.' + @tblname + ' rebuild with (DATA_COMPRESSION = PAGE)'
    print @sqlstatement

end
close idxcurs
deallocate idxcurs
```

Run the compression script

Review the output of the enumeration script that you ran in the last section. You can break this compression script up into smaller scripts and run them individually and in parallel.

⚠ Important

Make sure that the EnterpriseOne system is offline when you run this script against your EnterpriseOne database.

Here's an example of the compression script.

```
alter table PS920DTA.F07620 rebuild with (DATA_COMPRESSION = PAGE)
alter table PS920DTA.F760404A rebuild with (DATA_COMPRESSION = PAGE)
alter table PS920DTA.F31B93Z1 rebuild with (DATA_COMPRESSION = PAGE)
alter table PS920DTA.F31B65 rebuild with (DATA_COMPRESSION = PAGE)
alter table PS920DTA.F47156 rebuild with (DATA_COMPRESSION = PAGE)
alter table PS920DTA.F74F210 rebuild with (DATA_COMPRESSION = PAGE)
...
alter index F4611_16 on PS920DTA.F4611 rebuild with (DATA_COMPRESSION = PAGE)
alter index F4611_17 on PS920DTA.F4611 rebuild with (DATA_COMPRESSION = PAGE)
alter index F7000110_PK on PS920DTA.F7000110 rebuild with (DATA_COMPRESSION = PAGE)
alter index F7000110_3 on PS920DTA.F7000110 rebuild with (DATA_COMPRESSION = PAGE)
alter index F7000110_4 on PS920DTA.F7000110 rebuild with (DATA_COMPRESSION = PAGE)
alter index F76A801T_PK on PS920DTA.F76A801T rebuild with (DATA_COMPRESSION = PAGE)
...
```

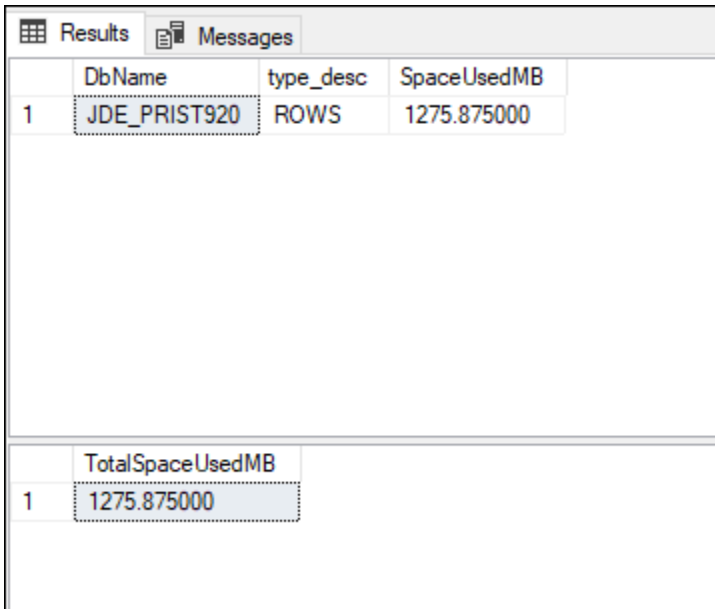
Check disk space utilization after compression

To check the current disk space utilization of the database after compression, run the following script.

```
USE JDE_PRIST920
SELECT DB_NAME() AS dbName,
       type_desc,
       CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS SpaceUsedMB
FROM sys.database_files
WHERE type IN (0,1)
AND type_desc = 'ROWS';

SELECT SUM(CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0) AS TotalSpaceUsedMB
FROM sys.database_files
WHERE type IN (0,1)
AND type_desc = 'ROWS'
```

The output should be similar to the following.



The screenshot shows a SQL Server Enterprise Manager interface with two tabs: 'Results' and 'Messages'. The 'Results' tab is active and displays a table with the following data:

	DbName	type_desc	SpaceUsedMB
1	JDE_PRIST920	ROWS	1275.875000

Below this table, there is a summary row with the following data:

	TotalSpaceUsedMB
1	1275.875000

In this example, you can see that space used dropped from 3,407 MB to 1,275 MB, which represents a 62 percent savings from compression. The savings for your database will vary based on how data is distributed among the tables in the database.

Add and balance data files

The SQL Server databases provided with EnterpriseOne can often benefit from additional files. Additional files enable optimal balancing across storage and processor cores. Balancing the files is a multi-step process. Many of these steps require exclusive access to the database objects, so you would have to take EnterpriseOne and other systems accessing the database offline.

1. [Complete file sizing calculations.](#)
2. [Create new files.](#)
3. [Temporarily empty the MDF file.](#)
4. [Resize the MDF file.](#)
5. [Clean up.](#)
6. [Validate the results.](#)

Complete file size calculations

To find the appropriate size of the database files, start by examining the size of the current ROW data by using the following query.

```
USE JDE_PRIST920
SELECT SUM(CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0) AS SpaceUsedMB
FROM sys.database_files
WHERE type IN (0,1)
AND type_desc = 'ROWS'
```

Then complete the following calculations and fill out the *Your value* column:

Line	Name	Example	Your value	Description
1	Current size of the row	1 TB		The results from the previous query.
2	Planned growth	20%		Expected growth over the next number of months, including a safety margin.
3	Required size	1.2 TB		Line 1 multiplied by line 2.
4	Number of files	8		The number of files targeted.
5	Size per file	150 GB		Line 3 divided by line 4.
6	Autogrowth percentage	10%		The size for automatic growth. For minimum fragmentation, 10% is a good target.

Line	Name	Example	Your value	Description
7	Autogrowth size	15 GB		Line 5 multiplied by line 6.

Create new files

Use the following script as a template to add files to the database. Modify the following parameters:

- Change JDE-PRIST920 to the name of the database you want to add files to.
- For NAME, specify the logical name of each file you want to add.
- For FILENAME, specify the physical name of each file you want to add.
- For FILEGROWTH, use the value you calculated in row 7 of the previous table.
- For SIZE, specify the value from row 5 of the previous table.

```
USE master;
GO

ALTER DATABASE JDE_PRIST920
MODIFY FILE (NAME = JDE_PRIST920_Data, FILEGROWTH = 15GB);
GO

ALTER DATABASE JDE_PRIST920
ADD FILE
(NAME = JDE_PRIST920_Data2, FILENAME = 'M:\DATA\PRIST920_Data2.ndf', SIZE=150GB,
FILEGROWTH = 15GB),
(NAME = JDE_PRIST920_Data3, FILENAME = 'M:\DATA\PRIST920_Data3.ndf', SIZE=150GB,
FILEGROWTH = 15GB),
(NAME = JDE_PRIST920_Data4, FILENAME = 'M:\DATA\PRIST920_Data4.ndf', SIZE=150GB,
FILEGROWTH = 15GB),
(NAME = JDE_PRIST920_Data5, FILENAME = 'M:\DATA\PRIST920_Data5.ndf', SIZE=150GB,
FILEGROWTH = 15GB),
(NAME = JDE_PRIST920_Data6, FILENAME = 'M:\DATA\PRIST920_Data6.ndf', SIZE=150GB,
FILEGROWTH = 15GB),
(NAME = JDE_PRIST920_Data7, FILENAME = 'M:\DATA\PRIST920_Data7.ndf', SIZE=150GB,
FILEGROWTH = 15GB),
```

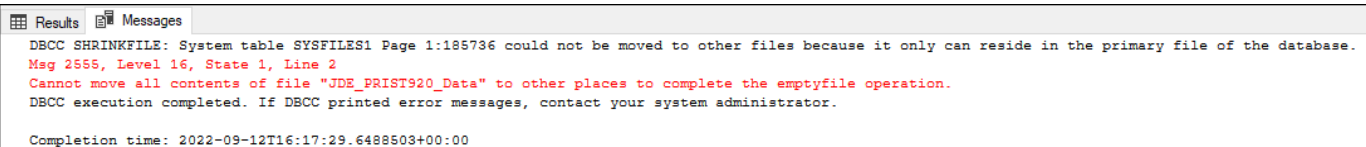
```
(NAME = JDE_PRIST920_Data8, FILENAME = 'M:\DATA\PRIST920_Data8.ndf', SIZE=150GB,  
FILEGROWTH = 15GB),  
(NAME = JDE_PRIST920_TEMP, FILENAME = 'M:\DATA\PRIST920_TEMP.ndf', SIZE=150GB,  
FILEGROWTH = 15GB)  
GO
```

Temporarily empty the MDF file

When the files have been created, migrate the data from the MDF to the NDF files by running the following command for each file. Adjust the file names to reflect the file names in your database.

```
USE JDE_PRIST920  
DBCC SHRINKFILE (JDE_PRIST920_Data, EMPTYFILE)
```

The EMPTYFILE command generates an error because some content can't be moved to an NDF file. You can ignore this error message.



```
Results Messages  
DBCC SHRINKFILE: System table SYSFILES1 Page 1:185736 could not be moved to other files because it only can reside in the primary file of the database.  
Msg 2555, Level 16, State 1, Line 2  
Cannot move all contents of file "JDE_PRIST920_Data" to other places to complete the emptyfile operation.  
DBCC execution completed. If DBCC printed error messages, contact your system administrator.  
  
Completion time: 2022-09-12T16:17:29.6488503+00:00
```

Resize the MDF file

To reduce the size of the MDF file to the target size, run the following command. Adjust the file size to reflect the value from line 5 of the calculation in the table.

```
JDE_PRIST920  
DBCC SHRINKFILE (JDE_PRIST920_Data, 150000);
```

Occasionally, the SHRINKFILE command will fail because of the placement of content that couldn't be moved to an NDF file. In this case, you might need to run the DBCC DBREINDEX command, rerun the process to empty the file, and try the SHRINKFILE operation again.

Clean up

When the target files have been created and the MDF file is correctly sized, use the following command to migrate the data from the TEMP file back to the MDF file. Adjust the file names to reflect the file names in your database.


```
DBCC SHRINKFILE (JDE_PRIST920_TEMP, EMPTYFILE)
```

When the file is empty, you can remove it by using the following command:

```
ALTER DATABASE JDE_PRIST920;  
REMOVE FILE JDE_PRIST920_TEMP;
```

Validate results

To check the current disk space utilization of the database after balancing, run the following scripts.

```
USE JDE_PRIST920  
SELECT DB_NAME() AS DbName,  
       type_desc,  
       CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS SpaceUsedMB  
FROM sys.database_files  
WHERE type IN (0,1)  
AND type_desc = 'ROWS';  
  
SELECT SUM(CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0) AS TotalSpaceUsedMB  
FROM sys.database_files  
WHERE type IN (0,1)  
AND type_desc = 'ROWS'
```

The output should be similar to the following. The files will seldom be perfectly balanced, because some content can exist only in the MDF file.

	DbName	type_desc	SpaceUsedMB
1	JDE_PRIST920	ROWS	330.687500
2	JDE_PRIST920	ROWS	144.812500
3	JDE_PRIST920	ROWS	146.125000
4	JDE_PRIST920	ROWS	144.125000
5	JDE_PRIST920	ROWS	142.500000
6	JDE_PRIST920	ROWS	142.625000
7	JDE_PRIST920	ROWS	143.312500
8	JDE_PRIST920	ROWS	150.812500

	TotalSpaceUsedMB
1	1345.000000

Configure RCSI

We recommend that you enable read-committed snapshot isolation (RCSI) for each EnterpriseOne database. Otherwise, you might encounter locking scenarios that affect system performance. In some cases, EnterpriseOne might revert to *NOLOCK, which can result in dirty read operations, if RCSI isn't in effect.

To enable RCSI, run the following command for each database, and revise the database name accordingly.

```
USE master
ALTER DATABASE JDE_Prist920 SET READ_COMMITTED_SNAPSHOT ON
GO
```

After you enable RCSI, adjust the EnterpriseOne configuration as described in My Oracle Support (MOS) note 2565588 (requires registration). The adjustments disable the SQL query timeout and disable retries with *NOLOCK. These settings are required when RCSI isn't enabled.

Move tempdb to instance storage

When RCSI is enabled, a significant IOPS and throughput load can be created in tempdb to maintain versions of records during transactions. Because of this load, you should move tempdb to NVMe instance storage. For information about how to move tempdb to the instance store, follow the steps in the [Best practices for deploying SQL Server on Amazon EC2](#) guide.

Enable delayed durability

Certain processes such as EnterpriseOne's Material Requirements Planning (MRP) encounter a bottleneck caused by the disk latency of committing transactions to the transaction log. Moving the transaction log (LDF) to `io2` or `io2 Block Express` storage often improves the performance sufficiently to meet business requirements. If this is insufficient, you can configure delayed durability in the database.

Important

You should enable delayed durability only if performance isn't acceptable and you fully understand how your transactions will behave across databases during system failure scenarios.

When you enable delayed durability, transaction commits are cached by using a write-back (instead of a write-through) operation. In the event of a system failure, the consistency of the database is still guaranteed. However, any transactions that haven't been committed to disk are lost. Also, additional functionality related to replication, including the functionality used by AWS Database Migration Service (AWS DMS), becomes unavailable when delayed durability is in effect.

During MRP testing in a specific configuration, we observed the following:

- Moving the LDF file to `io2 Block Express` dropped the runtime by 52 percent compared to a baseline with the LDF file on `gp3`.
- Enabling delayed durability dropped the runtime by 79 percent compared to a baseline with the LDF file on `gp3`.

To enable delayed durability, run the following command on the database.

```
USE master
ALTER DATABASE JDE_Prist920 SET DELAYED_DURABILITY = FORCE
```

Delayed durability typically flushes the log several times per second, but there can be an increased lag if there is a disk I/O bottleneck. To ensure a low recovery point objective (RPO), you can place the `sys.sp_flush_log` command on a scheduler to run at a high frequency. This procedure forces a flush of the log to disk.

The following script creates a job on the SQL job scheduler to run every minute. Adjust the job name and database name in the script to reflect your requirements.

```
USE msdb;
GO

DECLARE @myjob nvarchar(128);
DECLARE @mydb nvarchar(128);
DECLARE @mycommand nvarchar(max);
DECLARE @myschedule nvarchar(128);
DECLARE @jobId binary(16);
DECLARE @scheduleId binary(16);

SET @myjob = 'JDE_Prist920 Flush Log Cache';
SET @mydb = 'JDE_Prist920';
SET @mycommand = 'sys.sp_flush_log';
SET @myschedule = 'EveryMinute';

SELECT @scheduleId = schedule_id FROM msdb.dbo.sysschedules WHERE (name = @myschedule)
IF (@scheduleId IS NULL)
BEGIN
    EXEC sp_add_schedule
        @schedule_name = @myschedule,
        @freq_type = 4,
        @freq_interval = 1,
        @freq_subday_type = 0x2,
        @freq_subday_interval= 60
END

SELECT @jobId = job_id FROM msdb.dbo.sysjobs WHERE (name = @myjob)
IF (@jobId IS NULL)
BEGIN
    EXEC sp_add_job @job_name = @myjob
```

```
EXEC sp_add_jobstep
    @job_name = @myjob,
    @step_name = N'process step',
    @subsystem = N'TSQL',
    @command = @mycommand,
    @database_name = @mydb

EXEC sp_attach_schedule
    @job_name = @myjob,
    @schedule_name = @myschedule;

EXEC dbo.sp_add_jobserver
    @job_name = @myjob

END
```

Next steps

This guide focused on optimizing a SQL Server configuration for EnterpriseOne. Topics such as database disaster recovery are beyond the scope of this document but should be addressed as part of the configuration; see the [Resources](#) section for additional reading.

There are also many AWS services and features that you can use to optimize your EnterpriseOne system, including the following.

Service	Use
AWS Application Migration Service	You can use Application Migration Service to migrate EnterpriseOne from any source infrastructure that runs supported operating systems and databases, including Microsoft Windows, Red Hat Enterprise Linux (RHEL), Oracle Linux, SQL Server, and Oracle Database.
AWS Elastic Disaster Recovery	Elastic Disaster Recovery minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery.
AWS Database Migration Service (AWS DMS)	AWS DMS supports the migration of data between more than 20 different database platforms, including those used with EnterpriseOne. Common EnterpriseOne use cases include building data marts by using EnterpriseOne data.
Application Load Balancer	The Application Load Balancer allows you to spread your workload among multiple HTTP or HTTPS-based EnterpriseOne application services.
Amazon WorkSpaces	You can use Amazon WorkSpaces products to access high-performance workstations, on demand, for your JD Edwards development and administrative client applications.
AWS IAM Identity Center	You can use AWS Identity and Access Management (IAM) to provide authentication for EnterpriseOne. This service is

Service	Use
	available with EnterpriseOne Tools 9.2.5.4 or later, which supports JSON web tokens (JWTs).
Amazon Simple Email Service (Amazon SES)	Amazon SES provides a reliable and compliant way to manage your EnterpriseOne emails. This service is available for all EnterpriseOne releases by using a third-party utility for SMTP authentication. EnterpriseOne Tools 9.2.7 and later versions provide support for authenticated SMTP with EnterpriseOne.

Resources

- [Migrating Microsoft SQL Server databases to the AWS Cloud](#) (AWS Prescriptive Guidance)
- [Best practices for deploying SQL Server on Amazon EC2](#) (AWS Prescriptive Guidance)
- [Amazon EC2 instances](#) (Amazon EC2 documentation)
- [Amazon Elastic Block Store \(Amazon EBS\) documentation](#)
- [Placement groups](#) (Amazon EC2 documentation)
- [Hosting IBM i and AIX Systems with Low-Latency Connectivity to AWS with Connectria](#) (AWS blog post)

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an [RSS feed](#).

Change	Description	Date
Initial publication	—	December 6, 2022

AWS Prescriptive Guidance glossary

The following are commonly used terms in strategies, guides, and patterns provided by AWS Prescriptive Guidance. To suggest entries, please use the **Provide feedback** link at the end of the glossary.

Numbers

7 Rs

Seven common migration strategies for moving applications to the cloud. These strategies build upon the 5 Rs that Gartner identified in 2011 and consist of the following:

- **Refactor/re-architect** – Move an application and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This typically involves porting the operating system and database. Example: Migrate your on-premises Oracle database to the Amazon Aurora PostgreSQL-Compatible Edition.
- **Replatform (lift and reshape)** – Move an application to the cloud, and introduce some level of optimization to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Amazon Relational Database Service (Amazon RDS) for Oracle in the AWS Cloud.
- **Repurchase (drop and shop)** – Switch to a different product, typically by moving from a traditional license to a SaaS model. Example: Migrate your customer relationship management (CRM) system to Salesforce.com.
- **Rehost (lift and shift)** – Move an application to the cloud without making any changes to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Oracle on an EC2 instance in the AWS Cloud.
- **Relocate (hypervisor-level lift and shift)** – Move infrastructure to the cloud without purchasing new hardware, rewriting applications, or modifying your existing operations. You migrate servers from an on-premises platform to a cloud service for the same platform. Example: Migrate a Microsoft Hyper-V application to AWS.
- **Retain (revisit)** – Keep applications in your source environment. These might include applications that require major refactoring, and you want to postpone that work until a later time, and legacy applications that you want to retain, because there's no business justification for migrating them.

- Retire – Decommission or remove applications that are no longer needed in your source environment.

A

ABAC

See [attribute-based access control](#).

abstracted services

See [managed services](#).

ACID

See [atomicity, consistency, isolation, durability](#).

active-active migration

A database migration method in which the source and target databases are kept in sync (by using a bidirectional replication tool or dual write operations), and both databases handle transactions from connecting applications during migration. This method supports migration in small, controlled batches instead of requiring a one-time cutover. It's more flexible but requires more work than [active-passive migration](#).

active-passive migration

A database migration method in which the source and target databases are kept in sync, but only the source database handles transactions from connecting applications while data is replicated to the target database. The target database doesn't accept any transactions during migration.

aggregate function

A SQL function that operates on a group of rows and calculates a single return value for the group. Examples of aggregate functions include SUM and MAX.

AI

See [artificial intelligence](#).

AIOps

See [artificial intelligence operations](#).

anonymization

The process of permanently deleting personal information in a dataset. Anonymization can help protect personal privacy. Anonymized data is no longer considered to be personal data.

anti-pattern

A frequently used solution for a recurring issue where the solution is counter-productive, ineffective, or less effective than an alternative.

application control

A security approach that allows the use of only approved applications in order to help protect a system from malware.

application portfolio

A collection of detailed information about each application used by an organization, including the cost to build and maintain the application, and its business value. This information is key to [the portfolio discovery and analysis process](#) and helps identify and prioritize the applications to be migrated, modernized, and optimized.

artificial intelligence (AI)

The field of computer science that is dedicated to using computing technologies to perform cognitive functions that are typically associated with humans, such as learning, solving problems, and recognizing patterns. For more information, see [What is Artificial Intelligence?](#)

artificial intelligence operations (AIOps)

The process of using machine learning techniques to solve operational problems, reduce operational incidents and human intervention, and increase service quality. For more information about how AIOps is used in the AWS migration strategy, see the [operations integration guide](#).

asymmetric encryption

An encryption algorithm that uses a pair of keys, a public key for encryption and a private key for decryption. You can share the public key because it isn't used for decryption, but access to the private key should be highly restricted.

atomicity, consistency, isolation, durability (ACID)

A set of software properties that guarantee the data validity and operational reliability of a database, even in the case of errors, power failures, or other problems.

attribute-based access control (ABAC)

The practice of creating fine-grained permissions based on user attributes, such as department, job role, and team name. For more information, see [ABAC for AWS](#) in the AWS Identity and Access Management (IAM) documentation.

authoritative data source

A location where you store the primary version of data, which is considered to be the most reliable source of information. You can copy data from the authoritative data source to other locations for the purposes of processing or modifying the data, such as anonymizing, redacting, or pseudonymizing it.

Availability Zone

A distinct location within an AWS Region that is insulated from failures in other Availability Zones and provides inexpensive, low-latency network connectivity to other Availability Zones in the same Region.

AWS Cloud Adoption Framework (AWS CAF)

A framework of guidelines and best practices from AWS to help organizations develop an efficient and effective plan to move successfully to the cloud. AWS CAF organizes guidance into six focus areas called perspectives: business, people, governance, platform, security, and operations. The business, people, and governance perspectives focus on business skills and processes; the platform, security, and operations perspectives focus on technical skills and processes. For example, the people perspective targets stakeholders who handle human resources (HR), staffing functions, and people management. For this perspective, AWS CAF provides guidance for people development, training, and communications to help ready the organization for successful cloud adoption. For more information, see the [AWS CAF website](#) and the [AWS CAF whitepaper](#).

AWS Workload Qualification Framework (AWS WQF)

A tool that evaluates database migration workloads, recommends migration strategies, and provides work estimates. AWS WQF is included with AWS Schema Conversion Tool (AWS SCT). It analyzes database schemas and code objects, application code, dependencies, and performance characteristics, and provides assessment reports.

B

bad bot

A [bot](#) that is intended to disrupt or cause harm to individuals or organizations.

BCP

See [business continuity planning](#).

behavior graph

A unified, interactive view of resource behavior and interactions over time. You can use a behavior graph with Amazon Detective to examine failed logon attempts, suspicious API calls, and similar actions. For more information, see [Data in a behavior graph](#) in the Detective documentation.

big-endian system

A system that stores the most significant byte first. See also [endianness](#).

binary classification

A process that predicts a binary outcome (one of two possible classes). For example, your ML model might need to predict problems such as "Is this email spam or not spam?" or "Is this product a book or a car?"

bloom filter

A probabilistic, memory-efficient data structure that is used to test whether an element is a member of a set.

blue/green deployment

A deployment strategy where you create two separate but identical environments. You run the current application version in one environment (blue) and the new application version in the other environment (green). This strategy helps you quickly roll back with minimal impact.

bot

A software application that runs automated tasks over the internet and simulates human activity or interaction. Some bots are useful or beneficial, such as web crawlers that index information on the internet. Some other bots, known as *bad bots*, are intended to disrupt or cause harm to individuals or organizations.

botnet

Networks of [bots](#) that are infected by [malware](#) and are under the control of a single party, known as a *bot herder* or *bot operator*. Botnets are the best-known mechanism to scale bots and their impact.

branch

A contained area of a code repository. The first branch created in a repository is the *main branch*. You can create a new branch from an existing branch, and you can then develop features or fix bugs in the new branch. A branch you create to build a feature is commonly referred to as a *feature branch*. When the feature is ready for release, you merge the feature branch back into the main branch. For more information, see [About branches](#) (GitHub documentation).

break-glass access

In exceptional circumstances and through an approved process, a quick means for a user to gain access to an AWS account that they don't typically have permissions to access. For more information, see the [Implement break-glass procedures](#) indicator in the AWS Well-Architected guidance.

brownfield strategy

The existing infrastructure in your environment. When adopting a brownfield strategy for a system architecture, you design the architecture around the constraints of the current systems and infrastructure. If you are expanding the existing infrastructure, you might blend brownfield and [greenfield](#) strategies.

buffer cache

The memory area where the most frequently accessed data is stored.

business capability

What a business does to generate value (for example, sales, customer service, or marketing). Microservices architectures and development decisions can be driven by business capabilities. For more information, see the [Organized around business capabilities](#) section of the [Running containerized microservices on AWS](#) whitepaper.

business continuity planning (BCP)

A plan that addresses the potential impact of a disruptive event, such as a large-scale migration, on operations and enables a business to resume operations quickly.

C

CAF

See [AWS Cloud Adoption Framework](#).

canary deployment

The slow and incremental release of a version to end users. When you are confident, you deploy the new version and replace the current version in its entirety.

CCoE

See [Cloud Center of Excellence](#).

CDC

See [change data capture](#).

change data capture (CDC)

The process of tracking changes to a data source, such as a database table, and recording metadata about the change. You can use CDC for various purposes, such as auditing or replicating changes in a target system to maintain synchronization.

chaos engineering

Intentionally introducing failures or disruptive events to test a system's resilience. You can use [AWS Fault Injection Service \(AWS FIS\)](#) to perform experiments that stress your AWS workloads and evaluate their response.

CI/CD

See [continuous integration and continuous delivery](#).

classification

A categorization process that helps generate predictions. ML models for classification problems predict a discrete value. Discrete values are always distinct from one another. For example, a model might need to evaluate whether or not there is a car in an image.

client-side encryption

Encryption of data locally, before the target AWS service receives it.

Cloud Center of Excellence (CCoE)

A multi-disciplinary team that drives cloud adoption efforts across an organization, including developing cloud best practices, mobilizing resources, establishing migration timelines, and leading the organization through large-scale transformations. For more information, see the [CCoE posts](#) on the AWS Cloud Enterprise Strategy Blog.

cloud computing

The cloud technology that is typically used for remote data storage and IoT device management. Cloud computing is commonly connected to [edge computing](#) technology.

cloud operating model

In an IT organization, the operating model that is used to build, mature, and optimize one or more cloud environments. For more information, see [Building your Cloud Operating Model](#).

cloud stages of adoption

The four phases that organizations typically go through when they migrate to the AWS Cloud:

- Project – Running a few cloud-related projects for proof of concept and learning purposes
- Foundation – Making foundational investments to scale your cloud adoption (e.g., creating a landing zone, defining a CCoE, establishing an operations model)
- Migration – Migrating individual applications
- Re-invention – Optimizing products and services, and innovating in the cloud

These stages were defined by Stephen Orban in the blog post [The Journey Toward Cloud-First & the Stages of Adoption](#) on the AWS Cloud Enterprise Strategy blog. For information about how they relate to the AWS migration strategy, see the [migration readiness guide](#).

CMDB

See [configuration management database](#).

code repository

A location where source code and other assets, such as documentation, samples, and scripts, are stored and updated through version control processes. Common cloud repositories include GitHub or AWS CodeCommit. Each version of the code is called a *branch*. In a microservice structure, each repository is devoted to a single piece of functionality. A single CI/CD pipeline can use multiple repositories.

cold cache

A buffer cache that is empty, not well populated, or contains stale or irrelevant data. This affects performance because the database instance must read from the main memory or disk, which is slower than reading from the buffer cache.

cold data

Data that is rarely accessed and is typically historical. When querying this kind of data, slow queries are typically acceptable. Moving this data to lower-performing and less expensive storage tiers or classes can reduce costs.

computer vision (CV)

A field of [AI](#) that uses machine learning to analyze and extract information from visual formats such as digital images and videos. For example, AWS Panorama offers devices that add CV to on-premises camera networks, and Amazon SageMaker provides image processing algorithms for CV.

configuration drift

For a workload, a configuration change from the expected state. It might cause the workload to become noncompliant, and it's typically gradual and unintentional.

configuration management database (CMDB)

A repository that stores and manages information about a database and its IT environment, including both hardware and software components and their configurations. You typically use data from a CMDB in the portfolio discovery and analysis stage of migration.

conformance pack

A collection of AWS Config rules and remediation actions that you can assemble to customize your compliance and security checks. You can deploy a conformance pack as a single entity in an AWS account and Region, or across an organization, by using a YAML template. For more information, see [Conformance packs](#) in the AWS Config documentation.

continuous integration and continuous delivery (CI/CD)

The process of automating the source, build, test, staging, and production stages of the software release process. CI/CD is commonly described as a pipeline. CI/CD can help you automate processes, improve productivity, improve code quality, and deliver faster. For more information, see [Benefits of continuous delivery](#). CD can also stand for *continuous deployment*. For more information, see [Continuous Delivery vs. Continuous Deployment](#).

CV

See [computer vision](#).

D

data at rest

Data that is stationary in your network, such as data that is in storage.

data classification

A process for identifying and categorizing the data in your network based on its criticality and sensitivity. It is a critical component of any cybersecurity risk management strategy because it helps you determine the appropriate protection and retention controls for the data. Data classification is a component of the security pillar in the AWS Well-Architected Framework. For more information, see [Data classification](#).

data drift

A meaningful variation between the production data and the data that was used to train an ML model, or a meaningful change in the input data over time. Data drift can reduce the overall quality, accuracy, and fairness in ML model predictions.

data in transit

Data that is actively moving through your network, such as between network resources.

data mesh

An architectural framework that provides distributed, decentralized data ownership with centralized management and governance.

data minimization

The principle of collecting and processing only the data that is strictly necessary. Practicing data minimization in the AWS Cloud can reduce privacy risks, costs, and your analytics carbon footprint.

data perimeter

A set of preventive guardrails in your AWS environment that help make sure that only trusted identities are accessing trusted resources from expected networks. For more information, see [Building a data perimeter on AWS](#).

data preprocessing

To transform raw data into a format that is easily parsed by your ML model. Preprocessing data can mean removing certain columns or rows and addressing missing, inconsistent, or duplicate values.

data provenance

The process of tracking the origin and history of data throughout its lifecycle, such as how the data was generated, transmitted, and stored.

data subject

An individual whose data is being collected and processed.

data warehouse

A data management system that supports business intelligence, such as analytics. Data warehouses commonly contain large amounts of historical data, and they are typically used for queries and analysis.

database definition language (DDL)

Statements or commands for creating or modifying the structure of tables and objects in a database.

database manipulation language (DML)

Statements or commands for modifying (inserting, updating, and deleting) information in a database.

DDL

See [database definition language](#).

deep ensemble

To combine multiple deep learning models for prediction. You can use deep ensembles to obtain a more accurate prediction or for estimating uncertainty in predictions.

deep learning

An ML subfield that uses multiple layers of artificial neural networks to identify mapping between input data and target variables of interest.

defense-in-depth

An information security approach in which a series of security mechanisms and controls are thoughtfully layered throughout a computer network to protect the confidentiality, integrity, and availability of the network and the data within. When you adopt this strategy on AWS, you add multiple controls at different layers of the AWS Organizations structure to help secure resources. For example, a defense-in-depth approach might combine multi-factor authentication, network segmentation, and encryption.

delegated administrator

In AWS Organizations, a compatible service can register an AWS member account to administer the organization's accounts and manage permissions for that service. This account is called the *delegated administrator* for that service. For more information and a list of compatible services, see [Services that work with AWS Organizations](#) in the AWS Organizations documentation.

deployment

The process of making an application, new features, or code fixes available in the target environment. Deployment involves implementing changes in a code base and then building and running that code base in the application's environments.

development environment

See [environment](#).

detective control

A security control that is designed to detect, log, and alert after an event has occurred. These controls are a second line of defense, alerting you to security events that bypassed the preventative controls in place. For more information, see [Detective controls](#) in *Implementing security controls on AWS*.

development value stream mapping (DVSM)

A process used to identify and prioritize constraints that adversely affect speed and quality in a software development lifecycle. DVSM extends the value stream mapping process originally designed for lean manufacturing practices. It focuses on the steps and teams required to create and move value through the software development process.

digital twin

A virtual representation of a real-world system, such as a building, factory, industrial equipment, or production line. Digital twins support predictive maintenance, remote monitoring, and production optimization.

dimension table

In a [star schema](#), a smaller table that contains data attributes about quantitative data in a fact table. Dimension table attributes are typically text fields or discrete numbers that behave like text. These attributes are commonly used for query constraining, filtering, and result set labeling.

disaster

An event that prevents a workload or system from fulfilling its business objectives in its primary deployed location. These events can be natural disasters, technical failures, or the result of human actions, such as unintentional misconfiguration or a malware attack.

disaster recovery (DR)

The strategy and process you use to minimize downtime and data loss caused by a [disaster](#). For more information, see [Disaster Recovery of Workloads on AWS: Recovery in the Cloud](#) in the AWS Well-Architected Framework.

DML

See [database manipulation language](#).

domain-driven design

An approach to developing a complex software system by connecting its components to evolving domains, or core business goals, that each component serves. This concept was introduced by Eric Evans in his book, *Domain-Driven Design: Tackling Complexity in the Heart of Software* (Boston: Addison-Wesley Professional, 2003). For information about how you can use domain-driven design with the strangler fig pattern, see [Modernizing legacy Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway](#).

DR

See [disaster recovery](#).

drift detection

Tracking deviations from a baselined configuration. For example, you can use AWS CloudFormation to [detect drift in system resources](#), or you can use AWS Control Tower to [detect changes in your landing zone](#) that might affect compliance with governance requirements.

DVSM

See [development value stream mapping](#).

E

EDA

See [exploratory data analysis](#).

edge computing

The technology that increases the computing power for smart devices at the edges of an IoT network. When compared with [cloud computing](#), edge computing can reduce communication latency and improve response time.

encryption

A computing process that transforms plaintext data, which is human-readable, into ciphertext.

encryption key

A cryptographic string of randomized bits that is generated by an encryption algorithm. Keys can vary in length, and each key is designed to be unpredictable and unique.

endianness

The order in which bytes are stored in computer memory. Big-endian systems store the most significant byte first. Little-endian systems store the least significant byte first.

endpoint

See [service endpoint](#).

endpoint service

A service that you can host in a virtual private cloud (VPC) to share with other users. You can create an endpoint service with AWS PrivateLink and grant permissions to other AWS accounts

or to AWS Identity and Access Management (IAM) principals. These accounts or principals can connect to your endpoint service privately by creating interface VPC endpoints. For more information, see [Create an endpoint service](#) in the Amazon Virtual Private Cloud (Amazon VPC) documentation.

enterprise resource planning (ERP)

A system that automates and manages key business processes (such as accounting, [MES](#), and project management) for an enterprise.

envelope encryption

The process of encrypting an encryption key with another encryption key. For more information, see [Envelope encryption](#) in the AWS Key Management Service (AWS KMS) documentation.

environment

An instance of a running application. The following are common types of environments in cloud computing:

- development environment – An instance of a running application that is available only to the core team responsible for maintaining the application. Development environments are used to test changes before promoting them to upper environments. This type of environment is sometimes referred to as a *test environment*.
- lower environments – All development environments for an application, such as those used for initial builds and tests.
- production environment – An instance of a running application that end users can access. In a CI/CD pipeline, the production environment is the last deployment environment.
- upper environments – All environments that can be accessed by users other than the core development team. This can include a production environment, preproduction environments, and environments for user acceptance testing.

epic

In agile methodologies, functional categories that help organize and prioritize your work. Epics provide a high-level description of requirements and implementation tasks. For example, AWS CAF security epics include identity and access management, detective controls, infrastructure security, data protection, and incident response. For more information about epics in the AWS migration strategy, see the [program implementation guide](#).

ERP

See [enterprise resource planning](#).

exploratory data analysis (EDA)

The process of analyzing a dataset to understand its main characteristics. You collect or aggregate data and then perform initial investigations to find patterns, detect anomalies, and check assumptions. EDA is performed by calculating summary statistics and creating data visualizations.

F

fact table

The central table in a [star schema](#). It stores quantitative data about business operations. Typically, a fact table contains two types of columns: those that contain measures and those that contain a foreign key to a dimension table.

fail fast

A philosophy that uses frequent and incremental testing to reduce the development lifecycle. It is a critical part of an agile approach.

fault isolation boundary

In the AWS Cloud, a boundary such as an Availability Zone, AWS Region, control plane, or data plane that limits the effect of a failure and helps improve the resilience of workloads. For more information, see [AWS Fault Isolation Boundaries](#).

feature branch

See [branch](#).

features

The input data that you use to make a prediction. For example, in a manufacturing context, features could be images that are periodically captured from the manufacturing line.

feature importance

How significant a feature is for a model's predictions. This is usually expressed as a numerical score that can be calculated through various techniques, such as Shapley Additive Explanations

(SHAP) and integrated gradients. For more information, see [Machine learning model interpretability with :AWS](#).

feature transformation

To optimize data for the ML process, including enriching data with additional sources, scaling values, or extracting multiple sets of information from a single data field. This enables the ML model to benefit from the data. For example, if you break down the “2021-05-27 00:15:37” date into “2021”, “May”, “Thu”, and “15”, you can help the learning algorithm learn nuanced patterns associated with different data components.

FGAC

See [fine-grained access control](#).

fine-grained access control (FGAC)

The use of multiple conditions to allow or deny an access request.

flash-cut migration

A database migration method that uses continuous data replication through [change data capture](#) to migrate data in the shortest time possible, instead of using a phased approach. The objective is to keep downtime to a minimum.

G

geo blocking

See [geographic restrictions](#).

geographic restrictions (geo blocking)

In Amazon CloudFront, an option to prevent users in specific countries from accessing content distributions. You can use an allow list or block list to specify approved and banned countries. For more information, see [Restricting the geographic distribution of your content](#) in the CloudFront documentation.

Gitflow workflow

An approach in which lower and upper environments use different branches in a source code repository. The Gitflow workflow is considered legacy, and the [trunk-based workflow](#) is the modern, preferred approach.

greenfield strategy

The absence of existing infrastructure in a new environment. When adopting a greenfield strategy for a system architecture, you can select all new technologies without the restriction of compatibility with existing infrastructure, also known as [brownfield](#). If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

guardrail

A high-level rule that helps govern resources, policies, and compliance across organizational units (OUs). *Preventive guardrails* enforce policies to ensure alignment to compliance standards. They are implemented by using service control policies and IAM permissions boundaries. *Detective guardrails* detect policy violations and compliance issues, and generate alerts for remediation. They are implemented by using AWS Config, AWS Security Hub, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

H

HA

See [high availability](#).

heterogeneous database migration

Migrating your source database to a target database that uses a different database engine (for example, Oracle to Amazon Aurora). Heterogeneous migration is typically part of a re-architecting effort, and converting the schema can be a complex task. [AWS provides AWS SCT](#) that helps with schema conversions.

high availability (HA)

The ability of a workload to operate continuously, without intervention, in the event of challenges or disasters. HA systems are designed to automatically fail over, consistently deliver high-quality performance, and handle different loads and failures with minimal performance impact.

historian modernization

An approach used to modernize and upgrade operational technology (OT) systems to better serve the needs of the manufacturing industry. A *historian* is a type of database that is used to collect and store data from various sources in a factory.

homogeneous database migration

Migrating your source database to a target database that shares the same database engine (for example, Microsoft SQL Server to Amazon RDS for SQL Server). Homogeneous migration is typically part of a rehosting or replatforming effort. You can use native database utilities to migrate the schema.

hot data

Data that is frequently accessed, such as real-time data or recent translational data. This data typically requires a high-performance storage tier or class to provide fast query responses.

hotfix

An urgent fix for a critical issue in a production environment. Due to its urgency, a hotfix is usually made outside of the typical DevOps release workflow.

hypercare period

Immediately following cutover, the period of time when a migration team manages and monitors the migrated applications in the cloud in order to address any issues. Typically, this period is 1–4 days in length. At the end of the hypercare period, the migration team typically transfers responsibility for the applications to the cloud operations team.

I

laC

See [infrastructure as code](#).

identity-based policy

A policy attached to one or more IAM principals that defines their permissions within the AWS Cloud environment.

idle application

An application that has an average CPU and memory usage between 5 and 20 percent over a period of 90 days. In a migration project, it is common to retire these applications or retain them on premises.

IIoT

See [Industrial Internet of Things](#).

immutable infrastructure

A model that deploys new infrastructure for production workloads instead of updating, patching, or modifying the existing infrastructure. Immutable infrastructures are inherently more consistent, reliable, and predictable than [mutable infrastructure](#). For more information, see the [Deploy using immutable infrastructure](#) best practice in the AWS Well-Architected Framework.

inbound (ingress) VPC

In an AWS multi-account architecture, a VPC that accepts, inspects, and routes network connections from outside an application. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

incremental migration

A cutover strategy in which you migrate your application in small parts instead of performing a single, full cutover. For example, you might move only a few microservices or users to the new system initially. After you verify that everything is working properly, you can incrementally move additional microservices or users until you can decommission your legacy system. This strategy reduces the risks associated with large migrations.

Industry 4.0

A term that was introduced by [Klaus Schwab](#) in 2016 to refer to the modernization of manufacturing processes through advances in connectivity, real-time data, automation, analytics, and AI/ML.

infrastructure

All of the resources and assets contained within an application's environment.

infrastructure as code (IaC)

The process of provisioning and managing an application's infrastructure through a set of configuration files. IaC is designed to help you centralize infrastructure management, standardize resources, and scale quickly so that new environments are repeatable, reliable, and consistent.

industrial Internet of Things (IIoT)

The use of internet-connected sensors and devices in the industrial sectors, such as manufacturing, energy, automotive, healthcare, life sciences, and agriculture. For more information, see [Building an industrial Internet of Things \(IIoT\) digital transformation strategy](#).

inspection VPC

In an AWS multi-account architecture, a centralized VPC that manages inspections of network traffic between VPCs (in the same or different AWS Regions), the internet, and on-premises networks. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

Internet of Things (IoT)

The network of connected physical objects with embedded sensors or processors that communicate with other devices and systems through the internet or over a local communication network. For more information, see [What is IoT?](#)

interpretability

A characteristic of a machine learning model that describes the degree to which a human can understand how the model's predictions depend on its inputs. For more information, see [Machine learning model interpretability with AWS](#).

IoT

See [Internet of Things](#).

IT information library (ITIL)

A set of best practices for delivering IT services and aligning these services with business requirements. ITIL provides the foundation for ITSM.

IT service management (ITSM)

Activities associated with designing, implementing, managing, and supporting IT services for an organization. For information about integrating cloud operations with ITSM tools, see the [operations integration guide](#).

ITIL

See [IT information library](#).

ITSM

See [IT service management](#).

L

label-based access control (LBAC)

An implementation of mandatory access control (MAC) where the users and the data itself are each explicitly assigned a security label value. The intersection between the user security label and data security label determines which rows and columns can be seen by the user.

landing zone

A landing zone is a well-architected, multi-account AWS environment that is scalable and secure. This is a starting point from which your organizations can quickly launch and deploy workloads and applications with confidence in their security and infrastructure environment. For more information about landing zones, see [Setting up a secure and scalable multi-account AWS environment](#).

large migration

A migration of 300 or more servers.

LBAC

See [label-based access control](#).

least privilege

The security best practice of granting the minimum permissions required to perform a task. For more information, see [Apply least-privilege permissions](#) in the IAM documentation.

lift and shift

See [7 Rs](#).

little-endian system

A system that stores the least significant byte first. See also [endianness](#).

lower environments

See [environment](#).

M

machine learning (ML)

A type of artificial intelligence that uses algorithms and techniques for pattern recognition and learning. ML analyzes and learns from recorded data, such as Internet of Things (IoT) data, to generate a statistical model based on patterns. For more information, see [Machine Learning](#).

main branch

See [branch](#).

malware

Software that is designed to compromise computer security or privacy. Malware might disrupt computer systems, leak sensitive information, or gain unauthorized access. Examples of malware include viruses, worms, ransomware, Trojan horses, spyware, and keyloggers.

managed services

AWS services for which AWS operates the infrastructure layer, the operating system, and platforms, and you access the endpoints to store and retrieve data. Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB are examples of managed services. These are also known as *abstracted services*.

manufacturing execution system (MES)

A software system for tracking, monitoring, documenting, and controlling production processes that convert raw materials to finished products on the shop floor.

MAP

See [Migration Acceleration Program](#).

mechanism

A complete process in which you create a tool, drive adoption of the tool, and then inspect the results in order to make adjustments. A mechanism is a cycle that reinforces and improves itself as it operates. For more information, see [Building mechanisms](#) in the AWS Well-Architected Framework.

member account

All AWS accounts other than the management account that are part of an organization in AWS Organizations. An account can be a member of only one organization at a time.

MES

See [manufacturing execution system](#).

Message Queuing Telemetry Transport (MQTT)

A lightweight, machine-to-machine (M2M) communication protocol, based on the [publish/subscribe](#) pattern, for resource-constrained [IoT](#) devices.

microservice

A small, independent service that communicates over well-defined APIs and is typically owned by small, self-contained teams. For example, an insurance system might include microservices that map to business capabilities, such as sales or marketing, or subdomains, such as purchasing, claims, or analytics. The benefits of microservices include agility, flexible scaling, easy deployment, reusable code, and resilience. For more information, see [Integrating microservices by using AWS serverless services](#).

microservices architecture

An approach to building an application with independent components that run each application process as a microservice. These microservices communicate through a well-defined interface by using lightweight APIs. Each microservice in this architecture can be updated, deployed, and scaled to meet demand for specific functions of an application. For more information, see [Implementing microservices on AWS](#).

Migration Acceleration Program (MAP)

An AWS program that provides consulting support, training, and services to help organizations build a strong operational foundation for moving to the cloud, and to help offset the initial cost of migrations. MAP includes a migration methodology for executing legacy migrations in a methodical way and a set of tools to automate and accelerate common migration scenarios.

migration at scale

The process of moving the majority of the application portfolio to the cloud in waves, with more applications moved at a faster rate in each wave. This phase uses the best practices and lessons learned from the earlier phases to implement a *migration factory* of teams, tools, and processes to streamline the migration of workloads through automation and agile delivery. This is the third phase of the [AWS migration strategy](#).

migration factory

Cross-functional teams that streamline the migration of workloads through automated, agile approaches. Migration factory teams typically include operations, business analysts and owners, migration engineers, developers, and DevOps professionals working in sprints. Between 20 and 50 percent of an enterprise application portfolio consists of repeated patterns that can be optimized by a factory approach. For more information, see the [discussion of migration factories](#) and the [Cloud Migration Factory guide](#) in this content set.

migration metadata

The information about the application and server that is needed to complete the migration. Each migration pattern requires a different set of migration metadata. Examples of migration metadata include the target subnet, security group, and AWS account.

migration pattern

A repeatable migration task that details the migration strategy, the migration destination, and the migration application or service used. Example: Rehost migration to Amazon EC2 with AWS Application Migration Service.

Migration Portfolio Assessment (MPA)

An online tool that provides information for validating the business case for migrating to the AWS Cloud. MPA provides detailed portfolio assessment (server right-sizing, pricing, TCO comparisons, migration cost analysis) as well as migration planning (application data analysis and data collection, application grouping, migration prioritization, and wave planning). The [MPA tool](#) (requires login) is available free of charge to all AWS consultants and APN Partner consultants.

Migration Readiness Assessment (MRA)

The process of gaining insights about an organization's cloud readiness status, identifying strengths and weaknesses, and building an action plan to close identified gaps, using the AWS CAF. For more information, see the [migration readiness guide](#). MRA is the first phase of the [AWS migration strategy](#).

migration strategy

The approach used to migrate a workload to the AWS Cloud. For more information, see the [7 Rs](#) entry in this glossary and see [Mobilize your organization to accelerate large-scale migrations](#).

ML

See [machine learning](#).

modernization

Transforming an outdated (legacy or monolithic) application and its infrastructure into an agile, elastic, and highly available system in the cloud to reduce costs, gain efficiencies, and take advantage of innovations. For more information, see [Strategy for modernizing applications in the AWS Cloud](#).

modernization readiness assessment

An evaluation that helps determine the modernization readiness of an organization's applications; identifies benefits, risks, and dependencies; and determines how well the organization can support the future state of those applications. The outcome of the assessment is a blueprint of the target architecture, a roadmap that details development phases and milestones for the modernization process, and an action plan for addressing identified gaps. For more information, see [Evaluating modernization readiness for applications in the AWS Cloud](#).

monolithic applications (monoliths)

Applications that run as a single service with tightly coupled processes. Monolithic applications have several drawbacks. If one application feature experiences a spike in demand, the entire architecture must be scaled. Adding or improving a monolithic application's features also becomes more complex when the code base grows. To address these issues, you can use a microservices architecture. For more information, see [Decomposing monoliths into microservices](#).

MPA

See [Migration Portfolio Assessment](#).

MQTT

See [Message Queuing Telemetry Transport](#).

multiclass classification

A process that helps generate predictions for multiple classes (predicting one of more than two outcomes). For example, an ML model might ask "Is this product a book, car, or phone?" or "Which product category is most interesting to this customer?"

mutable infrastructure

A model that updates and modifies the existing infrastructure for production workloads. For improved consistency, reliability, and predictability, the AWS Well-Architected Framework recommends the use of [immutable infrastructure](#) as a best practice.

O

OAC

See [origin access control](#).

OAI

See [origin access identity](#).

OCM

See [organizational change management](#).

offline migration

A migration method in which the source workload is taken down during the migration process. This method involves extended downtime and is typically used for small, non-critical workloads.

OI

See [operations integration](#).

OLA

See [operational-level agreement](#).

online migration

A migration method in which the source workload is copied to the target system without being taken offline. Applications that are connected to the workload can continue to function during the migration. This method involves zero to minimal downtime and is typically used for critical production workloads.

OPC-UA

See [Open Process Communications - Unified Architecture](#).

Open Process Communications - Unified Architecture (OPC-UA)

A machine-to-machine (M2M) communication protocol for industrial automation. OPC-UA provides an interoperability standard with data encryption, authentication, and authorization schemes.

operational-level agreement (OLA)

An agreement that clarifies what functional IT groups promise to deliver to each other, to support a service-level agreement (SLA).

operational readiness review (ORR)

A checklist of questions and associated best practices that help you understand, evaluate, prevent, or reduce the scope of incidents and possible failures. For more information, see [Operational Readiness Reviews \(ORR\)](#) in the AWS Well-Architected Framework.

operational technology (OT)

Hardware and software systems that work with the physical environment to control industrial operations, equipment, and infrastructure. In manufacturing, the integration of OT and information technology (IT) systems is a key focus for [Industry 4.0](#) transformations.

operations integration (OI)

The process of modernizing operations in the cloud, which involves readiness planning, automation, and integration. For more information, see the [operations integration guide](#).

organization trail

A trail that's created by AWS CloudTrail that logs all events for all AWS accounts in an organization in AWS Organizations. This trail is created in each AWS account that's part of the organization and tracks the activity in each account. For more information, see [Creating a trail for an organization](#) in the CloudTrail documentation.

organizational change management (OCM)

A framework for managing major, disruptive business transformations from a people, culture, and leadership perspective. OCM helps organizations prepare for, and transition to, new systems and strategies by accelerating change adoption, addressing transitional issues, and driving cultural and organizational changes. In the AWS migration strategy, this framework is called *people acceleration*, because of the speed of change required in cloud adoption projects. For more information, see the [OCM guide](#).

origin access control (OAC)

In CloudFront, an enhanced option for restricting access to secure your Amazon Simple Storage Service (Amazon S3) content. OAC supports all S3 buckets in all AWS Regions, server-side encryption with AWS KMS (SSE-KMS), and dynamic PUT and DELETE requests to the S3 bucket.

origin access identity (OAI)

In CloudFront, an option for restricting access to secure your Amazon S3 content. When you use OAI, CloudFront creates a principal that Amazon S3 can authenticate with. Authenticated principals can access content in an S3 bucket only through a specific CloudFront distribution. See also [OAC](#), which provides more granular and enhanced access control.

ORR

See [operational readiness review](#).

OT

See [operational technology](#).

outbound (egress) VPC

In an AWS multi-account architecture, a VPC that handles network connections that are initiated from within an application. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

P

permissions boundary

An IAM management policy that is attached to IAM principals to set the maximum permissions that the user or role can have. For more information, see [Permissions boundaries](#) in the IAM documentation.

personally identifiable information (PII)

Information that, when viewed directly or paired with other related data, can be used to reasonably infer the identity of an individual. Examples of PII include names, addresses, and contact information.

PII

See [personally identifiable information](#).

playbook

A set of predefined steps that capture the work associated with migrations, such as delivering core operations functions in the cloud. A playbook can take the form of scripts, automated runbooks, or a summary of processes or steps required to operate your modernized environment.

PLC

See [programmable logic controller](#).

PLM

See [product lifecycle management](#).

policy

An object that can define permissions (see [identity-based policy](#)), specify access conditions (see [resource-based policy](#)), or define the maximum permissions for all accounts in an organization in AWS Organizations (see [service control policy](#)).

polyglot persistence

Independently choosing a microservice's data storage technology based on data access patterns and other requirements. If your microservices have the same data storage technology, they can encounter implementation challenges or experience poor performance. Microservices are more easily implemented and achieve better performance and scalability if they use the data store best adapted to their requirements. For more information, see [Enabling data persistence in microservices](#).

portfolio assessment

A process of discovering, analyzing, and prioritizing the application portfolio in order to plan the migration. For more information, see [Evaluating migration readiness](#).

predicate

A query condition that returns true or false, commonly located in a WHERE clause.

predicate pushdown

A database query optimization technique that filters the data in the query before transfer. This reduces the amount of data that must be retrieved and processed from the relational database, and it improves query performance.

preventative control

A security control that is designed to prevent an event from occurring. These controls are a first line of defense to help prevent unauthorized access or unwanted changes to your network. For more information, see [Preventative controls](#) in *Implementing security controls on AWS*.

principal

An entity in AWS that can perform actions and access resources. This entity is typically a root user for an AWS account, an IAM role, or a user. For more information, see *Principal* in [Roles terms and concepts](#) in the IAM documentation.

Privacy by Design

An approach in system engineering that takes privacy into account throughout the whole engineering process.

private hosted zones

A container that holds information about how you want Amazon Route 53 to respond to DNS queries for a domain and its subdomains within one or more VPCs. For more information, see [Working with private hosted zones](#) in the Route 53 documentation.

proactive control

A [security control](#) designed to prevent the deployment of noncompliant resources. These controls scan resources before they are provisioned. If the resource is not compliant with the control, then it isn't provisioned. For more information, see the [Controls reference guide](#) in the AWS Control Tower documentation and see [Proactive controls](#) in *Implementing security controls on AWS*.

product lifecycle management (PLM)

The management of data and processes for a product throughout its entire lifecycle, from design, development, and launch, through growth and maturity, to decline and removal.

production environment

See [environment](#).

programmable logic controller (PLC)

In manufacturing, a highly reliable, adaptable computer that monitors machines and automates manufacturing processes.

pseudonymization

The process of replacing personal identifiers in a dataset with placeholder values. Pseudonymization can help protect personal privacy. Pseudonymized data is still considered to be personal data.

publish/subscribe (pub/sub)

A pattern that enables asynchronous communications among microservices to improve scalability and responsiveness. For example, in a microservices-based [MES](#), a microservice can publish event messages to a channel that other microservices can subscribe to. The system can add new microservices without changing the publishing service.

Q

query plan

A series of steps, like instructions, that are used to access the data in a SQL relational database system.

query plan regression

When a database service optimizer chooses a less optimal plan than it did before a given change to the database environment. This can be caused by changes to statistics, constraints, environment settings, query parameter bindings, and updates to the database engine.

R

RACI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

ransomware

A malicious software that is designed to block access to a computer system or data until a payment is made.

RASCI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

RCAC

See [row and column access control](#).

read replica

A copy of a database that's used for read-only purposes. You can route queries to the read replica to reduce the load on your primary database.

re-architect

See [7 Rs](#).

recovery point objective (RPO)

The maximum acceptable amount of time since the last data recovery point. This determines what is considered an acceptable loss of data between the last recovery point and the interruption of service.

recovery time objective (RTO)

The maximum acceptable delay between the interruption of service and restoration of service.

refactor

See [7 Rs](#).

Region

A collection of AWS resources in a geographic area. Each AWS Region is isolated and independent of the others to provide fault tolerance, stability, and resilience. For more information, see [Specify which AWS Regions your account can use](#).

regression

An ML technique that predicts a numeric value. For example, to solve the problem of "What price will this house sell for?" an ML model could use a linear regression model to predict a house's sale price based on known facts about the house (for example, the square footage).

rehost

See [7 Rs](#).

release

In a deployment process, the act of promoting changes to a production environment.

relocate

See [7 Rs](#).

replatform

See [7 Rs](#).

repurchase

See [7 Rs](#).

resiliency

An application's ability to resist or recover from disruptions. [High availability](#) and [disaster recovery](#) are common considerations when planning for resiliency in the AWS Cloud. For more information, see [AWS Cloud Resilience](#).

resource-based policy

A policy attached to a resource, such as an Amazon S3 bucket, an endpoint, or an encryption key. This type of policy specifies which principals are allowed access, supported actions, and any other conditions that must be met.

responsible, accountable, consulted, informed (RACI) matrix

A matrix that defines the roles and responsibilities for all parties involved in migration activities and cloud operations. The matrix name is derived from the responsibility types defined in the matrix: responsible (R), accountable (A), consulted (C), and informed (I). The support (S) type is optional. If you include support, the matrix is called a *RASCI matrix*, and if you exclude it, it's called a *RACI matrix*.

responsive control

A security control that is designed to drive remediation of adverse events or deviations from your security baseline. For more information, see [Responsive controls](#) in *Implementing security controls on AWS*.

retain

See [7 Rs](#).

retire

See [7 Rs](#).

rotation

The process of periodically updating a [secret](#) to make it more difficult for an attacker to access the credentials.

row and column access control (RCAC)

The use of basic, flexible SQL expressions that have defined access rules. RCAC consists of row permissions and column masks.

RPO

See [recovery point objective](#).

RTO

See [recovery time objective](#).

runbook

A set of manual or automated procedures required to perform a specific task. These are typically built to streamline repetitive operations or procedures with high error rates.

S

SAML 2.0

An open standard that many identity providers (IdPs) use. This feature enables federated single sign-on (SSO), so users can log into the AWS Management Console or call the AWS API operations without you having to create user in IAM for everyone in your organization. For more information about SAML 2.0-based federation, see [About SAML 2.0-based federation](#) in the IAM documentation.

SCADA

See [supervisory control and data acquisition](#).

SCP

See [service control policy](#).

secret

In AWS Secrets Manager, confidential or restricted information, such as a password or user credentials, that you store in encrypted form. It consists of the secret value and its metadata. The secret value can be binary, a single string, or multiple strings. For more information, see [What's in a Secrets Manager secret?](#) in the Secrets Manager documentation.

security control

A technical or administrative guardrail that prevents, detects, or reduces the ability of a threat actor to exploit a security vulnerability. There are four primary types of security controls: [preventative](#), [detective](#), [responsive](#), and [proactive](#).

security hardening

The process of reducing the attack surface to make it more resistant to attacks. This can include actions such as removing resources that are no longer needed, implementing the security best practice of granting least privilege, or deactivating unnecessary features in configuration files.

security information and event management (SIEM) system

Tools and services that combine security information management (SIM) and security event management (SEM) systems. A SIEM system collects, monitors, and analyzes data from servers, networks, devices, and other sources to detect threats and security breaches, and to generate alerts.

security response automation

A predefined and programmed action that is designed to automatically respond to or remediate a security event. These automations serve as [detective](#) or [responsive](#) security controls that help you implement AWS security best practices. Examples of automated response actions include modifying a VPC security group, patching an Amazon EC2 instance, or rotating credentials.

server-side encryption

Encryption of data at its destination, by the AWS service that receives it.

service control policy (SCP)

A policy that provides centralized control over permissions for all accounts in an organization in AWS Organizations. SCPs define guardrails or set limits on actions that an administrator can delegate to users or roles. You can use SCPs as allow lists or deny lists, to specify which services or actions are permitted or prohibited. For more information, see [Service control policies](#) in the AWS Organizations documentation.

service endpoint

The URL of the entry point for an AWS service. You can use the endpoint to connect programmatically to the target service. For more information, see [AWS service endpoints](#) in *AWS General Reference*.

service-level agreement (SLA)

An agreement that clarifies what an IT team promises to deliver to their customers, such as service uptime and performance.

service-level indicator (SLI)

A measurement of a performance aspect of a service, such as its error rate, availability, or throughput.

service-level objective (SLO)

A target metric that represents the health of a service, as measured by a [service-level indicator](#).

shared responsibility model

A model describing the responsibility you share with AWS for cloud security and compliance. AWS is responsible for security *of* the cloud, whereas you are responsible for security *in* the cloud. For more information, see [Shared responsibility model](#).

SIEM

See [security information and event management system](#).

single point of failure (SPOF)

A failure in a single, critical component of an application that can disrupt the system.

SLA

See [service-level agreement](#).

SLI

See [service-level indicator](#).

SLO

See [service-level objective](#).

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your organization's capabilities and services, improves developer productivity, and supports rapid innovation. For more information, see [Phased approach to modernizing applications in the AWS Cloud](#).

SPOF

See [single point of failure](#).

star schema

A database organizational structure that uses one large fact table to store transactional or measured data and uses one or more smaller dimensional tables to store data attributes. This structure is designed for use in a [data warehouse](#) or for business intelligence purposes.

strangler fig pattern

An approach to modernizing monolithic systems by incrementally rewriting and replacing system functionality until the legacy system can be decommissioned. This pattern uses the analogy of a fig vine that grows into an established tree and eventually overcomes and replaces its host. The pattern was [introduced by Martin Fowler](#) as a way to manage risk when rewriting monolithic systems. For an example of how to apply this pattern, see [Modernizing legacy Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway](#).

subnet

A range of IP addresses in your VPC. A subnet must reside in a single Availability Zone.

supervisory control and data acquisition (SCADA)

In manufacturing, a system that uses hardware and software to monitor physical assets and production operations.

symmetric encryption

An encryption algorithm that uses the same key to encrypt and decrypt the data.

synthetic testing

Testing a system in a way that simulates user interactions to detect potential issues or to monitor performance. You can use [Amazon CloudWatch Synthetics](#) to create these tests.

T

tags

Key-value pairs that act as metadata for organizing your AWS resources. Tags can help you manage, identify, organize, search for, and filter resources. For more information, see [Tagging your AWS resources](#).

target variable

The value that you are trying to predict in supervised ML. This is also referred to as an *outcome variable*. For example, in a manufacturing setting the target variable could be a product defect.

task list

A tool that is used to track progress through a runbook. A task list contains an overview of the runbook and a list of general tasks to be completed. For each general task, it includes the estimated amount of time required, the owner, and the progress.

test environment

See [environment](#).

training

To provide data for your ML model to learn from. The training data must contain the correct answer. The learning algorithm finds patterns in the training data that map the input data attributes to the target (the answer that you want to predict). It outputs an ML model that captures these patterns. You can then use the ML model to make predictions on new data for which you don't know the target.

transit gateway

A network transit hub that you can use to interconnect your VPCs and on-premises networks. For more information, see [What is a transit gateway](#) in the AWS Transit Gateway documentation.

trunk-based workflow

An approach in which developers build and test features locally in a feature branch and then merge those changes into the main branch. The main branch is then built to the development, preproduction, and production environments, sequentially.

trusted access

Granting permissions to a service that you specify to perform tasks in your organization in AWS Organizations and in its accounts on your behalf. The trusted service creates a service-linked role in each account, when that role is needed, to perform management tasks for you. For more information, see [Using AWS Organizations with other AWS services](#) in the AWS Organizations documentation.

tuning

To change aspects of your training process to improve the ML model's accuracy. For example, you can train the ML model by generating a labeling set, adding labels, and then repeating these steps several times under different settings to optimize the model.

two-pizza team

A small DevOps team that you can feed with two pizzas. A two-pizza team size ensures the best possible opportunity for collaboration in software development.

U

uncertainty

A concept that refers to imprecise, incomplete, or unknown information that can undermine the reliability of predictive ML models. There are two types of uncertainty: *Epistemic uncertainty* is caused by limited, incomplete data, whereas *aleatoric uncertainty* is caused by the noise and randomness inherent in the data. For more information, see the [Quantifying uncertainty in deep learning systems](#) guide.

undifferentiated tasks

Also known as *heavy lifting*, work that is necessary to create and operate an application but that doesn't provide direct value to the end user or provide competitive advantage. Examples of undifferentiated tasks include procurement, maintenance, and capacity planning.

upper environments

See [environment](#).

V

vacuuming

A database maintenance operation that involves cleaning up after incremental updates to reclaim storage and improve performance.

version control

Processes and tools that track changes, such as changes to source code in a repository.

VPC peering

A connection between two VPCs that allows you to route traffic by using private IP addresses. For more information, see [What is VPC peering](#) in the Amazon VPC documentation.

vulnerability

A software or hardware flaw that compromises the security of the system.

W

warm cache

A buffer cache that contains current, relevant data that is frequently accessed. The database instance can read from the buffer cache, which is faster than reading from the main memory or disk.

warm data

Data that is infrequently accessed. When querying this kind of data, moderately slow queries are typically acceptable.

window function

A SQL function that performs a calculation on a group of rows that relate in some way to the current record. Window functions are useful for processing tasks, such as calculating a moving average or accessing the value of rows based on the relative position of the current row.

workload

A collection of resources and code that delivers business value, such as a customer-facing application or backend process.

workstream

Functional groups in a migration project that are responsible for a specific set of tasks. Each workstream is independent but supports the other workstreams in the project. For example, the portfolio workstream is responsible for prioritizing applications, wave planning, and collecting migration metadata. The portfolio workstream delivers these assets to the migration workstream, which then migrates the servers and applications.

WORM

See [write once, read many](#).

WQF

See [AWS Workload Qualification Framework](#).

write once, read many (WORM)

A storage model that writes data a single time and prevents the data from being deleted or modified. Authorized users can read the data as many times as needed, but they cannot change it. This data storage infrastructure is considered [immutable](#).

Z

zero-day exploit

An attack, typically malware, that takes advantage of a [zero-day vulnerability](#).

zero-day vulnerability

An unmitigated flaw or vulnerability in a production system. Threat actors can use this type of vulnerability to attack the system. Developers frequently become aware of the vulnerability as a result of the attack.

zombie application

An application that has an average CPU and memory usage below 5 percent. In a migration project, it is common to retire these applications.