

Tigase Advanced Clustering Strategy (ACS)

Tigase Team

Tigase Advanced Clustering Strategy (ACS)

Tigase Team

Table of Contents

1. Design and implementation	1
ACS	1
Design	1
2. Tigase ACS SM Installation	2
3. Tigase ACS SM Configuration	3
4. Tigase Advanced Clustering Strategy (ACS) Release Notes	4
Tigase Advanced Clustering Strategy (ACS) 3.2.0 Release Notes	4
Major Changes	4
All Changes	4
5. Supported components	5
Tigase Advanced Clustering Strategy for Multi User Chat (ACS-MUC)	5
Overview	5
Tigase ACS MUC Configuration	5
ACS MUC Strategies	5
Tigase Advanced Clustering Strategy for PubSub (ACS-PubSub)	7
Overview	8
Tigase ACS PubSub Configuration	8
ACS PubSub Strategies	8
Tigase Advanced Clustering Strategy for WorkGroup (ACS-WG)	11
Overview	11
Tigase ACS WorkGroup Configuration	11

Chapter 1. Design and implementation

ACS

ACS is our general purpose, commercial clustering strategy designed for more or less typical XMPP installations easily scaling to millions and beyond of online users without limit on cluster nodes. The load tests we have run over the code were included a user database with 100mln accounts and an average roster size of up to 150, but that's not the limit.

Design

The clustering strategy is based on sharing information between cluster nodes about online users. Who is online and where the user is connected. Communication between cluster nodes is processed with the highest priority to ensure minimal delays with online user data population. An efficient synchronization mechanism allows for a minimal traffic between cluster nodes and distributing accurate data about connecting and disconnecting users.

Chapter 2. Tigase ACS SM Installation

Tigase ACS SM component is by default provided with Tigase XMPP Server release (@-dist-max@ flavour of archive) so it's enough to enable it in the configuration. It can be also obtained from `tigase-ac` distribution package.

After downloading the archive it's simply matter of extracting it and copying contents of `jars/` directory of extracted archive to the `jars/` directory in `tigase-server/` installation directory, eg. under *nix systems (assuming the archive was downloaded to main Tigase Server directory):

```
tar --xf tigase-ac-${version}.tar.gz
cp --R tigase-ac-${version}/jars/ tigase-server/jars/
```

Chapter 3. Tigase ACS SM Configuration

In order to use Advanced Clustering Strategy, clustering mode first needs to be turned on:

```
'cluster-mode' = true
```

and then an ACS strategy needs to be enabled:

```
'sess-man' {  
    strategy (class: tigase.server.cluster.strategy.OnlineUsersCachingStrategy) {}  
}
```

Chapter 4. Tigase Advanced Clustering Strategy (ACS) Release Notes

Welcome to Tigase Advanced Clustering Strategy (ACS)! This is a feature release for with a number of fixes and updates.

Tigase Advanced Clustering Strategy (ACS) 3.2.0 Release Notes

Major Changes

- Deprecate `PartitionedStrategy` in ACS-PubSub

All Changes

- #acs-8 [<https://projects.tigase.net/issue/acs-8>]: Fix `NotAuthorizedException`: Session has not been yet authorised. in `OnlineUsersCachingStrategy`
- #acsmix-1 [<https://projects.tigase.net/issue/acsmix-1>]: Implement clustering support for MIX
- #acsmix-3 [<https://projects.tigase.net/issue/acsmix-3>]: Fix NPE in `DefaultPubSubLogic`
- #acsmix-4 [<https://projects.tigase.net/issue/acsmix-4>]: Fix NPE in `DefaultPubSubLogic.subscribersOfNotifications()`
- #acsmuc-23 [<https://projects.tigase.net/issue/acsmuc-23>]: Fix NPE in `ClusteredRoomStrategyV2`
- #acsmuc-25 [<https://projects.tigase.net/issue/acsmuc-25>]: Fix NPE in `OccupantChangedPresenceCmd`
- #acspubsub-20 [<https://projects.tigase.net/issue/acspubsub-20>]: Fix NPE in `pubsub-nodes-changed-cmd`
- #acspubsub-21 [<https://projects.tigase.net/issue/acspubsub-21>]: Fix Multiple notifications for single publication
- #acspubsub-22 [<https://projects.tigase.net/issue/acspubsub-22>]: Fix Presences informations are kept indefinitely
- #acspubsub-24 [<https://projects.tigase.net/issue/acspubsub-24>]: Fix `caps-changed-cmd` not processed correctly
- #acspubsub-25 [<https://projects.tigase.net/issue/acspubsub-25>]: Deprecate `PartitionedStrategy`
- #acspubsub-27 [<https://projects.tigase.net/issue/acspubsub-27>]: Review and improve clustering documentation

Chapter 5. Supported components

Tigase Advanced Clustering Strategy for Multi User Chat (ACS-MUC)

Tigase Team <team@tigase.com [mailto:team@tigase.com]> v3.4.0-SNAPSHOT, 2023-11-05 :numbered:

Overview

ACS for MUC allows seamless clustering of MUC rooms across Tigase XMPP server cluster installation. ACS for MUC is required for clustered MUC deployments. It offers various strategies to handle distribution of traffic and rooms across the cluster, which allows fine-tune configuration of the deployment to individual needs.

Tigase ACS MUC Configuration

In order to use ACS for MUC, main Advance Clustering Strategy is required. Once it's enabled, clustered version of MUC component will be selected by default during startup therefore it's not required to configure it explicitly (make sure no class is configured).

```
muc () { }
```

It's also possible to explicitly configure the class with the following configuration:

```
muc (class: tigase.muc.cluster.MUCComponentClustered) { }
```

With the above configuration default MUC clustering strategy will be used. In order to select different strategy you have to configure it's class in strategy bean within muc component bean:

```
muc () {  
    strategy (class: tigase.muc.cluster.ShardingStrategy) { }  
}
```

ACS MUC Strategies

ShardingStrategy

This is default clustering strategy used by Tigase ACS - MUC component. It should be used in most cases when we do not have small number of rooms with many occupants.

Short description

This is default clustering strategy used by Tigase ACS - MUC component. In this strategy MUC rooms are partitioned so each room is hosted only on one node. Every node contains full list of rooms with list of occupants available in each room and map which contains room address as a key and node as a value. If room is already opened (hosted) on some node, then node hosting this room is resolved using map of room to node. In other case node is selected by hash of room address using following algorithm

```
Math.abs(roomJid.hashCode()) % connectedNodes.size()
```


which gives a index of node on connected nodes list. So if node is not opened, then every cluster node should forward packets related to this room to the same node.

Connection to cluster

Once node connects to a cluster then map of hosted rooms and it's occupants is synchronized between connecting node and any other already connected node.

Disconnection from cluster

If node is disconnected from a cluster due to server failure, then every other node will send "kick" stanzas to every occupant of every room hosted on disconnected node.

If node is disconnected due to node shutdown then node which is shutting down will send "kick" stanzas on it's own, but it will notify every node about shutdown before disconnection from cluster.

Configuration

This is default strategy thus it's used if no strategy configuration is present, but if you wish to enable this strategy and keep it enabled even if default clustering strategy will change in the future then you need to set class property of strategy bean within muc component to `tigase.muc.cluster.ShardingStrategy`.

Example:

```
muc () {  
    strategy (class: tigase.muc.cluster.ShardingStrategy) {}  
}
```

ClusteredRoomStrategy

This is clustering strategy which can be used for by Tigase ACS - MUC component, which is recommended for installations with relatively few rooms but rooms itself having a lot of occupants.

Short description

In this strategy MUC rooms are persistent and each room is hosted on every node. Every node contains full list of rooms (as they are persistent). Room on each node has knowledge only about occupants which joined room on this node. If remote user has joined room on one node and then it's packets are delivered by S2S to other node, then packets will be forwarded to node to on which user joined room. Packets from user are processed by node on which they joined to room and notifications about joining room, leaving room, change of presences or about new message are sent to all other nodes and those other nodes are responsible for delivering proper notifications to users which joined room on them.

Connection to cluster

Once node connects to a cluster then map of occupants and their rooms are synchronized with other nodes.

Disconnection from cluster

If node is disconnected from a cluster, then every other node will send "kick" stanzas to every occupant of every room hosted on disconnected node.

Configuration

To enable this strategy you have to set class property of strategy bean within muc component to `tigase.muc.cluster.ClusteredRoomStrategy`.

Example:

```
muc () {  
    strategy (class: tigase.muc.cluster.ClusteredRoomStrategy) {}  
}
```

ClusteredRoomStrategyV2

This is clustering strategy which can be used for by Tigase ACS - MUC component, which is recommended for installations with relatively few rooms but rooms itself having a lot of occupants - contains improvements over ClusteredRoomStrategy

Short description

In this strategy MUC rooms are persistent and each room is hosted on every node. Every node contains full list of rooms (as they are persistent). Room on each node has knowledge only about occupants which joined room on this node. If remote user has joined room on one node and then it's packets are delivered by S2S to other node, then packets will be forwarded to node to on which user joined room. Packets from user are processed by node on which they joined to room and notifications about joining room, leaving room, change of presences or about new message are sent to all other nodes and those other nodes are responsible for delivering proper notifications to users which joined room on them.

This version contains improvements over the section called "ClusteredRoomStrategy" which leads to improved performance and reduced traffic over cluster connection due to changes in logic of processing packets.

Connection to cluster

Once node connects to a cluster then map of occupants and their rooms are synchronized with other nodes.

Disconnection from cluster

If node is disconnected from a cluster, then every other node will send "kick" stanzas to every occupant of every room hosted on disconnected node.

Configuration

To enable this strategy you have to set class property of strategy bean within muc component to `tigase.muc.cluster.ClusteredRoomStrategyV2`.

Example:

```
muc () {  
    strategy (class: tigase.muc.cluster.ClusteredRoomStrategyV2) {}  
}
```

Tigase Advanced Clustering Strategy for Pub-Sub (ACS-PubSub)

Tigase Team <team@tigase.com [mailto:team@tigase.com]> v3.4.0-SNAPSHOT, 2023-11-05 :numbered:

Overview

ACS for PubSub allows seamless clustering of PubSub nodes across Tigase XMPP server cluster installation. ACS for PubSub is required for clustered PubSub deployments. It offers various strategies to handle distribution of traffic and nodes across the cluster, which allows fine-tune configuration of the deployment to individual needs.

Tigase ACS PubSub Configuration

In order to use ACS for PubSub, main Advance Clustering Strategy (ACS) is required. Once it's enabled, clustered version of PubSub component will be selected by default during startup therefore it's not required to configure it explicitly (make sure no other class is configured).

```
pubsub () {}
```

It's also possible to explicitly configure the class with the following configuration:

```
pubsub (class: tigase.pubsub.cluster.PubSubComponentClustered) {}
```

With the above configuration default ACS PubSub clustering strategy will be used. In order to select different strategy you have to configure its class for strategy bean within pubsub component bean:

```
pubsub () {  
    strategy (class: tigase.pubsub.cluster.PartitionedStrategy) {}  
}
```

ACS PubSub Strategies

ACS-PubSub Partitioned Strategy

Warning

This strategy is now deprecated.

Short description

This is the simplest strategy that can be used by Tigase ACS - PubSub component in which particular PubSub node is handled by only one cluster node.

Description of processing

In this strategy all configuration of nodes of the same PubSub service JID is done by the same node of a cluster which is dynamically selected based on hash of service JID and number of connected nodes within cluster. After any change to node configuration is done, then node established for processing manipulation of this particular PubSub node is notified about details of the PubSub node and detailed changes made to its configuration. Remaining nodes do not require to know about changes.

Presences sent from users to PubSub service will be handled on the local node of the cluster and then distributed to every node of a cluster as events.

Messages sent from users to PubSub service will be handled on the local node of the cluster.

Presence and IQ stanzas will be processed according to rules outlined below.

Rules of processing packets:

- `presence` - packets are processed on the local cluster node and then other cluster nodes are notified about changed presence as each cluster node is responsible for handling different PubSub nodes
- `message` - is always processed locally
- `iq` - cluster node which will process this packet is selected based on following rules:
 - CAPS query responses will be processed on the local node
 - non-PubSub packets (no pubsub subelement) are processed on local node
 - PubSub related packets without PubSub node name or packets that change PubSub node configuration (and contains node name and one of the following subelements: `create`, `configure`, `default`, `delete`) are processed on cluster node selected on hashcode derived from `to` attribute (service JID) and number of cluster nodes (this is done deal with concurrency issues between configuration changes)
 - remaining PubSub packets (non-configuration and containing node name) are processed on cluster node selected based on `to` attribute of packet and name of PubSub node

Note

This strategy for every packet should force processing of packet on only one cluster node

Note

This strategy will react on PubSub configuration node change and will send notification to cluster node responsible for processing items for PubSub node (selected based on `to` attribute of packet, name of PubSub node) to trigger cached PubSub node configuration refresh

Note

Result of packets generated on remote node should **not be** filtered, so if packet from one cluster node was forwarded to other cluster node for processing, then response packet should not be filtered when this PubSub clustering strategy is used.

ACS-PubSub Clustered Node Strategy

Short description

This strategy is used by default by Tigase ACS - PubSub component in which each PubSub node is handled on every cluster node but each cluster node will contain only partial information about user connections. This way strategy is better suited for deployments with PubSub nodes having a lot of subscribers. The benefit of using ClusterNodeStrategy in this case is reduced network traffic on cluster connections, as most of notifications and retrieval of items will be handled on the same cluster node to which user is connected.

Description of processing

In this strategy almost all packets are processed on the local cluster node. Only packets related to settings default options (such as pubsub packets containing subelements `options` or `default`) will be forwarded to all cluster nodes and processed on each of them. If packets will result in the node configuration, subscriptions or affiliations changes other nodes will be notified to refresh node configuration/subscriptions/affiliations.

Presences sent from users to PubSub service will be handled on the local node of the cluster and then distributed to every node of a cluster as events.

Messages sent from users to PubSub service will be handled on the local node of the cluster.

If every other IQ stanza sent to PubSub service which is does not change configuration of a node and is item publication stanza then this stanza will always be processed on local node (as data retrieval/removal may be done only on one node as items are not cached).

IQ stanzas which are stanza responsible for publication are processed on the local cluster node. In this case PubSub will be able to properly send notifications (as notification are generated always on local user node, if user is connected, which reduces cluster network traffic).

Rules of processing packets:

- `presence` - packets are delivered to every cluster node as each cluster node is responsible for handling different PubSub nodes
- `message` - is always processed locally
- `iq` - cluster node which will process this packet is selected based on following rules:
 - CAPS query responses will be processed on the local node
 - non-PubSub packets (no `pubsub` subelement) are processed on local node
 - PubSub related packets without PubSub node name are processed on the local node. If packets will result in the node configuration changes, other nodes will be notified to refresh node configuration.
 - PubSub related packets containing subelements named `options` or `default` will be forwarded to all connected cluster nodes
 - remaining PubSub packets are processed on the local cluster node. If packets will result in the node configuration changes, other nodes will be notified to refresh node configuration.

Note

This strategy will react to PubSub configuration node change and will send notification to every cluster node (as every cluster node is responsible for processing items for every PubSub node) to refresh particular PubSub node configuration.

Note

Result of processing of packets generated on remote node should **be** filtered if packet was also processed on local node, so if packet from one cluster node was forwarded to other cluster node for processing but was not processed locally, then response packet should be filtered when this PubSub clustering strategy is used.

In a nutshell

`Clustered Node Strategy` distributes processing of nodes across cluster by processing all requests locally - each cluster node has complete knowledge of all PubSub nodes and processed requests for all of them, but generates notifications only for users connected to this particular node.

`Partitioned Strategy` assigns PubSub node to cluster node and handle all processing on that particular cluster node (configuration change, generating notification packets).

Tigase Advanced Clustering Strategy for Work-Group (ACS-WG)

Tigase Team <team@tigase.com [mailto:team@tigase.com]> v3.4.0-SNAPSHOT, 2023-11-05

Overview

ACS for WorkGroup allows seamless clustering of WorkGroup nodes across Tigase XMPP server cluster installation. ACS for WorkGroup is required for clustered WorkGroup deployments. It offers various strategies to handle distribution of traffic and nodes across the cluster, which allows fine-tune configuration of the deployment to individual needs.

Tigase ACS WorkGroup Configuration

In order to use ACS for WorkGroup, main Advance Clustering Strategy (ACS) is required. Once it's enabled, clustered version of WorkGroup component will be selected by default during startup therefore it's not required to configure it explicitly (make sure no other class is configured).

```
wg () {}
```

It's also possible to explicitly configure the class with the following configuration:

```
wg (class: tigase.workgroupqueues.cluster.WorkgroupQueuesClusteredComponent) {}
```

With the above configuration default ACS WorkGroup clustering strategy will be used. In order to select different strategy you have to configure it's class in strategy bean within wg component bean:

```
wg () {  
    strategy (class: tigase.workgroupqueues.cluster.ClusteredStrategy) {}  
}
```