HANA BW Scale Out

Customer

HANA BW Scale Out Concept: Overview 1/2 – Setup with 512 GB Nodes

Master Node

- 512 GB of memory
- holds system tables, all row tables
- DDL execution, global locks, database metadata

Up to 14 Slave Nodes

- 512 GB of memory each
- master data tables will be distributed across slave nodes
- fact, DSO and PSA tables are partitioned and distributed

Standby Node(s) (optional)

- 512 GB of memory
- can take over role of master node or of one of the slaves for failover

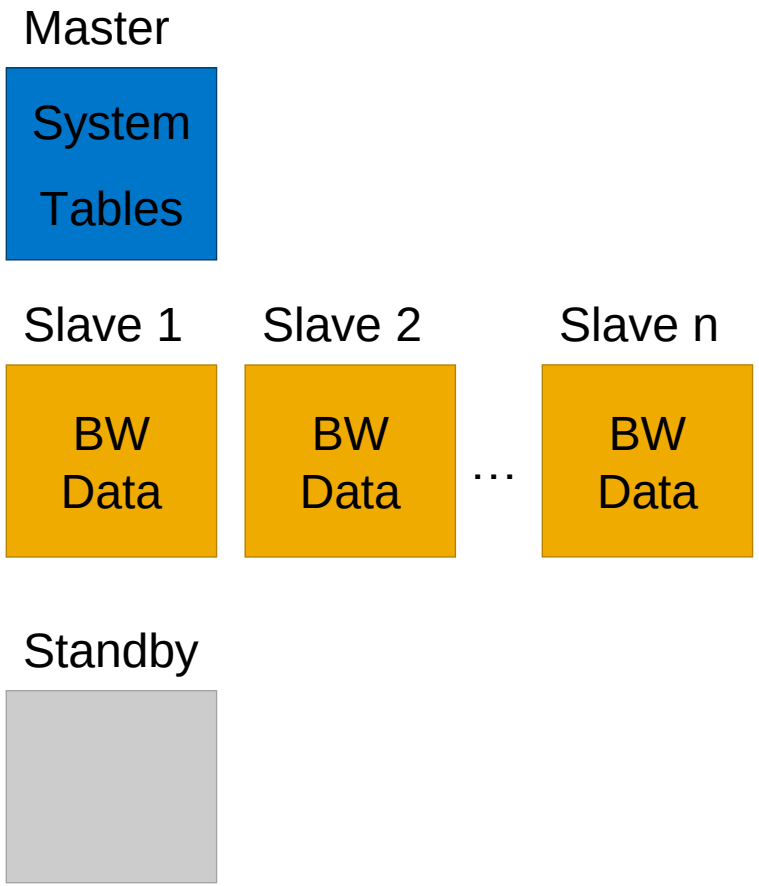

HANA BW Scale Out Concept: Overview 2/2 – Setup with 1 TB Nodes*

Master Node

- 1 TB of memory
- holds system tables, all row tables
- DDL execution, global locks, database metadata

Up to 14 Slave Nodes

- 1 TB of memory each
- master data tables will be distributed across slave nodes
- fact, DSO and PSA tables are partitioned and distributed

Standby Node(s) (optional)

- 1 TB of memory
- can take over role of master node or of one of the slaves for failover

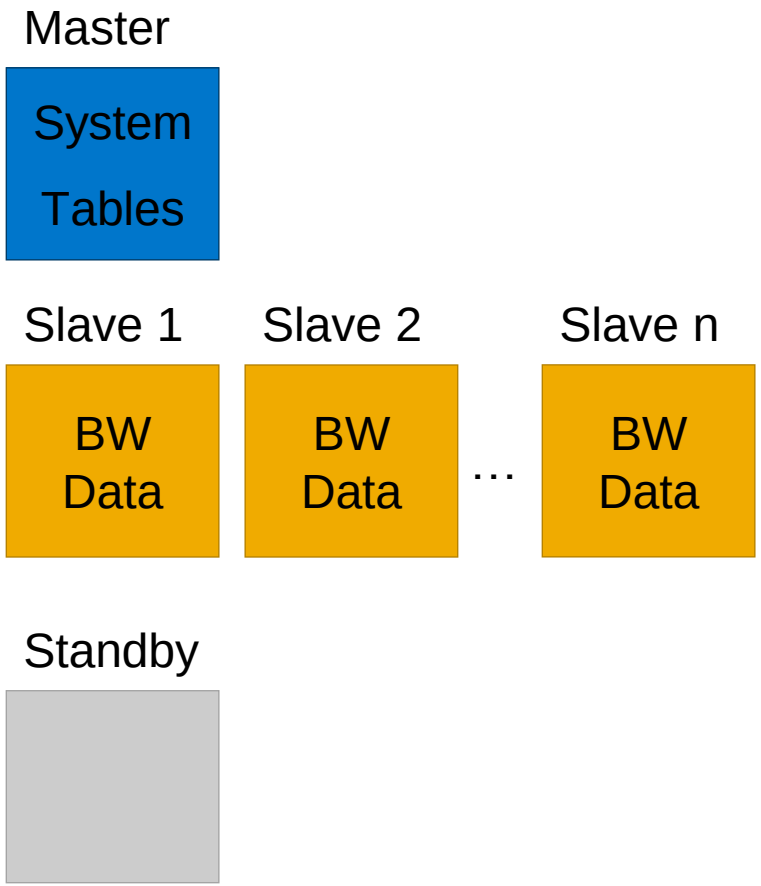

* For Hardware vendors that support it

HANA BW Scale Out Concept: Facts & Features

- **Master node** will handle system load and transactional load: ABAP system tables and general operational data of the BW are stored on the master node. DDL statements are executed on this node, global locks are acquired here.
- **Slave nodes** will handle OLAP queries as well as loading/staging/activation/merging. BW data (master data + cubes/DSOs/PSAs) is distributed across all slaves. This ensures a balanced utilization of the available CPU and memory resources.
- With **statement routing** and **statement distribution**, all slaves nodes are used directly by the application server, minimizing network latency and bandwidth exhaustion.
- **Easy extensibility**: When a new slave is added, BW data can quickly be reorganized to fit the new system.
- **Symmetric hardware** across the system: same hardware can be used as master, slave or standby node. Standby node can take over the role of failing master or failing slave node.

HANA BW Scale Out Concept: Table Distribution on Slave Nodes

Master data tables are distributed across slave nodes. Huge master data tables can be partitioned.

Fact, DSO and PSA tables are partitioned and distributed across slave nodes. To accommodate more data, the distributed system can easily be extended by adding new slave nodes.

If another HANA scenario is used on the same system, its tables are also distributed across the slaves with the same benefits.

HANA BW Scale Out Concept:
Pure High Availability Solution with a Standby Node

Worker Node

- 512 GB or 1 TB of memory
- single worker node

Standby Node

- same amount of memory as worker node
- can take over role of worker node for failover

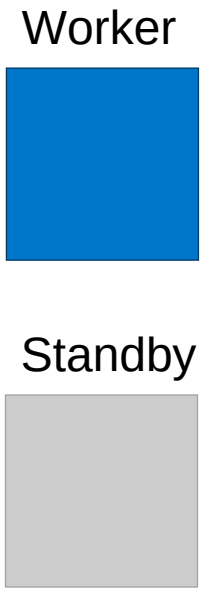

HANA BW Scale Out Concept: Exception for systems with a single slave node

Please see note 1702409 why this setup is not recommended.

Master Node

- holds system tables, all row tables
- DDL execution, global locks, database metadata

Slave Node

- holds master data, fact, DSO and PSA tables

Standby Node (optional)

- can take over role of master or slave node for failover

For special requirements, master data tables may optionally be moved to the master node. This can lead to performance degradation. Master data tables may be moved manually (there is no official tool support) or with the help of a consulting note.