

A vertical strip on the left side of the slide showing a close-up of a dark, layered rock face with some blue mineral veins.

# See What's New in MemSQL 6.7: The No-Limits Database

Available as of 11/6/18

# Today's Speaker & Agenda



**Eric Hanson**

Principal Product Manager  
MemSQL


## Agenda

- What is MemSQL?
- New monitoring and management tooling
- Introduction to MemSQL Studio
- How to take advantage of the new query improvements
- New storage and loading features
- MemSQL free tier overview

## Questions?

Submit via the Q&A or chat function in the lower right hand side of the console

# MemSQL Release Timeline



# MemSQL 6.5

(available 7/18)

- Faster Queries
- Automated Workload Management
- Extensible Pipelines
- Full-Text Search
- Scalable Multi-Tenancy

# MemSQL 6.7

Improve usability, compatibility, and performance

## MemSQL 6.7

- Free tier
- Developer forums
- Improved ease of use
- Faster queries
- New manageability



# New Free Tier for Simplified Evaluations

- Full featured product
- Deploy to production
- Install on clusters up to 128GB of RAM
- Unlimited disk usage
- Update to Enterprise Edition with license key change



# Portal and Forums

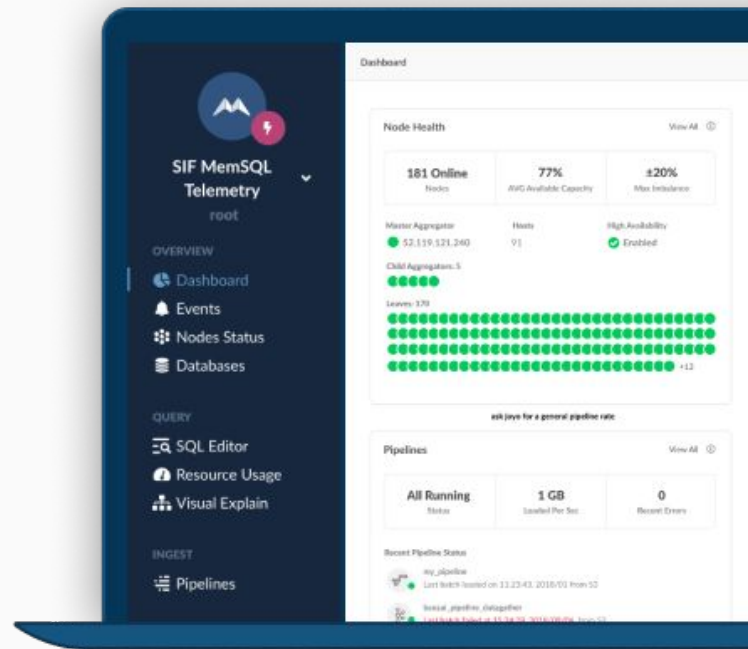
[portal.memsql.com](https://portal.memsql.com)

[memsql.com/forum](https://memsql.com/forum)



# MemSQL Studio

- Primary UX tool for managing and monitoring
- Delivers at-a-glance cluster health statistics, query troubleshooting, and query editor
- Requires MemSQL 6.5 or greater
- Independent of the cluster
- Provides multi-cluster access
- Faster update cadence



# Management Tools



MemSQL has a new installable package that contains a set of tools for deploying, managing, and troubleshooting MemSQL

## **memsql-deploy**

Installs MemSQL software

## **memsql-admin**

Run commands such as create-node, set-license, add-leaf

## **memsql-report**

Runs a set of checks on hardware, configurations, and clusters

## **memsql-toolbox-config**

Tool to simplify setup of ip-addresses of nodes

# Faster Query Performance

## Star Join Improvements

- MemSQL 6.0 delivered faster queries on compressed data - ie. Trillion row per second scan with SIMD (on-chip query vectorization)
- New 6.7 improvements extend SIMD to star join queries, yielding up to 3-100x faster queries over 6.5



## 101X speedup -- 780 million rows/sec/core

```
select d_daynuminweek, d_dayofweek, count(*) as c
from media_view_fact f, date_dim d
where f_datekey = d_datekey
group by 1, 2
order by 1 asc;
```


## 3X speedup

```
select d_dayofweek, count(*)
from media_view_fact, date_dim, user_dim, item_dim
where f_datekey = d_datekey
and f_userkey = u_userkey
and f_itemkey = i_itemkey
and i_itemuri regexp 'http://www.c000[0-5]. '
and u_zipcode like '002%' and
d_year = 1997 and d_month = "January"
group by d_dayofweek;
```

# Demo

## Star Join Performance in Studio

## 6.7 and 6.5 Time (seconds) on Star Schema Benchmark





# Faster Query Performance

## Additional improvements

- Just-in-time compilation of queries (experimental)
- IN-list performance improvement
- Columnstore range filter improvements
- Fast sampled rowstore scans
- INSERT\_ALL() function



# Query Optimization and Statistics Improvements

- Row store sampling with true random sample
- IN-list filter cardinality estimation with histograms
- Faster compilation for some queries (up to 6X)
- New information\_schema views to help with tuning
  - mv\_query\_column\_usage
  - mv\_aggregated\_column\_usage

## New Stats Views

### `mv_query_prospective_histograms`

Shows what queries could have used histograms on a particular column but they weren't available

### `optimizer_statistics`

Shows what statistics are available for each column that has any kind of statistics, and when they were last updated

## Additional Usability Improvements

- Expanded Resource Governor adds CPU and concurrency, in addition to RAM
- Additional query syntax, new extensibility, and loading
  - Dynamic SQL
  - Added ROLLUP and CUBE syntax

# Load, Backup, and Restore Improvements

- Adaptive data compression for LOAD DATA (up to 2X performance gain)
- Tar backup to Amazon S3
- Backup/restore to/from Azure Blob Store
- Enhanced metadata for backups in MV\_BACKUP\_HISTORY view as well as backup files
- Enhanced validation prior to restore or disaster recovery replicate operations, to ensure there is enough space on the target system
- SELECT INTO OUTFILE on S3
- JSON and Avro support for LOAD DATA and pipelines

# Expanded Compatibility with Other Technology

- Microstrategy support
- SAS Access certification
- Informatica Connector
- Streamsets support
- New Oracle-compatible functions
  - `nvl()`, `to_char()`, `to_timestamp()`, `regexp_instr()` etc.





# Q&A

[memsql.com](https://memsql.com)

# Thank You



Learn more about 6.7 at [memsql.com/blog/memsql67](https://memsql.com/blog/memsql67)



Try us for free at [portal.memsql.com](https://portal.memsql.com)

Questions? Email us at [team@memsql.com](mailto:team@memsql.com)