

ORACLE®

# Oracle Database 12c Rel. 2 Cluster Health Advisor - Deep Dive How it Works and How to Use It

Mark V. Scardina – Director of Product Management  
Oracle Autonomous Health Framework and  
Oracle Quality of Service Management  
Ankita Khandelwal – Product Manager  
Oracle Autonomous Health Framework

# Agenda

- 1 Introduction
- 2 CHA Architecture and Operation Details
- 3 Using CHA from the Command Line
- 4 Using CHA from EMCC for Alerts and Corrective Actions
- 5 Using the CHA GUI to Perform Root-Cause Analysis
- 6 Calibrating CHA to your RAC deployment
- 7 Q & A – Further Information

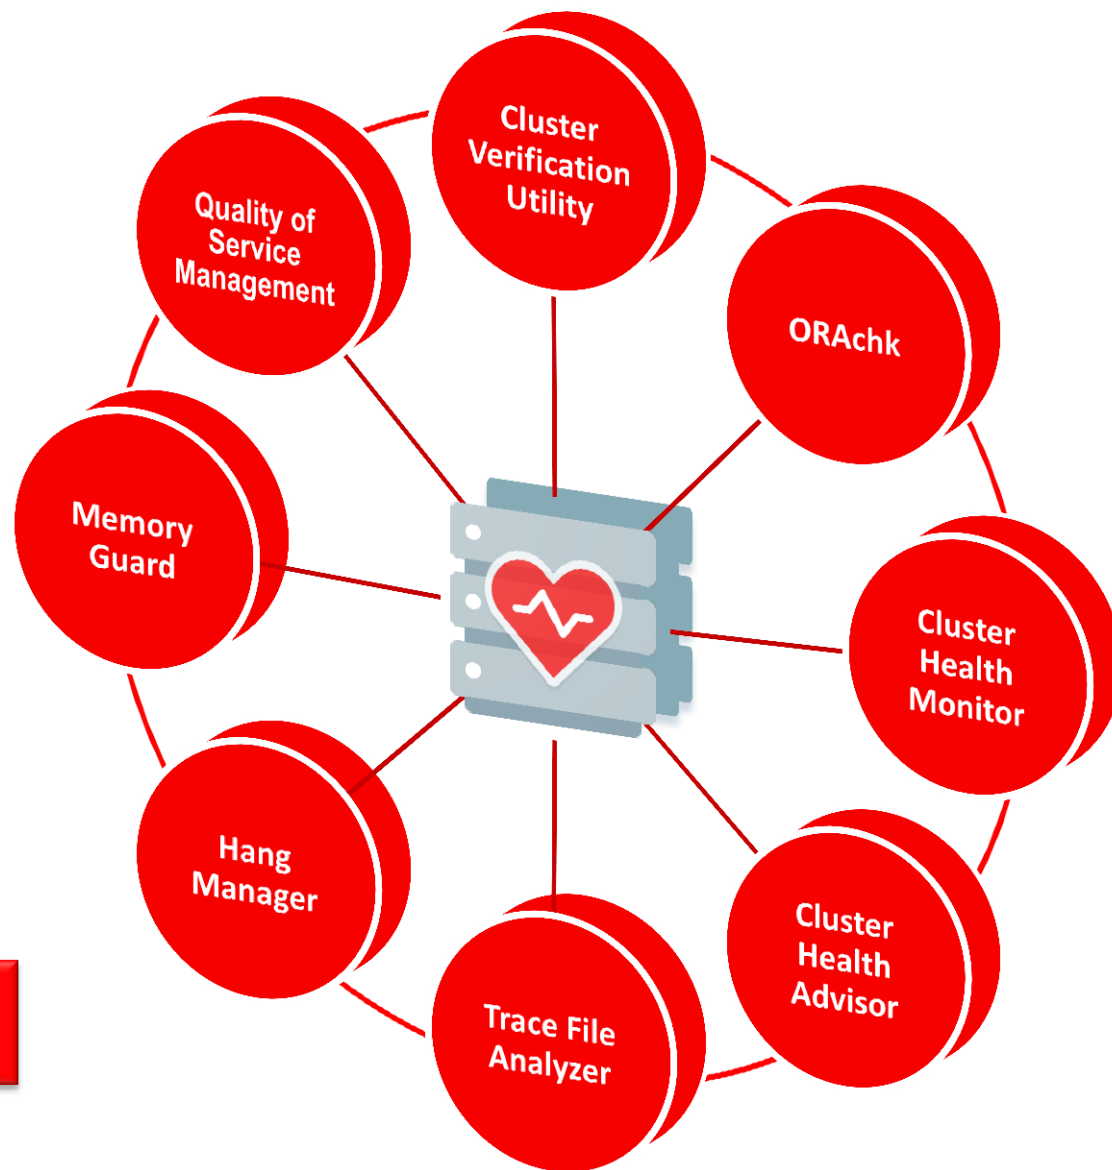
# Agenda

- 1 Introduction
- 2 CHA Architecture and Operation Details
- 3 Using CHA from the Command Line
- 4 Using CHA from EMCC for Alerts and Corrective Actions
- 5 Using the CHA GUI to Perform Root-Cause Analysis
- 6 Calibrating CHA to your RAC deployment
- 7 Q & A – Further Information

Introducing...

# Oracle 12c Autonomous Health Framework

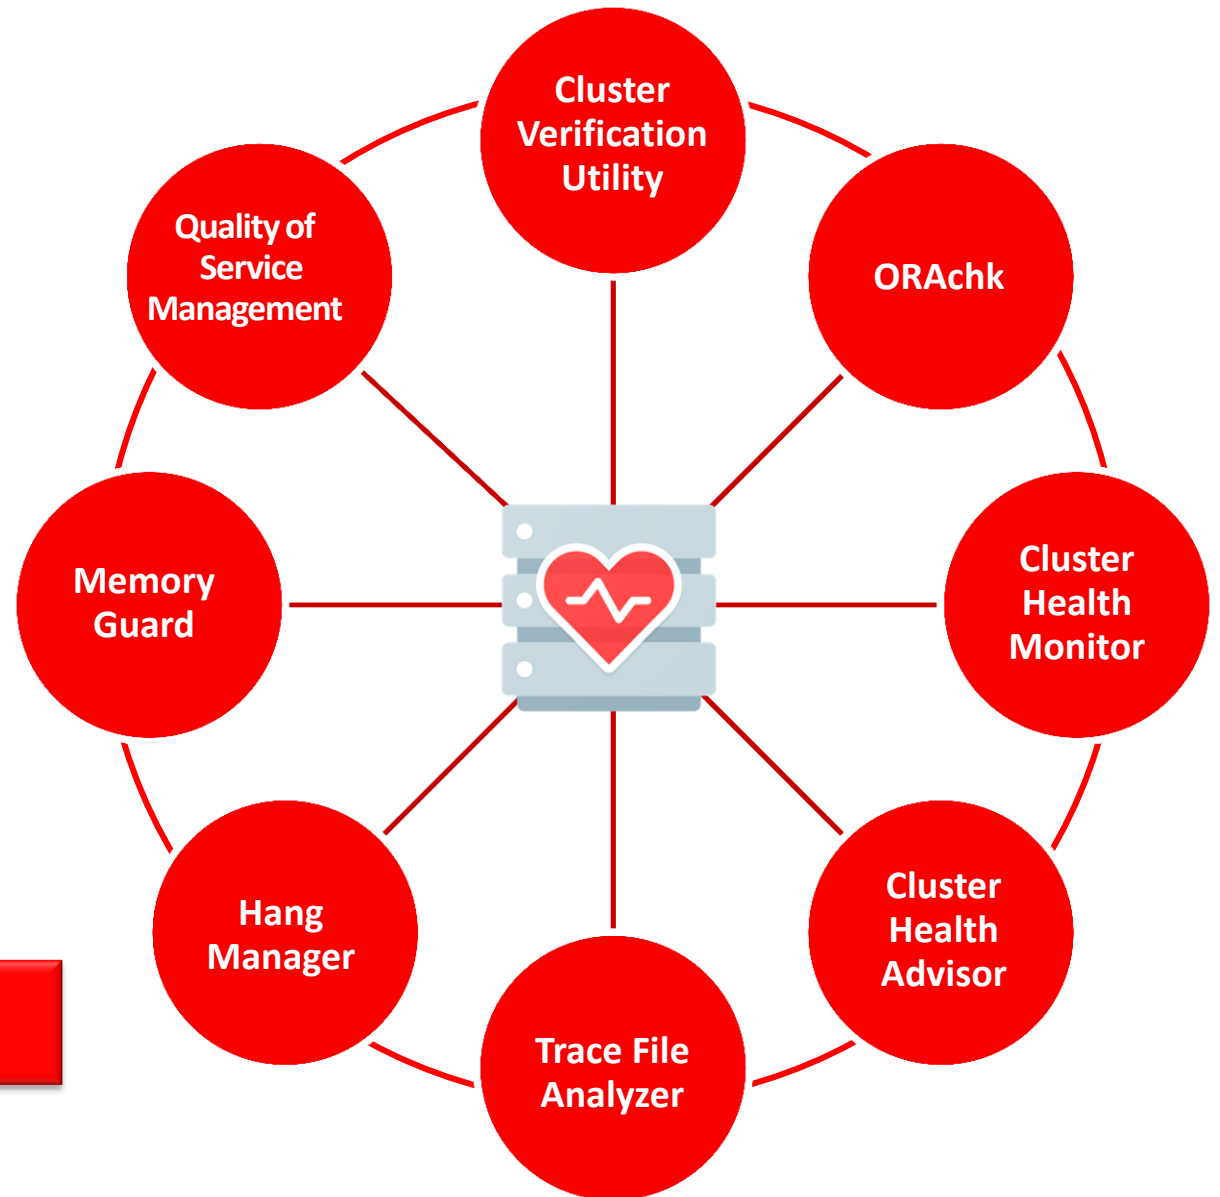
**Powered By Machine Learning**



Introducing...


# Oracle 12c Autonomous Health Framework

**Working For You Continuously**



Discovers Potential Cluster  
& DB Problems - Notifies  
with Corrective Actions

## Oracle 12c Cluster Health Advisor



# Introducing Oracle 12c Cluster Health Advisor

## Proactive Health Prognostics System

- Real time monitoring of Oracle RAC database systems and their hosts
- Early detection of impending as well as ongoing system faults
- Diagnoses and identifies the most likely root causes
- Provides corrective actions for targeted triage.
- Generates alerts and notifications for rapid recovery

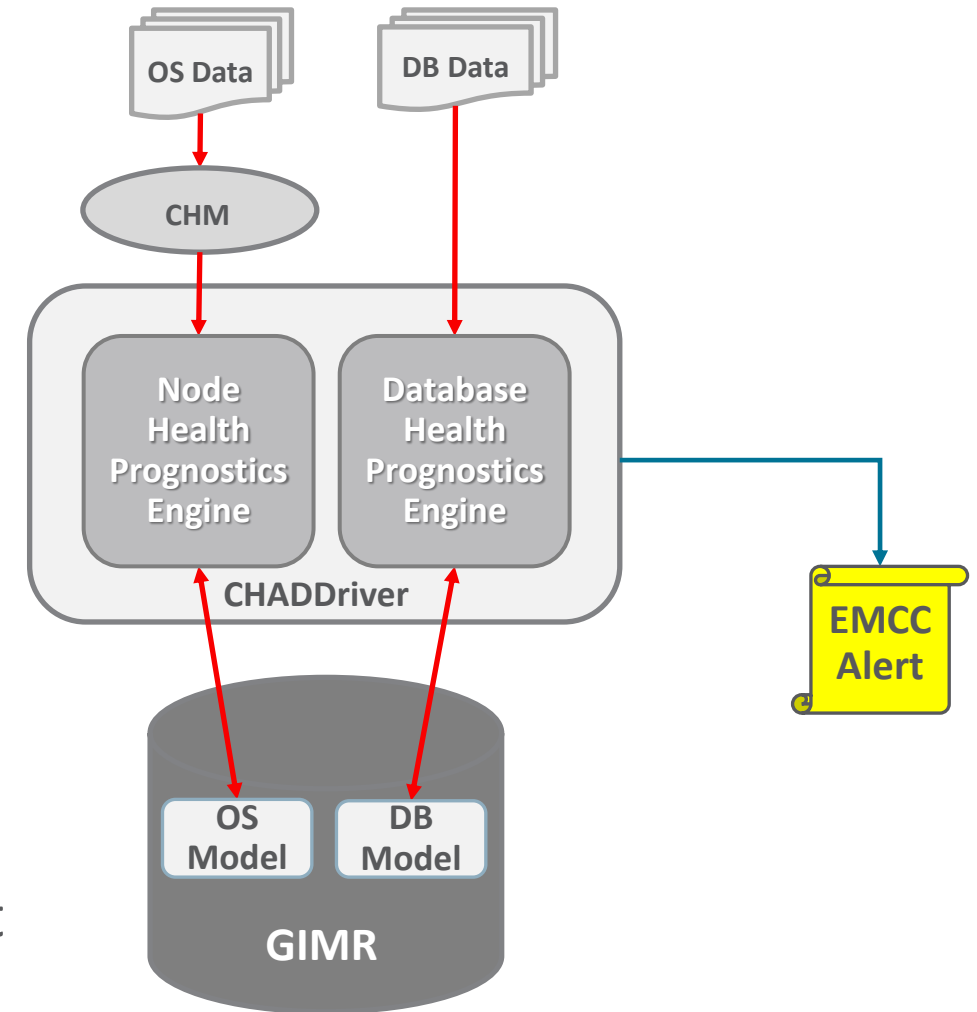


# Agenda

- 1 Introduction
- 2 CHA Architecture and Operation Details
- 3 Using CHA from the Command Line
- 4 Using CHA from EMCC for Alerts and Corrective Actions
- 5 Using the CHA GUI to Perform Root-Cause Analysis
- 6 Calibrating CHA to your RAC deployment
- 7 Q & A – Further Information

# Cluster Health Advisor (CHA) Architecture Overview

- **cha** – Cluster node resource
- Single Java `oracle.cha.server.CHADriver` daemon per node
- Reads Cluster Health Monitor data directly from memory
- Reads DB ASH data from SMR w/o DB connection
- Uses OS and DB models and data to perform prognostics
- Stores analysis and evidence in the GI Management Repository
- Sends alerts to EMCC Incident Manager per target



# Cluster Health Advisor - Scope of Problem Detection

## Best Effort Immediate Guided Diagnosis

- Over 30 node and database problems have been modeled
- Over 150 OS and DB metric predictors identified
- Problem network model created based upon its signature
- Problem Detection in 12.2.0.1 includes
  - Interconnect , Global Cache and Cluster Problems
  - Host CPU and Memory , PGA Memory stress
  - IO and Storage Performance issues
  - Reconfiguration and Recovery issues
  - Workload and Session abnormal variations

# Cluster Health Advisor

## *Data Sources and Data Points*

A CHA *Data Point* contains > 150 signals (statistics and events) from *multiple sources*


OS, ASM , Network         DB ( ASH, AWR session, system and PDB statistics )

Time	CPU	ASM IOPS	Network % util	Network_Packets Dropped	Log file sync	Log file parallel write	GC CR request	GC current request	GC current block 2-way	GC current block busy	Enq: CF - contention	...
15:16:00	0.90	4100	13%	0	2 ms	600 us	0	0	300 us	1.5 ms	0	

Statistics are collected at a **1 second internal sampling** rate , synchronized, smoothed and aggregated to a Data Point **every 5 seconds**

# Models Capture all Normal Operating Modes

*Models Capture the Dynamic Behavior of all Normal Operation*



In-Memory Reference Matrix  
(Part of “Normality” Model)

IOPS	####	2500	4900	800	####
User Commits	####	10000	21000	4400	####
Log File Parallel Write	####	2350	4100	22050	####
Log File Sync	####	5100	9025	4024	####
...	...	...	...	...	...

A model captures *the normal load phases* and their statistics over time , and thus the characteristics for all load intensities and profiles . During monitoring , *any data point similar* to one of the vectors is NORMAL. One could say *that the model REMEMBERS the normal operational dynamics over time*

# Cluster Health Advisor

## CHA Model: Find Similarity with Normal Values

In-Memory Reference Matrix  
(Part of “Normality” Model)

IOPS	####	2500	4900	800	####
User Commits	####	10000	21000	4400	####
Log File Parallel Write	####	2350	4100	22050	####
Log File Sync	####	5100	9025	4024	####
...	...	...	...	...	...

Observed values  
(Part of a Data Point)

10500
20000
4050
10250
...

Observed -  
Predicted =

Residual Values  
(Part of a Data Point)

5600
-1000
-50
325
...

CHA estimator/predictor (ESEE): *“based on my normality model, the value of IOPS should be in the vicinity of ~ 4900, but it is reported as 10500, this is causing a residual of ~ 5600 in magnitude”,*

CHA fault detector: *“such high magnitude of residuals should be tracked carefully! I’ll keep an eye on the incoming sequence of this signal IOPS and if it remains deviant I’ll generate a fault on it”.*

# Cluster Health Advisor

## *Inline and Immediate Fault Detection and Diagnostic Inference*

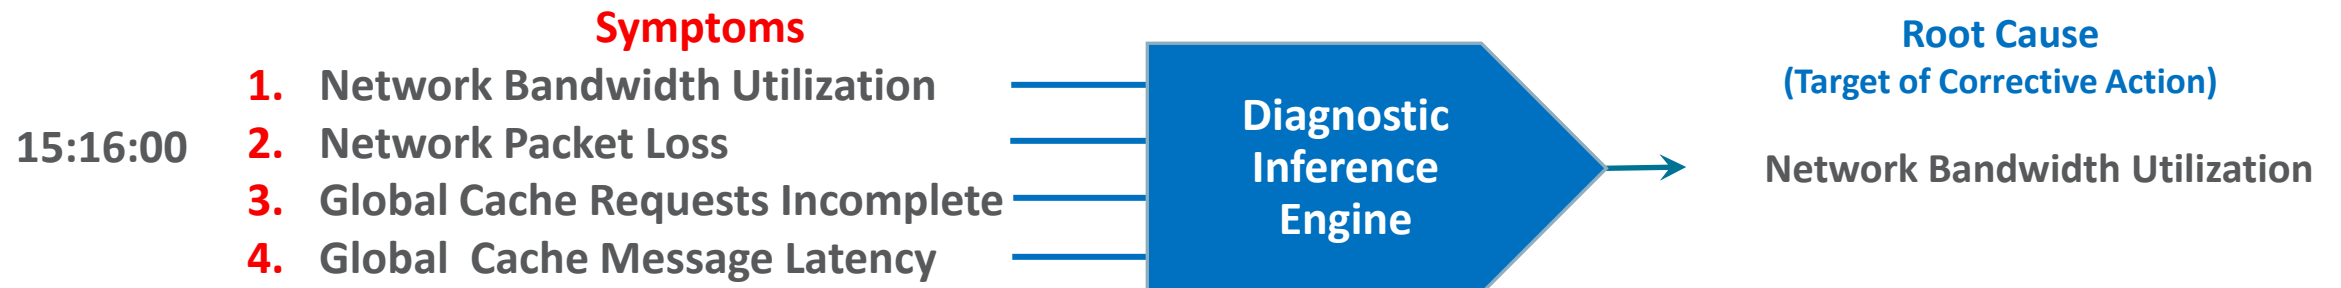
Input : Data Point at Time  $t$

Time	CPU	ASM IOPS	Network % util	Network_Packets Dropped	Log file sync	Log file parallel write	GC CR request	GC current request	GC current block 2-way	GC current block busy	Enq: CF - contention	...
15:16:00	0.90	4100	88%	105	2 ms	600 us	504 ms	513 ms	2 ms	5.9 ms	0	

### Fault Detection and Classification

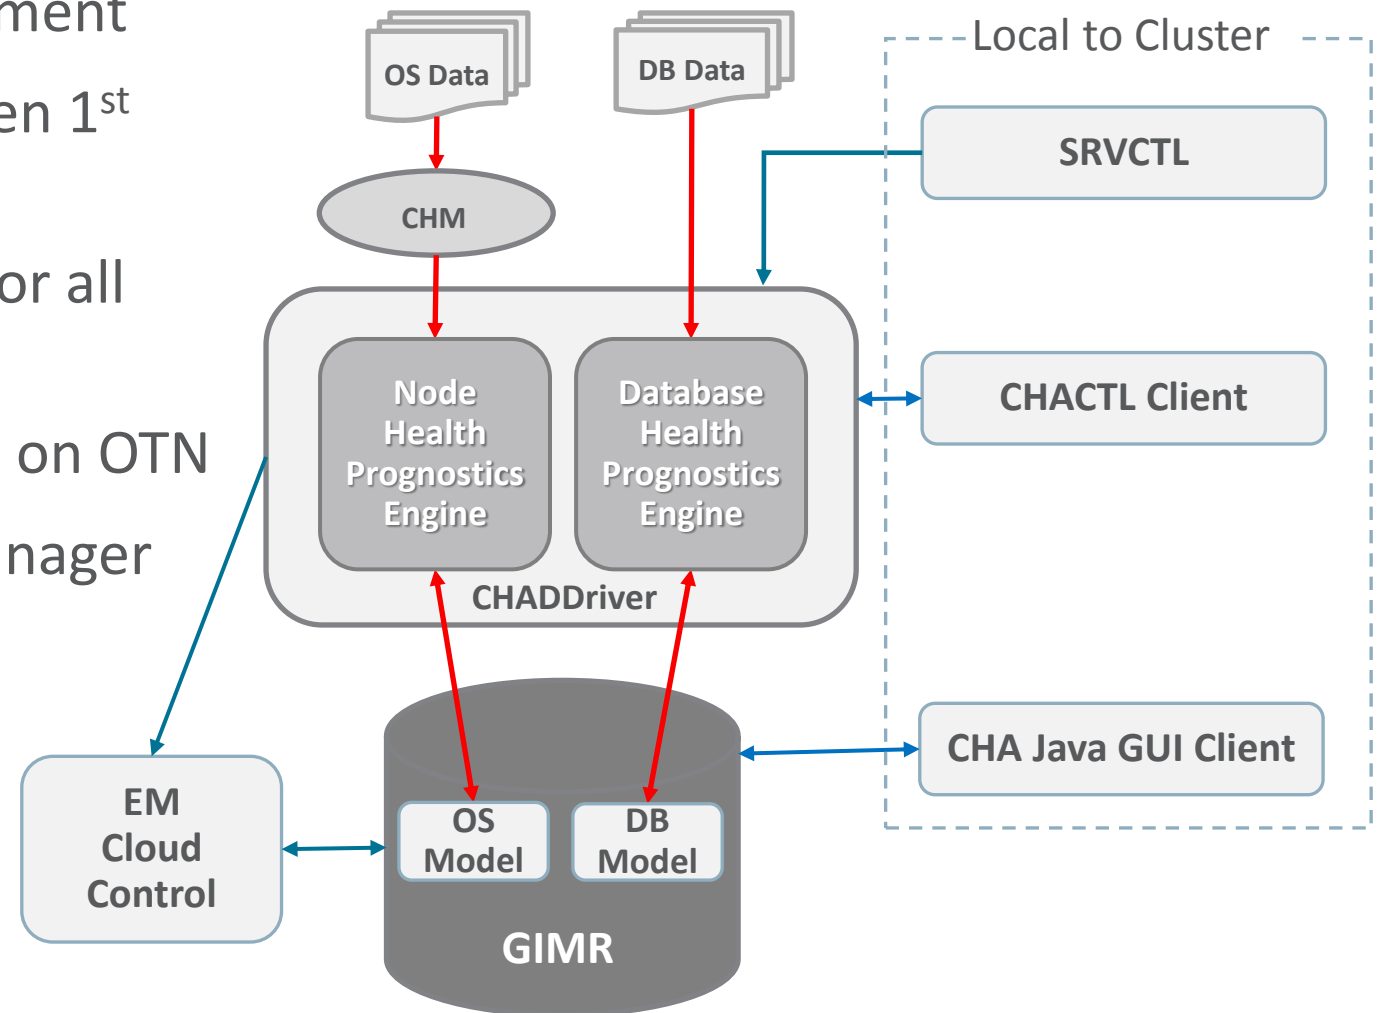
15:16:00	OK	OK	HIGH 1	HIGH 2	OK	OK	HIGH 3	HIGH 3	HIGH 4	HIGH 4	OK	
----------	----	----	-----------	-----------	----	----	-----------	-----------	-----------	-----------	----	--

### Diagnostic Inference



# Cluster Health Advisor (CHA) Operation Overview

- SRVCTL lifecycle daemon management
- Enabled by default - Activates when 1<sup>st</sup> RAC instance starts
- New CHACTL command line tool for all local operations
- Java Interactive GUI Tool available on OTN
- Integrated into EMCC Incident Manager and notifications
- Monitoring has no impact on DB performance or availability





# Agenda

- 1 Introduction
- 2 CHA Architecture and Operation Details
- 3 Using CHA from the Command Line
- 4 Using CHA from EMCC for Alerts and Corrective Actions
- 5 Using the CHA GUI to Perform Root-Cause Analysis
- 6 Calibrating CHA to your RAC deployment
- 7 Q & A – Further Information

# Using CHA From the Command Line

## Overview

- CHA is enabled by default
- Autonomously monitors nodes once a RAC DB starts in the cluster
- RAC or RAC One Node Database must be explicitly monitored – opt in
- CHACTL Command Line supports
  - Start/Stop Monitoring
  - Diagnosing Health Issues and Corrective Actions
  - Model Calibration, Activation and Lifecycle Management
  - CHA Repository Data Lifecycle Management

# Cluster Health Advisor – Command Line Operations

## Monitoring Your Databases and Nodes with CHACTL

Enable CHA monitoring on RAC database with optional model

```
$ chactl monitor database -db oltpacdb [-model model_name]
```

Enable CHA monitoring on RAC database with optional verbose

```
$ chactl status -verbose
```

```
monitoring nodes svr01, svr02 using model DEFAULT_CLUSTER
```

```
monitoring database oltpacdb, instances oltpacdb_1, oltpacdb_2 using model DEFAULT_DB
```

# CHA Command Line Operations

## Checking for Health Issues and Corrective Actions with CHACTL QUERY DIAGNOSIS

```
$ chactl query diagnosis -db oltpacdb -start "2016-10-28 01:52:50" -end "2016-10-28 03:19:15"
```

```
2016-10-28 01:47:10.0 Database oltpacdb DB Control File IO Performance (oltpacdb_1) [detected]
2016-10-28 01:47:10.0 Database oltpacdb DB Control File IO Performance (oltpacdb_2) [detected]
2016-10-28 02:59:35.0 Database oltpacdb DB Log File Switch (oltpacdb_1) [detected]
2016-10-28 02:59:45.0 Database oltpacdb DB Log File Switch (oltpacdb_2) [detected]
```

### **Problem: DB Control File IO Performance**

Description: CHA has detected that reads or writes to the control files are slower than expected.

Cause: The Cluster Health Advisor (CHA) detected that reads or writes to the control files were slow because of an increase in disk IO.

The slow control file reads and writes may have an impact on checkpoint and Log Writer (LGWR) performance.

Action: Separate the control files from other database files and move them to faster disks or Solid State Devices.

### **Problem: DB Log File Switch**

Description: CHA detected that database sessions are waiting longer than expected for log switch completions.

Cause: The Cluster Health Advisor (CHA) detected high contention during log switches because the redo log files were small and the redo logs switched frequently.

Action: Increase the size of the redo logs.

# Cluster Health Advisor – Command Line Operations

## HTML Diagnostic Health Output Available (-html <file\_name>)

Timestamp	Target Information	Event Name	Detected/Cleared
2016-07-03 01:49:30.0	Host svr02	<a href="#">Host CPU Utilization</a>	detected
2016-07-03 01:49:50.0	Host svr01	<a href="#">Host CPU Utilization</a>	detected
2016-07-03 05:54:55.0	Host svr01	<a href="#">Host Memory Consumption</a>	detected
2016-07-04 03:40:00.0	Host svr02	<a href="#">Host CPU Utilization</a>	cleared
2016-07-04 03:40:05.0	Host svr01	<a href="#">Host CPU Utilization</a>	cleared
2016-07-04 03:40:05.0	Host svr01	<a href="#">Host Memory Consumption</a>	cleared

Problem	Description	Cause	Action
Host CPU Utilization	CHA detected larger than expected CPU utilization on this node. The available CPU resource may not be sufficient to support application failover or relocation of databases to this node.	The Cluster Health Advisor (CHA) detected an unexpected increase in CPU utilization by databases or applications on this node.	Identify CPU intensive processes and databases by reviewing Cluster Health Monitoring (CHM) data. Relocate databases to less busy machines, or limit the number of connections to databases on this node. Add nodes if more resources are required.
Host Memory Consumption	CHA detected that more memory than expected is consumed on this server. The memory is not allocated by sessions of this database.	The Cluster Health Advisor (CHA) detected an increase in memory consumption by other databases or by applications not connected to a database on this node.	Identify the top memory consumers by using the Cluster Health Monitor (CHM).

# Agenda

- 1 Introduction
- 2 CHA Architecture and Operation Details
- 3 Using CHA from the Command Line
- 4 Using CHA from EMCC for Alerts and Corrective Actions
- 5 Using the CHA GUI to Perform Root-Cause Analysis
- 6 Calibrating CHA to your RAC deployment
- 7 Q & A – Further Information

## Incident Manager

[Incident Manager](#) > Incident Details

Page Refreshed Sep 12, 2016 6:43:01 PM PDT

**CHA has detected a service degradation due to higher than expected I/O latencies.**

[Open in new tab](#)

Unassigned, Not acknowledged

**General** Events Notifications My Oracle Support Knowledge All Updates Related Events Related Metrics


### Incident Details

ID	5505
Metric Group	Cluster Health Advisor
Metric	CHA/DB Health
Target	prod (Cluster Database) <span></span>
Incident Created	Sep 13, 2016 1:33:57 AM GMT
Last Updated	Sep 13, 2016 1:33:57 AM GMT
Summary	CHA detected a for service degradation due to higher than expected I/O latencies.
Internal Event Name	CHA/DB Health: I/O problem
Event Type	Metric Alert
Category	Business

[Show internal values for attributes ...](#)

### Metric Data

Critical Threshold	600
Warning Threshold	300
Number of Occurrences	3
Last Known Value	1,178
Last Collection Timestamp	Sep 13, 2016 1:40:57 AM GMT



CHA/DB Health: I/O problem

### Tracking

[Acknowledge](#) [Add Comment ...](#) [Manage ...](#) **More**

Escalated	No	Owner	-
Priority	None	Acknowledged	No
Status	New		

**Last Comment** Incident created by rule (Name = Incident management rule set for all targets, Create incident for critical metric alerts [System generated rule]).: on Sep 13, 2016 1:33:57 AM GMT

☒ This incident will be automatically cleared when the underlying issue is resolved.

### Guided Resolution

#### Diagnostics

[Quality of Service Management](#)  
[View Performance Class Quality of Service](#)  
[Problem Analysis](#)  
[View Metric Help](#)

#### Actions

[Reevaluate Alert](#)  
[Edit Thresholds](#)

#### Corrective Actions

No corrective action defined.  
[Add corrective action](#)

Incident Manager

Incident Manager > Incident Details

Page Refreshed Sep 12, 2016 6:43:01 PM PDT

CHA has detected a service degradation due to higher than expected I/O latencies.

Open in new tab

Unassigned, Not acknowledged

- General
- Events
- Notifications
- My Oracle Support Knowledge
- All Updates
- Related Events
- Related Metrics


Incident Details

ID	5505
Metric Group	Cluster Health Advisor
Metric	CHA/DB Health
Target	prod (Cluster Database)
Incident Created	Sep 13, 2016 1:33:57 AM GMT
Last Updated	Sep 13, 2016 1:33:57 AM GMT
Summary	CHA detected a for service degradation due to higher than expected I/O latencies.
Internal Event Name	CHA/DB Health: I/O problem
Event Type	Metric Alert
Category	Business

Show internal values for attributes ...

Metric Data

Critical Threshold	600
Warning Threshold	300
Number of Occurrences	3
Last Known Value	1,178
Last Collection Timestamp	Sep 13, 2016 1:40:57 AM GMT



Tracking

Acknowledge Add Comment ... Manage ... More

Escalated	No	Owner	-
Priority	None	Acknowledged	No
Status	New		

Last Incident created by rule (Name = Incident management rule set for all targets, Create incident for critical metric Comment alerts [System generated rule]).: on Sep 13, 2016 1:33:57 AM GMT


This incident will be automatically cleared when the underlying issue is resolved.

Guided Resolution

Diagnostics	Actions	Corrective Actions
Quality of Service Management	Reevaluate Alert	No corrective action defined.
View Performance Class Quality of Service	Edit Thresholds	Add corrective action
Problem Analysis		
View Metric Help		

Cluster Health Advisor

Problem	The degradation is caused by a higher than expected utilization of shared storage devices for this database. No evidence of significant increase in I/O demand on the local node.
Confidence	95.17%
Action	Validate whether there is increase in I/O demand on other nodes than the local and find I/O intensive SQL . Add more disks to disk group or move database to faster disks.






# Using EMCC for Alerts and Corrective Actions

## Cluster Health Advisor

**Problem** The degradation is caused by a higher than expected utilization of shared storage devices for this database. No evidence of significant increase in I/O demand on the local node.

**Confidence** 95.17%

**Action** Validate whether there is increase in I/O demand on other nodes than the local and find I/O intensive SQL . Add more disks to disk group or move database to faster disks.



# Agenda

- 1 Introduction
- 2 CHA Architecture and Operation Details
- 3 Using CHA from the Command Line
- 4 Using CHA from EMCC for Alerts and Corrective Actions
- 5 Using the CHA GUI to Perform Root-Cause Analysis
- 6 Calibrating CHA to your RAC deployment
- 7 Q & A – Further Information

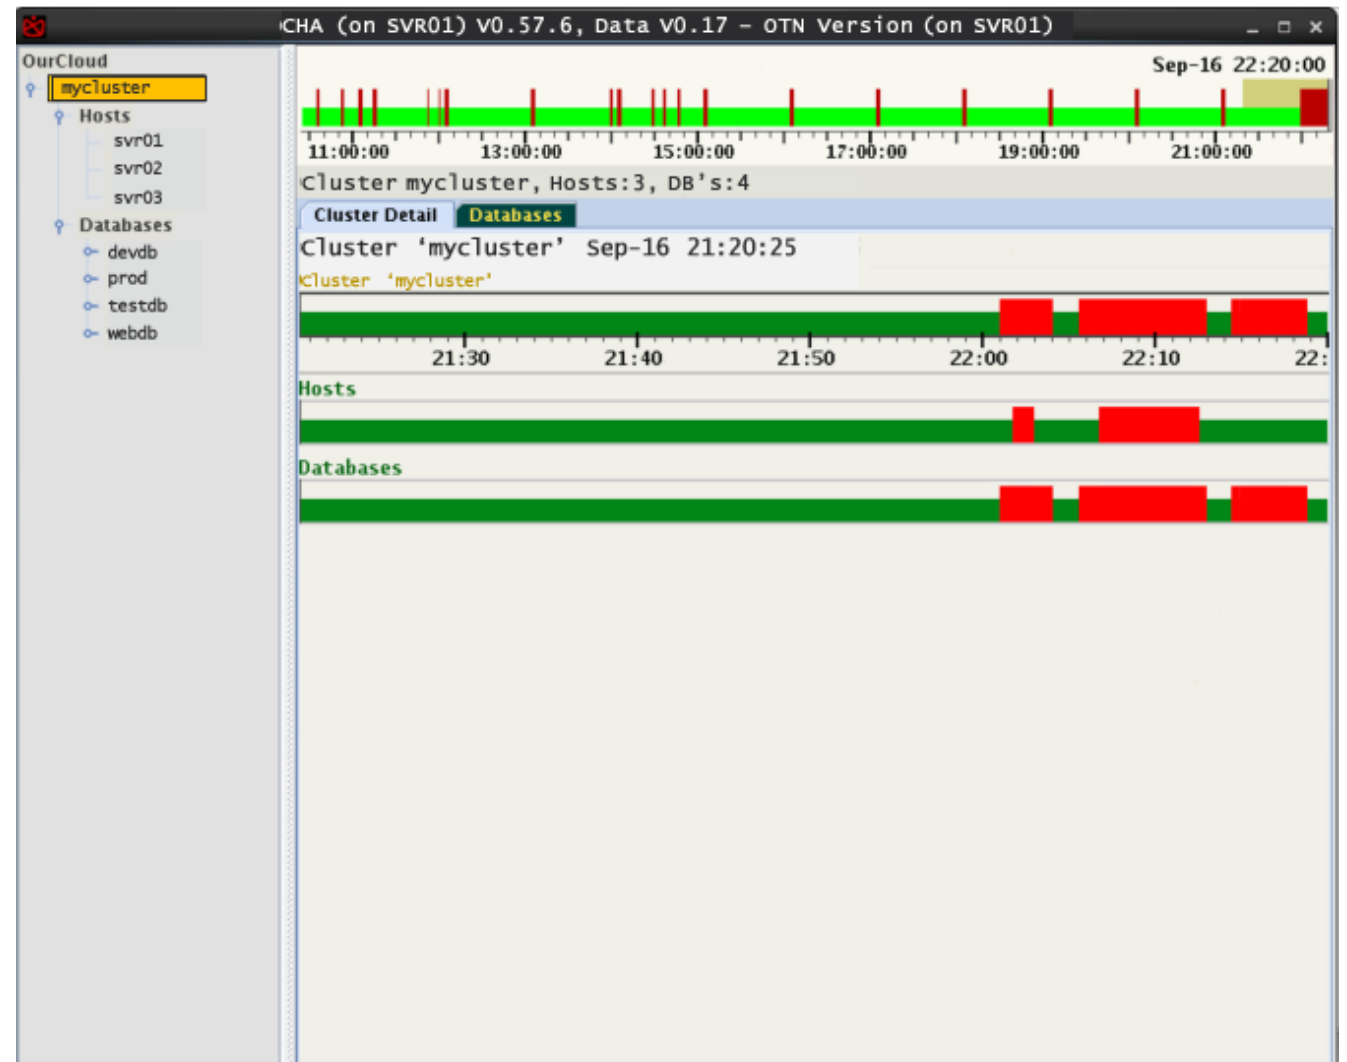
# Using the CHA GUI to Perform Root-Cause Analysis

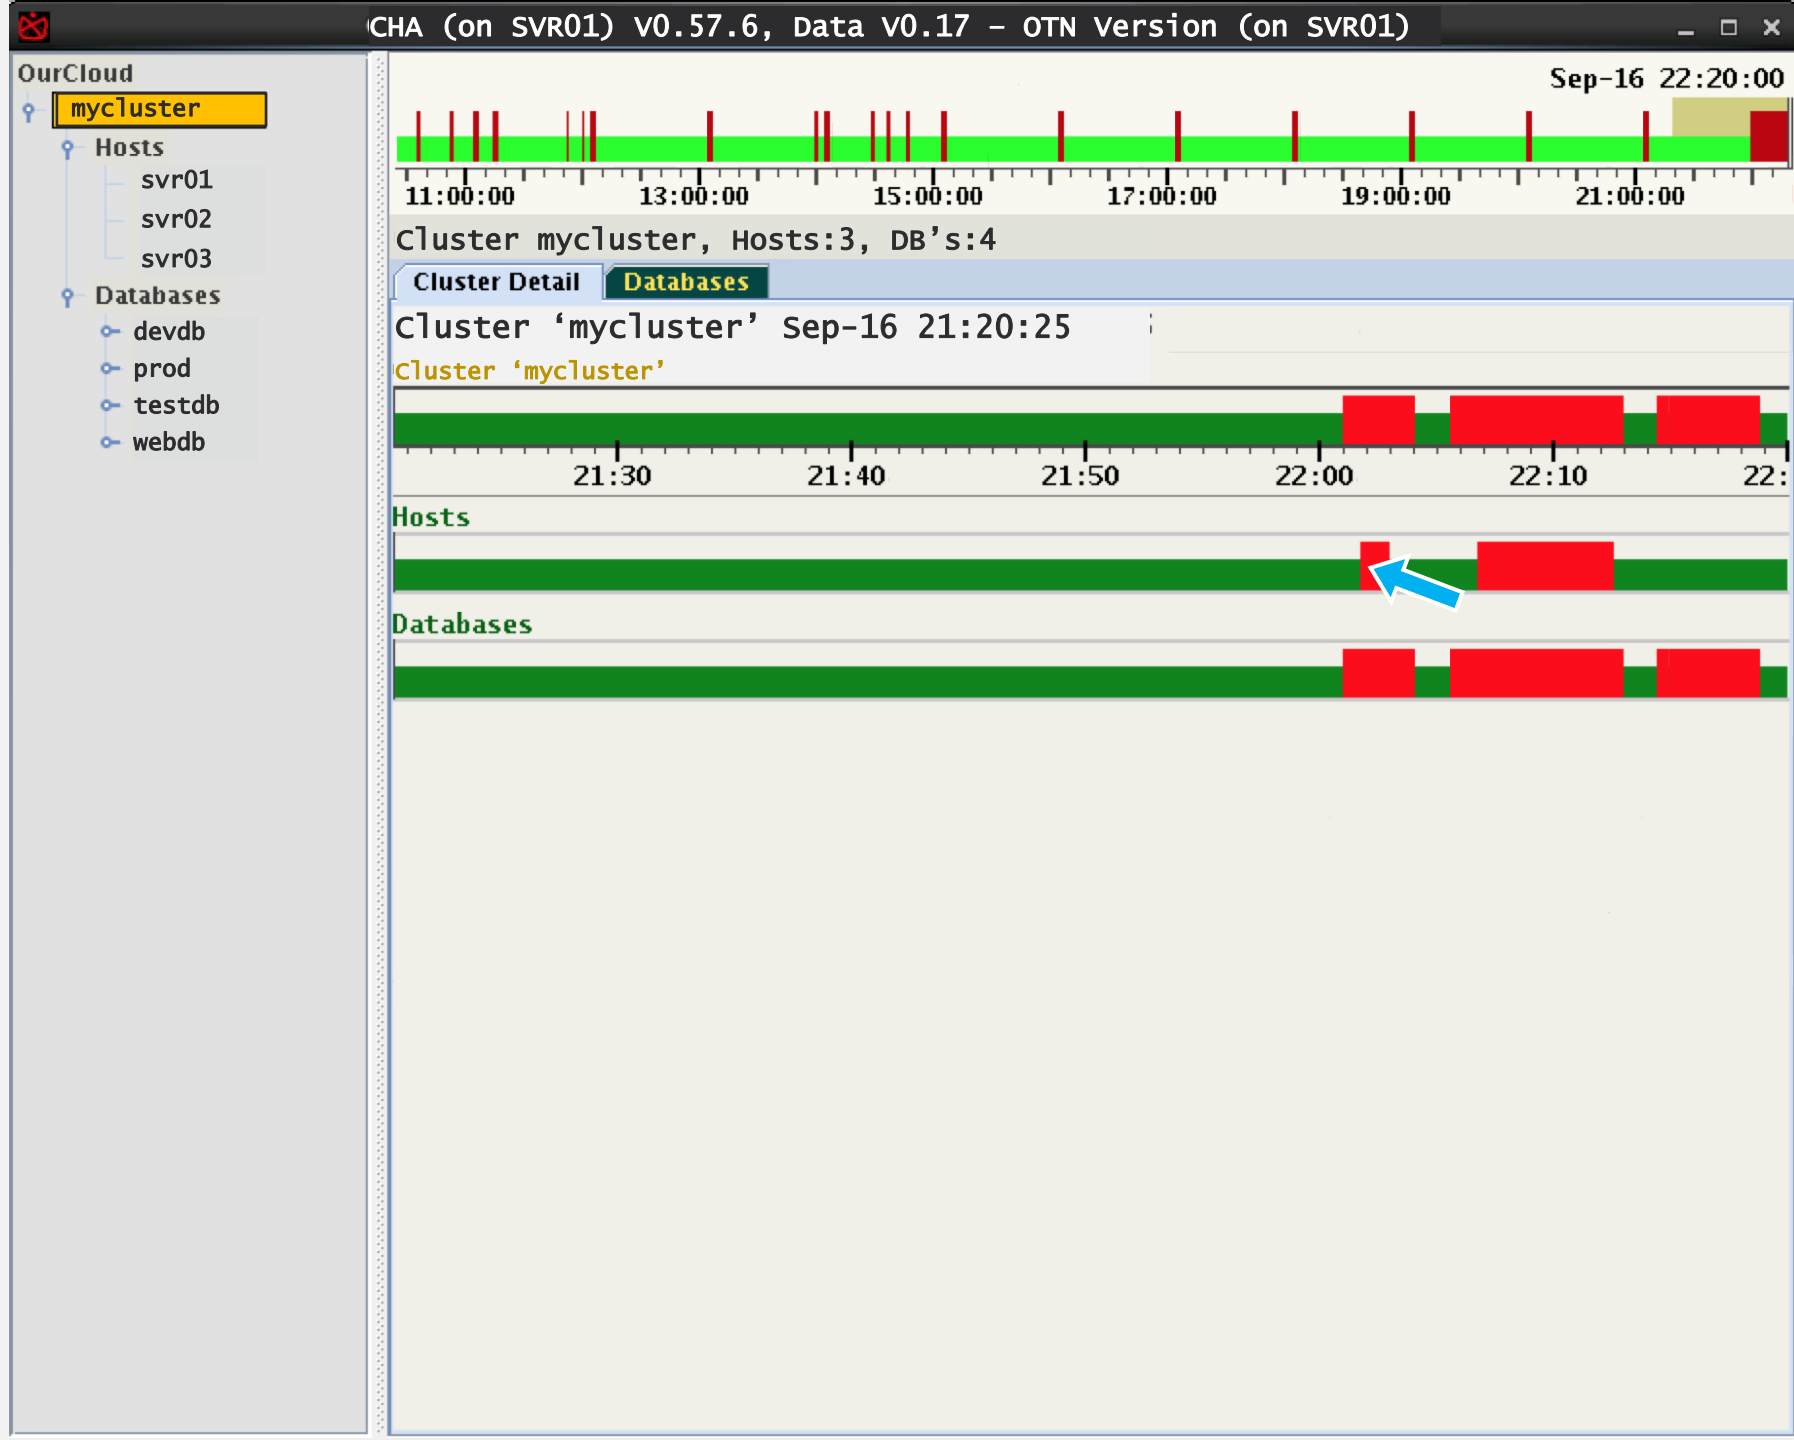
## Overview

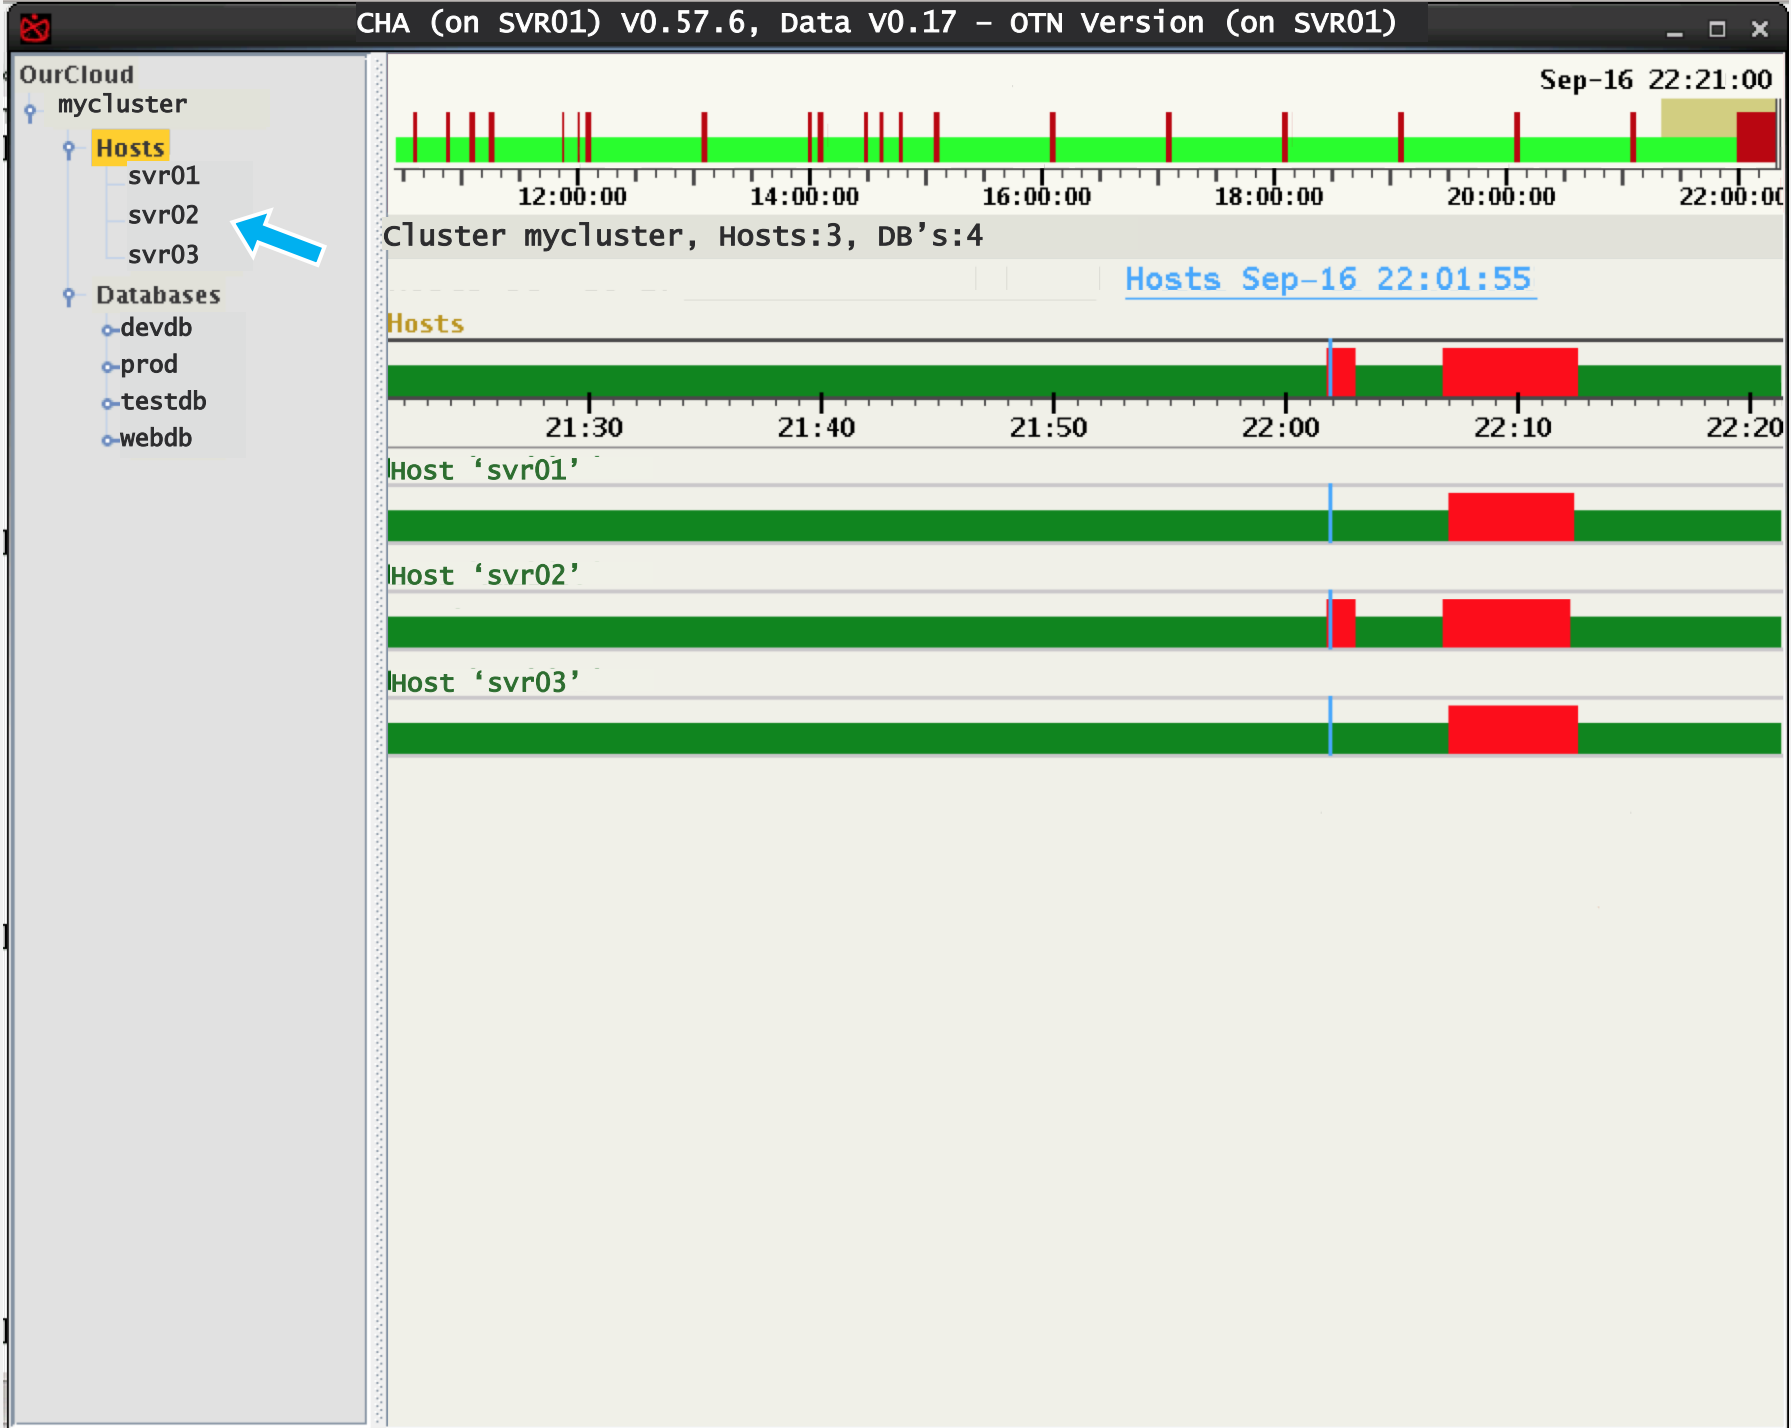
- Standalone Java GUI Client
- Must be run on local cluster node
- Can be run against live GIMR or MDB (dump) file

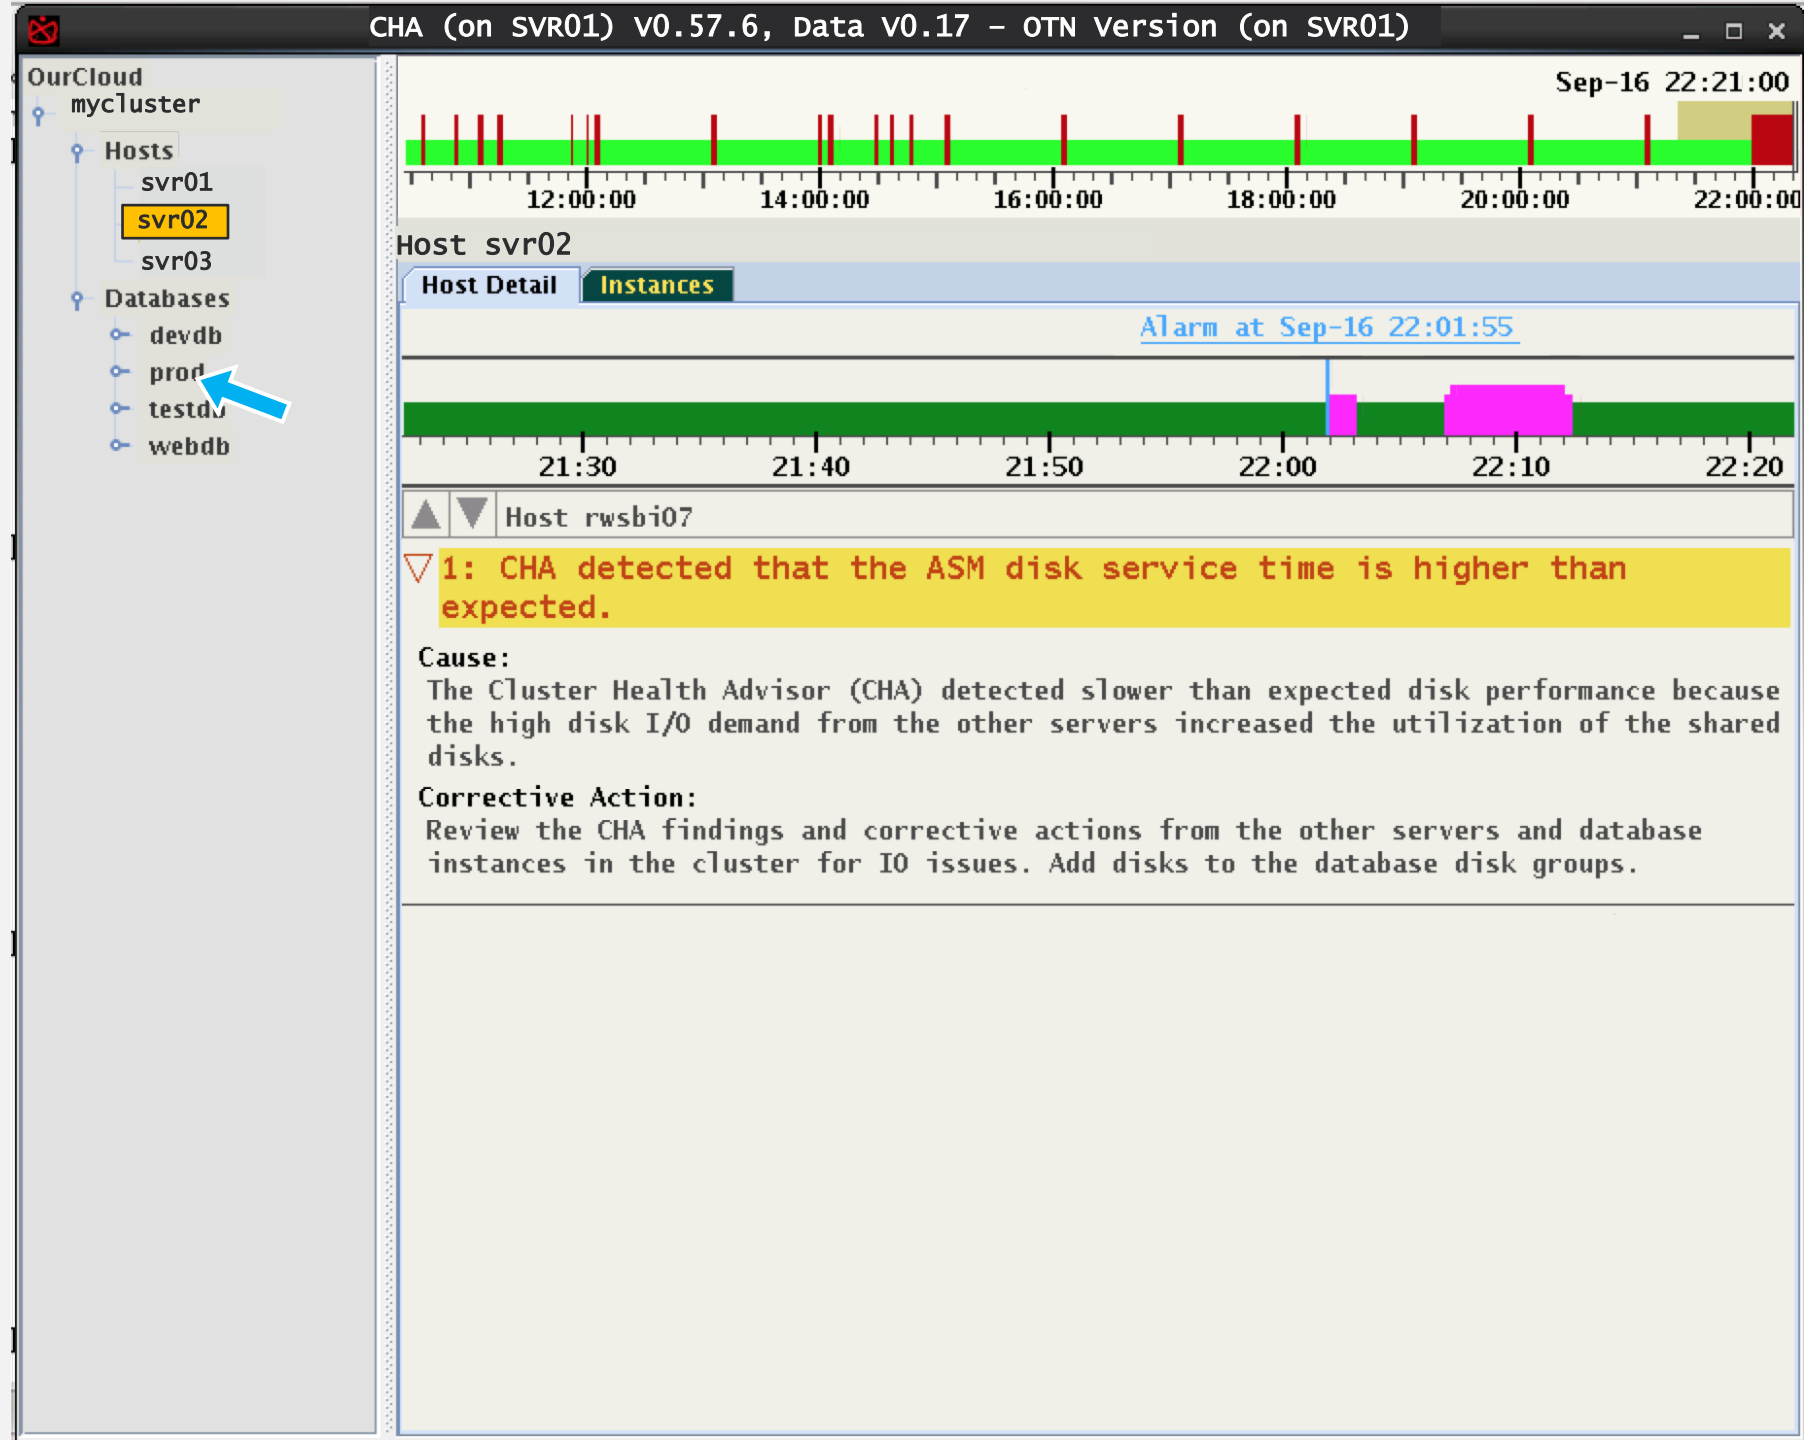
```
chactl export repository -format  
mdb -start '2017-05-01 00:00:00'  
-end '2017-05-10 00:00:00'
```

- Used internally for development
- Will be available and maintained on Oracle Technology Network










- OurCloud
  - mycluster
    - Hosts
      - svr01
      - svr02
      - svr03
    - Databases
      - devdb
      - prod
        - prod\_1
        - prod\_2
      - testdb
      - webdb
        - webdb\_1
        - webdb\_2




Database prod, Instances: 2

Database

DB prod in mycluster Sep-16 22:16:35

DB prod in mycluster




Instance 'prod\_1'



Instance 'prod\_2'

- OurCloud
- mycluster
    - Hosts
      - svr01
      - svr02
      - svr03
    - Databases
      - devdb
      - prod
        - prod\_1
        - prod\_2
      - testdb
      - webdb


Sep-16 22:23:00



Instance prod\_1

Instance Detail Host

[Alarm at Sep-16 22:01:20](#)

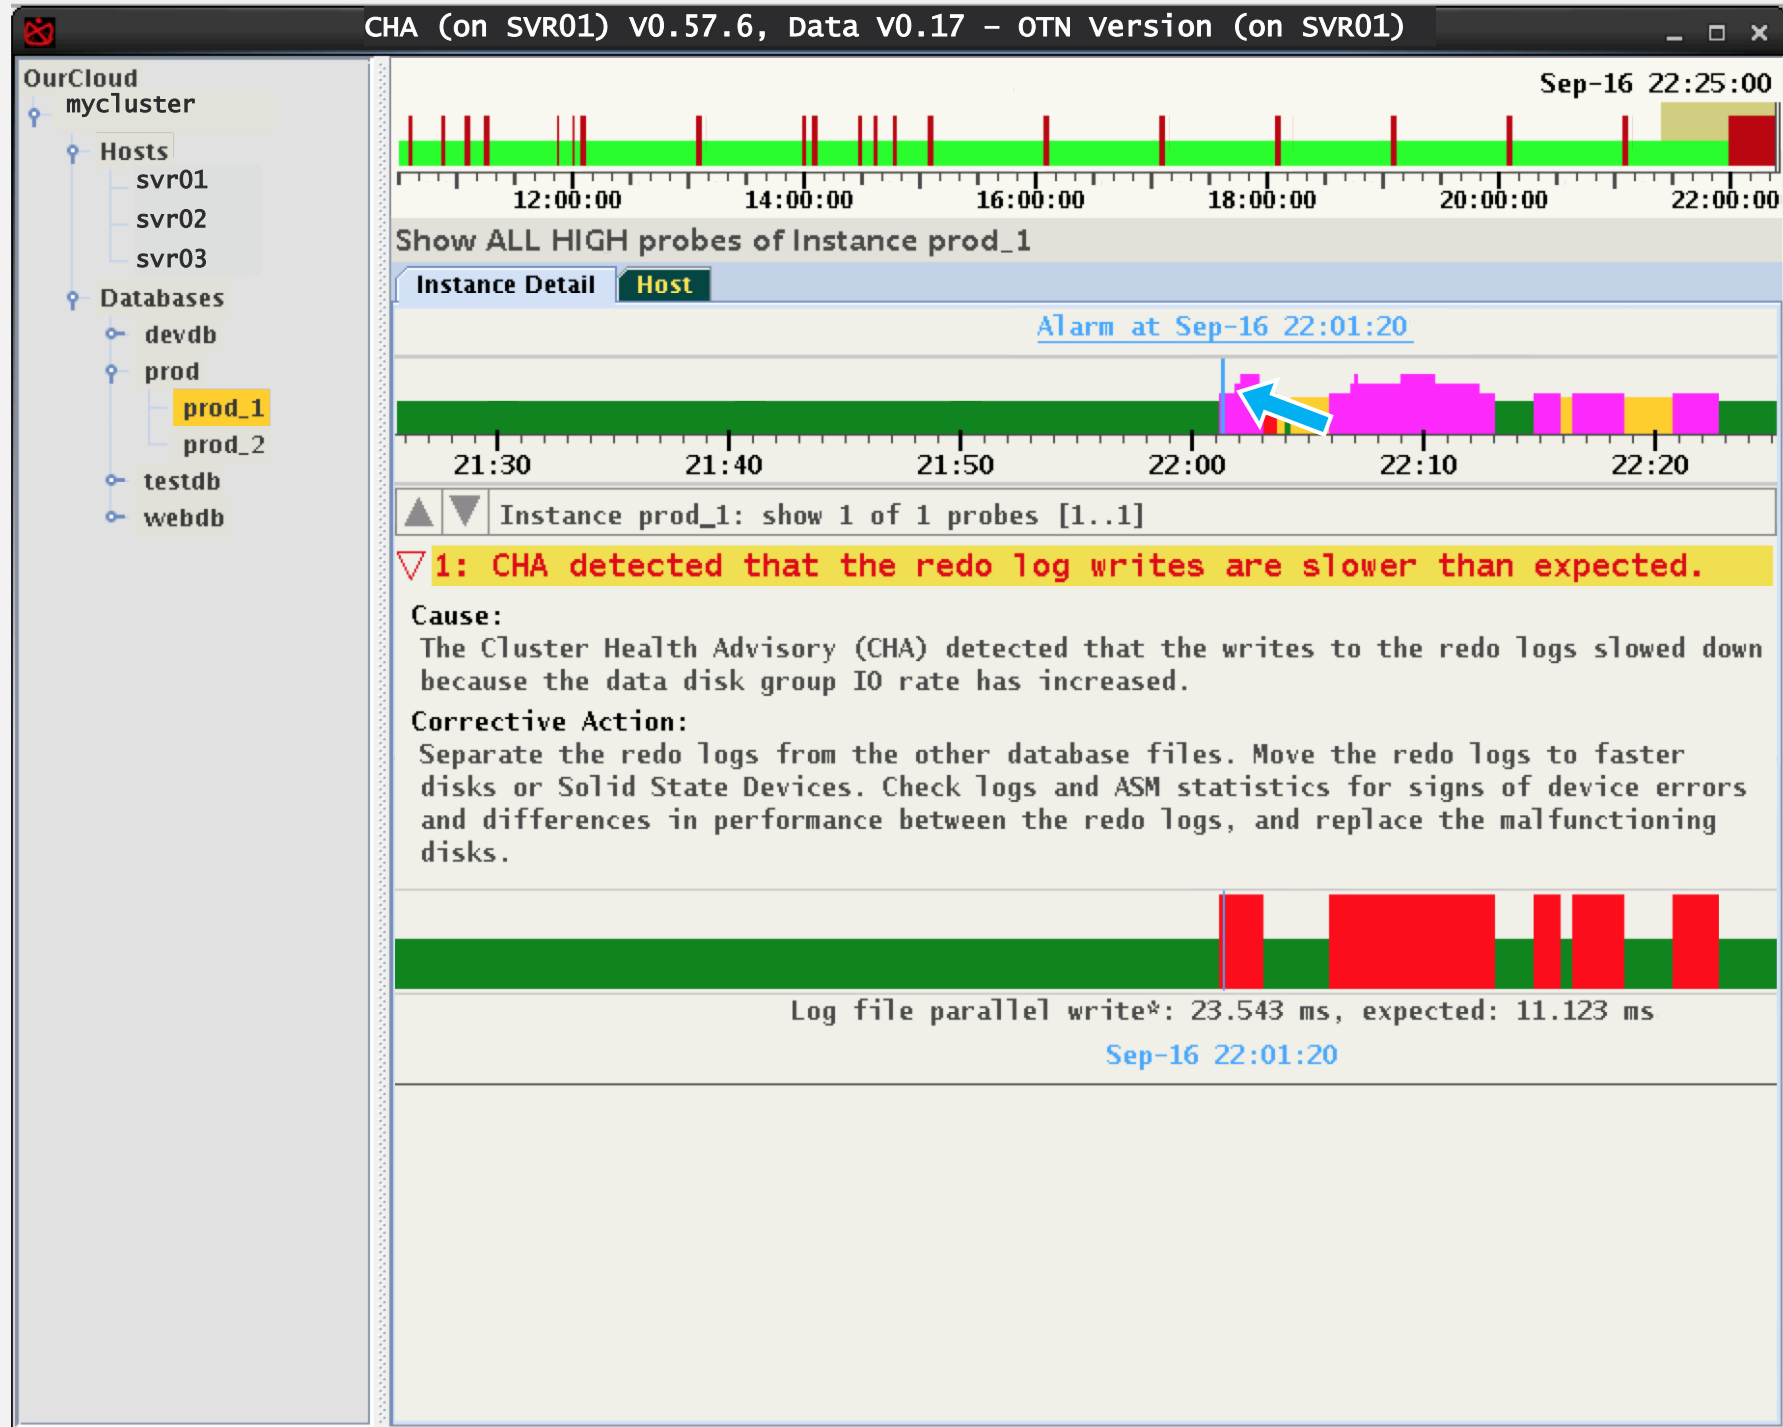


▲ ▼ Instance prod\_1

1: CHA detected that the redo log writes are slower than expected.







OurCloud

mycluster

Hosts

svr01

svr02

svr03

Databases

devdb

prod


prod\_1

prod\_2

testdb

webdb

Sep-16 22:22:00




Instance prod\_1

Instance Detail

Host

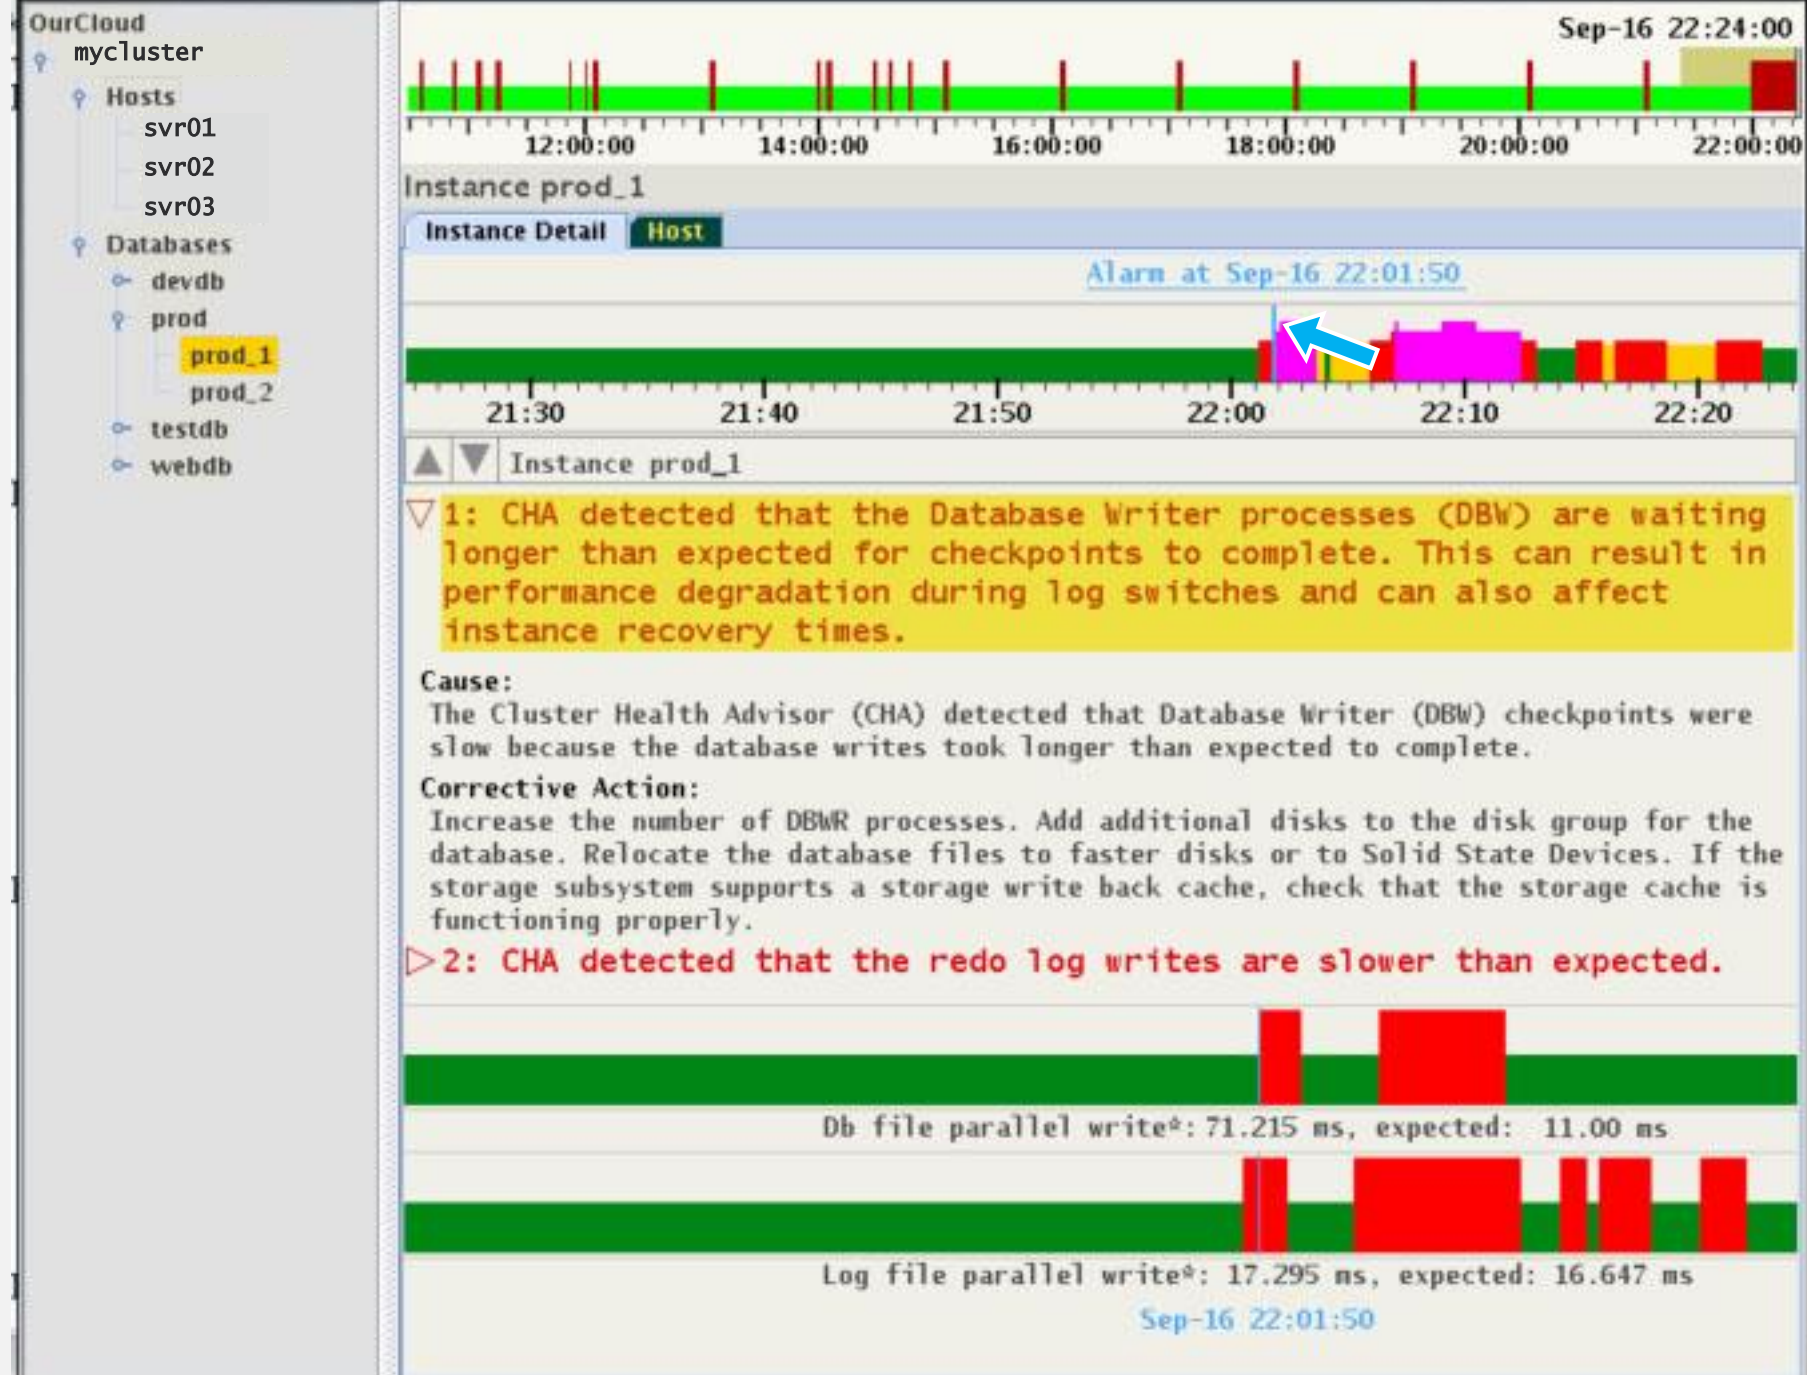
Alarm at Sep-16 22:01:55



▲ ▼ Instance prod\_1

1: CHA detected that the Database Writer processes (DBW) are waiting longer than expected for checkpoints to complete. This can result in performance degradation during log switches and can also affect instance recovery times.

2: CHA detected that the redo log writes are slower than expected.



OurCloud

mycluster

Hosts

svr01

svr02

svr03

Databases

devdb

prod


prod\_1

prod\_2

testdb

webdb


Sep-16 22:25:00



Instance prod\_1

Instance Detail

Host

[Alarm at Sep-16 22:02:05](#)

Instance prod\_1

1: CHA detected that the Database Writer processes (DBW) are waiting longer than expected for checkpoints to complete. This can result in performance degradation during log switches and can also affect instance recovery times.

**Cause:**

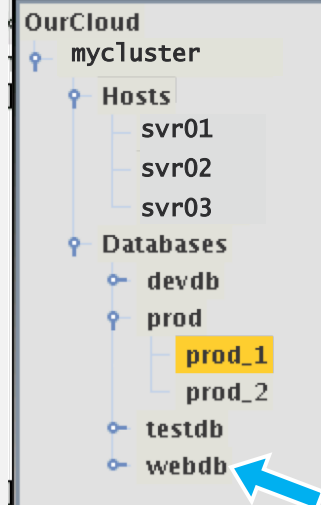
The Cluster Health Advisor (CHA) detected that Database Writer (DBW) checkpoints were slow because the database writes took longer than expected to complete.

**Corrective Action:**


Increase the number of DBWR processes. Add additional disks to the disk group for the database. Relocate the database files to faster disks or to Solid State Devices. If the storage subsystem supports a storage write back cache, check that the storage cache is functioning properly.

2: CHA detected that the redo log writes are slower than expected.

3: CHA detected that the ASM disk service time is higher than expected.



Sep-16 22:26:00




Show ALL HIGH probes of Instance prod\_1

Instance Detail

Host

Alarm at Sep-16 22:02:05



▲ ▼ Instance prod\_1: show 1 of 4 probes [1..1]


- ▷ 1: CHA detected that the Database Writer processes (DBW) are waiting longer than expected for checkpoints to complete. This can result in performance degradation during log switches and can also affect instance recovery times.
- 2: CHA detected that the redo log writes are slower than expected.
- ▽ 3: CHA detected that the ASM disk service time is higher than expected.

**Cause:**

The Cluster Health Advisor (CHA) detected slower than expected disk performance because the high disk I/O demand from the other servers increased the utilization of the shared disks.

**Corrective Action:**

Review the CHA findings and corrective actions from the other servers and database instances in the cluster for IO issues. Add disks to the database disk groups.

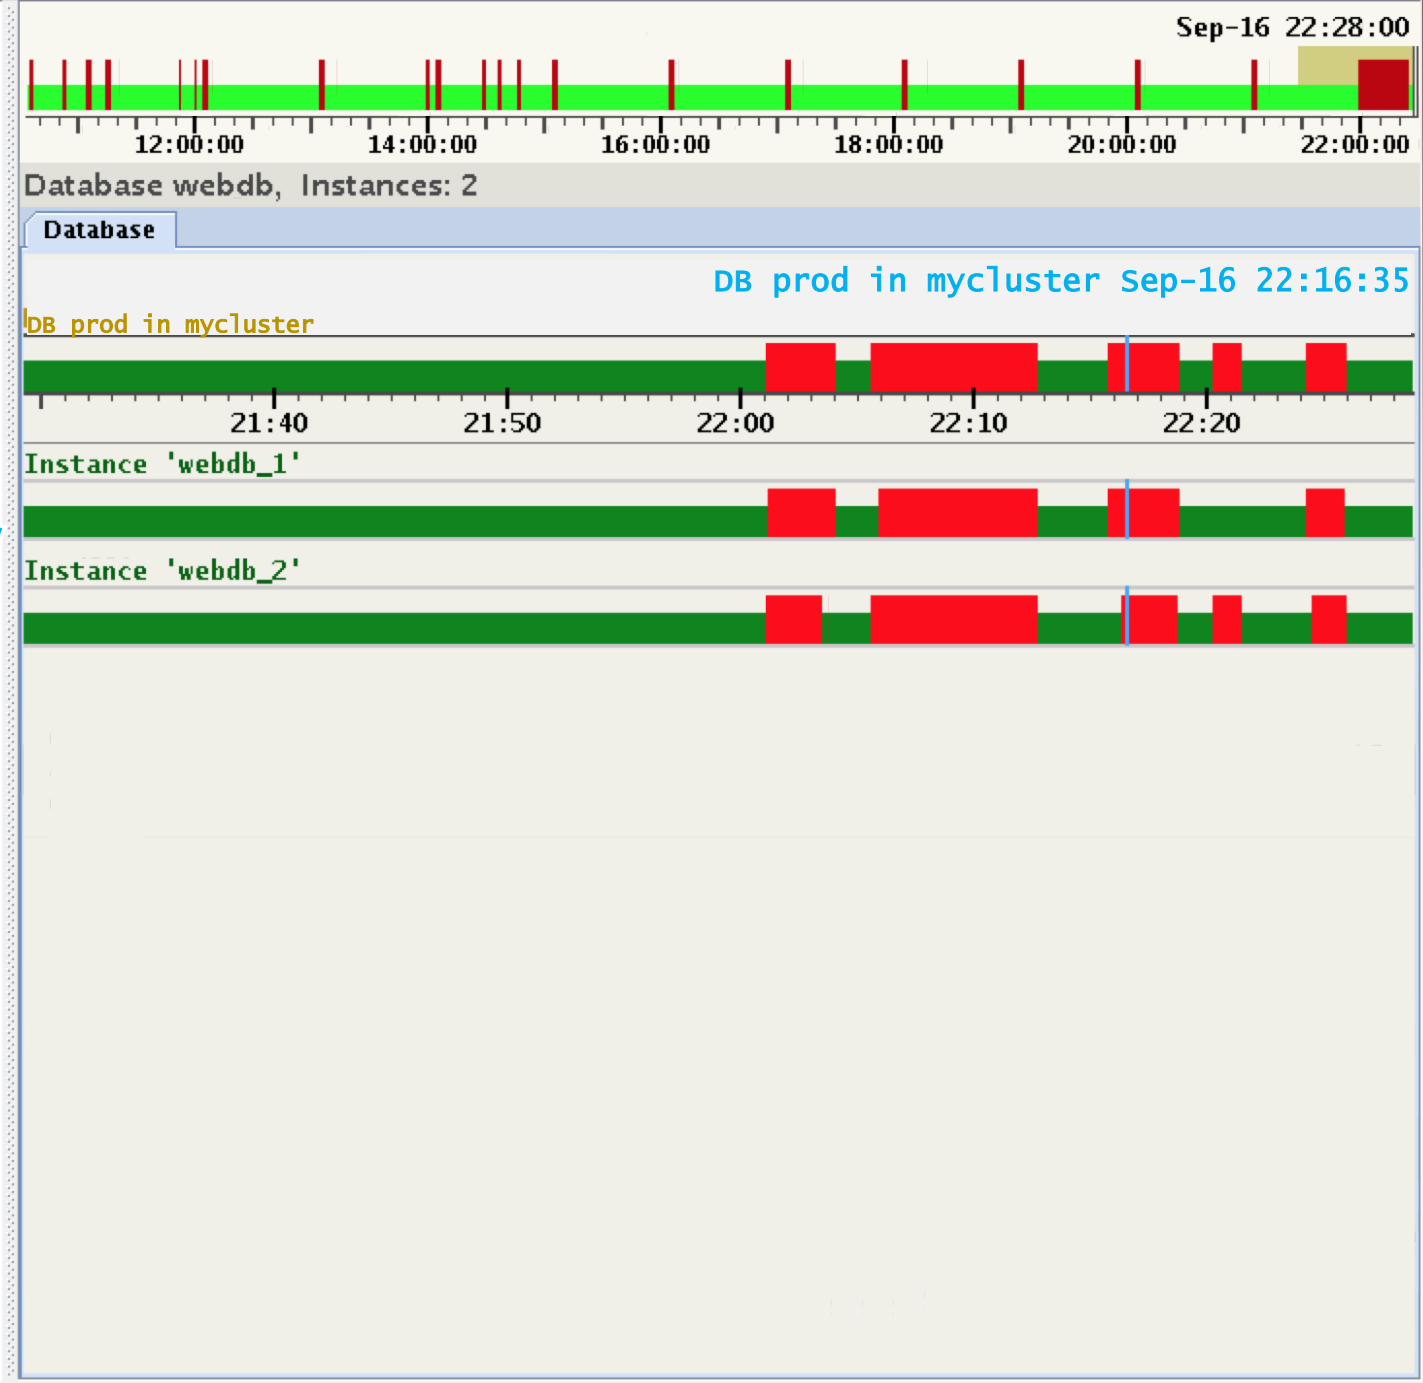


Disk service time (ASM)\*: 11.00 ms/I/O, expected: 8.04 ms/I/O

Sep-16 22:02:05


OurCloud

- mycluster
  - Hosts
    - svr01
    - svr02
    - svr03
  - Databases
    - devdb
    - prod
      - prod\_1
      - prod\_2
    - testdb
    - webdb
      - webdb\_1
      - webdb\_2



- OurCloud
- mycluster
    - Hosts
      - svr01
      - svr02
      - svr03
    - Databases
      - devdb
      - prod
        - prod\_1
        - prod\_2
      - testdb
      - webdb
        - webdb\_1**
        - webdb\_2


Sep-16 22:27:00



Show ALL HIGH probes of Instance webdb\_1

Instance Detail Host

Alarm at Sep-16 22:01:30



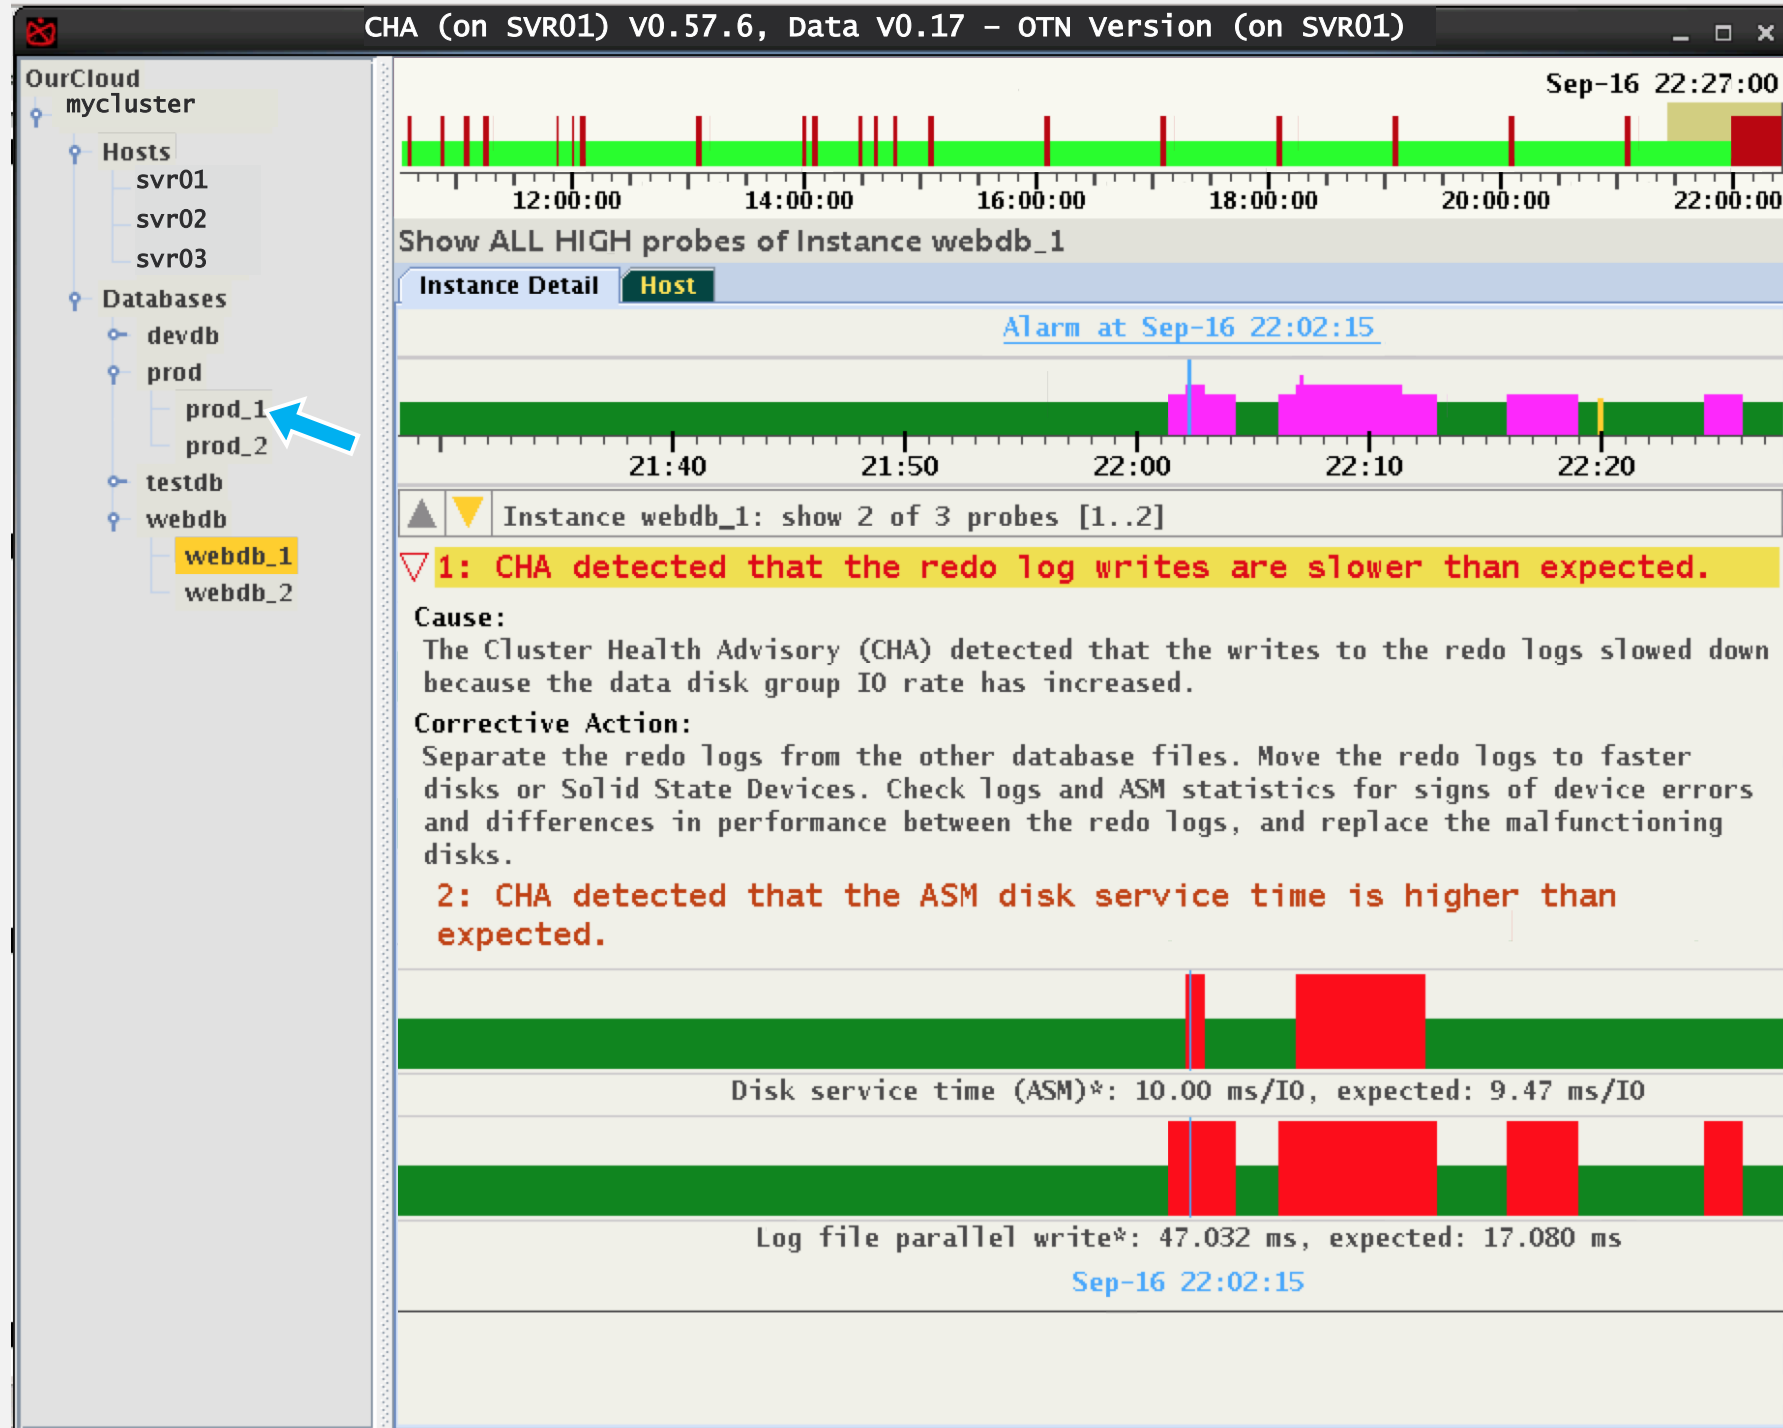
▲ ▼ Instance webdb\_1: show 1 of 1 probes [1..1]

1 CHA detected that the redo log writes are slower than expected.



Log file parallel write\*: 8.769 ms, expected: 7.668 ms

Sep-16 22:01:30





- OurCloud
- rwbsi0508-mb1
    - Hosts
    - Databases
      - prod
        - prod1
        - prod2

May-11 10:57:00


Instance prod2

Probe Selection


Instance Detail

Host

High probes at May-11 10:46:35




Instance prod2



Consistent gets\*: 983599.81 gets/s, expected: 694050.56 gets/s

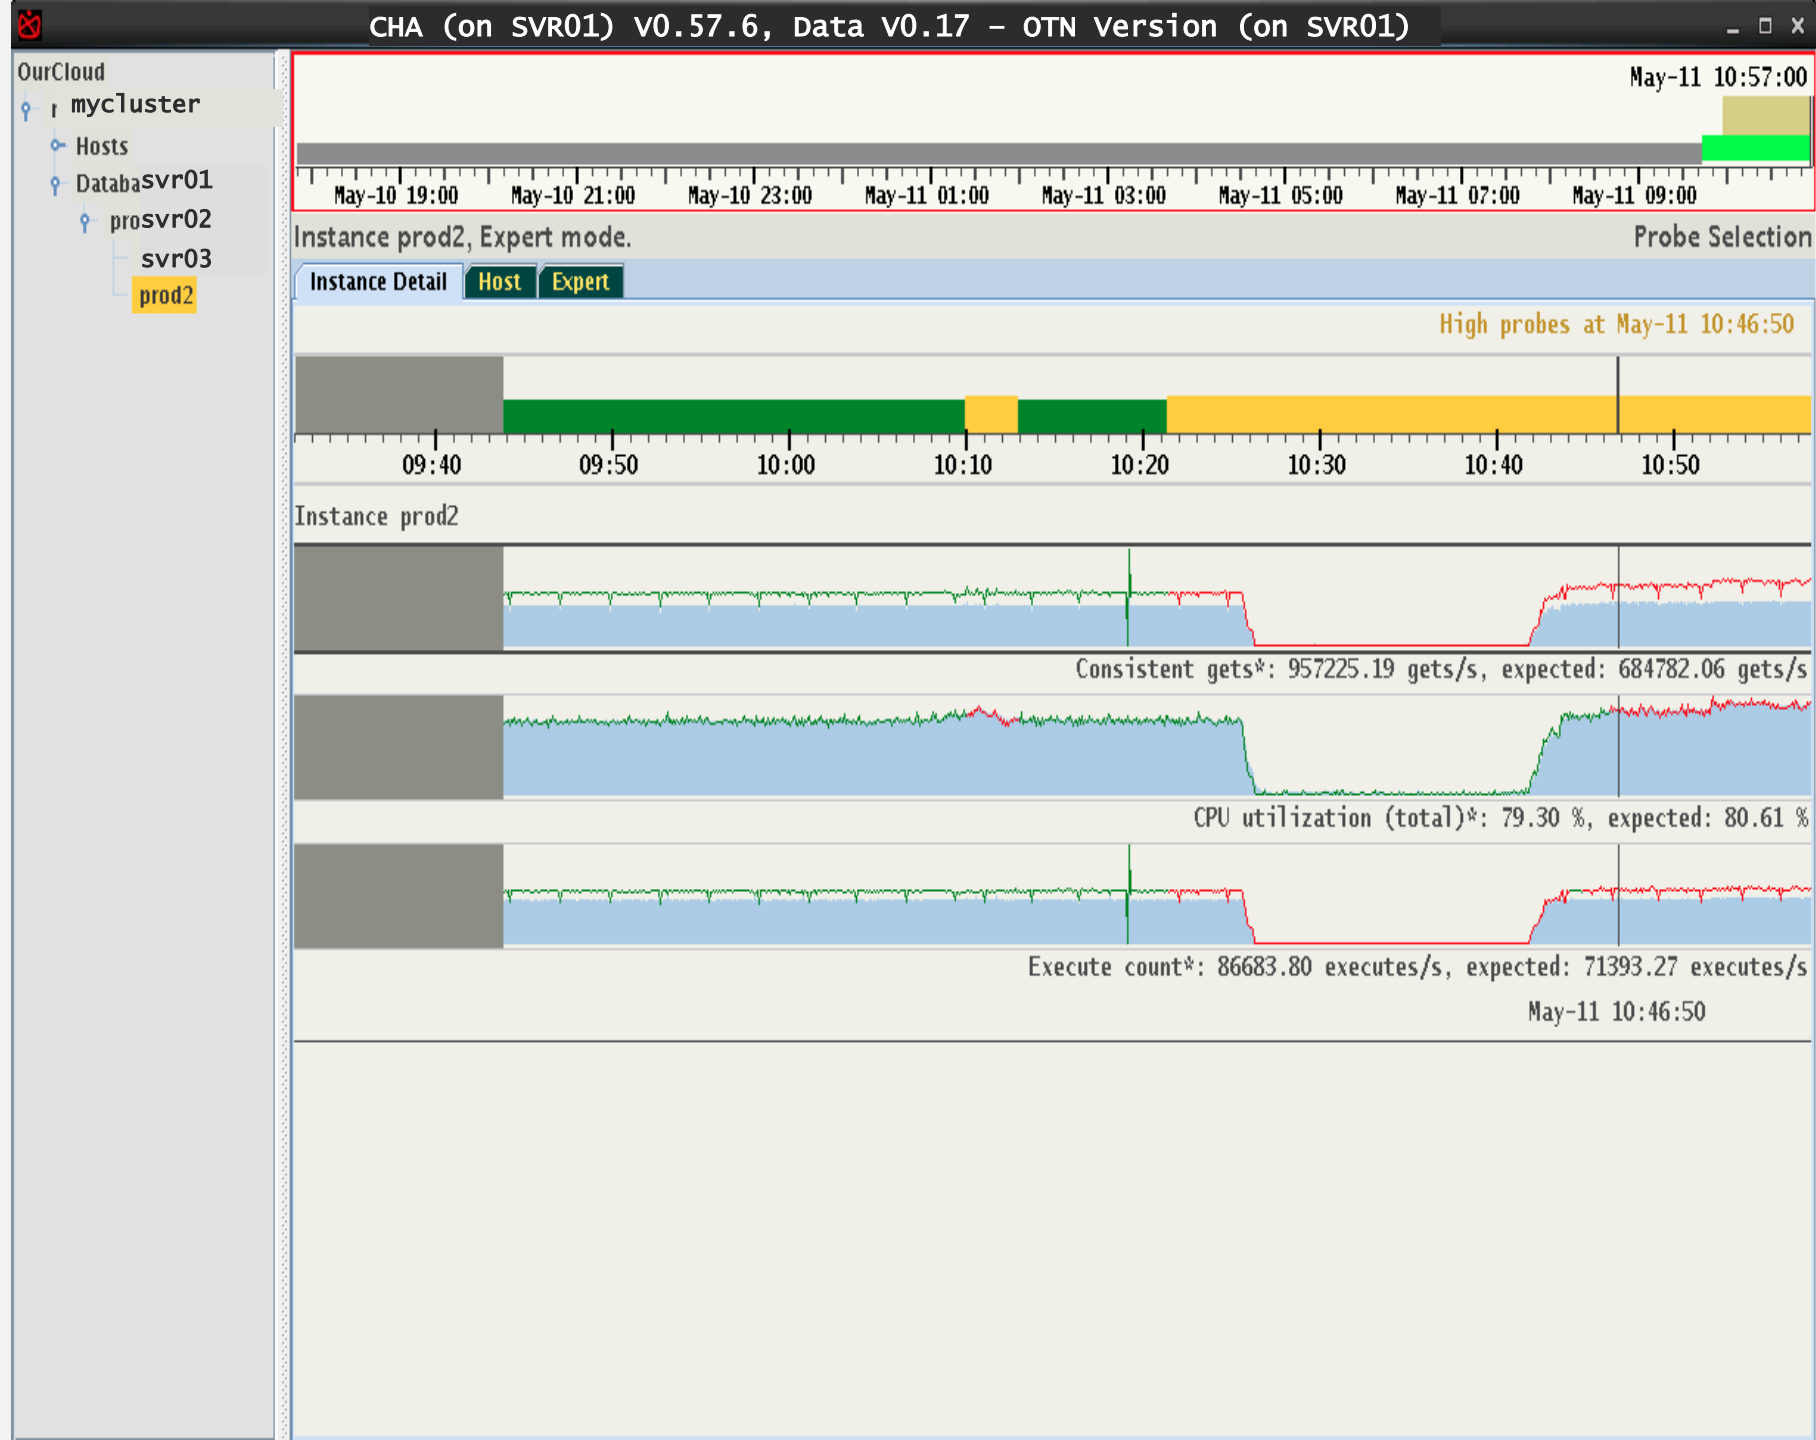


CPU utilization (total)\*: 82.90 %, expected: 82.51 %



Execute count\*: 88759.40 executes/s, expected: 71916.57 executes/s

May-11 10:46:35



“We also deployed the CHA GUI which has been designed specifically to visualize abnormalities in OS and RDBMS statistics data streams which are detected and diagnosed by CHA in real time. It stands out as a successful visualization of the real time and historical monitoring data and diagnostics in our Oracle RAC database environment.”

— Jens-Christian Pokolm, Head of Databases  
Postbank Systems AG

# Agenda

- 1 Introduction
- 2 CHA Architecture and Operation Details
- 3 Using CHA from the Command Line
- 4 Using CHA from EMCC for Alerts and Corrective Actions
- 5 Using the CHA GUI to Perform Root-Cause Analysis
- 6 Calibrating CHA to your RAC Deployment
- 7 Q & A – Further Information

# Calibrating CHA to your RAC Deployment

## Overview

- Calibration Goal: Increase sensitivity and accuracy with sufficient warning
- Release ships with conservative models to minimize false warnings
  - DEFAULT\_CLUSTER for each cluster node
  - DEFAULT\_DB for each database instance
- Use your own data for periods of “normal operations” to increase sensitivity
  - Recommended minimum 6 hour period
  - Should include all normal workload phases for that model
- Models may be changed dynamically online using CHACTL

# Calibrating CHA to your RAC deployment

## Choosing a Data Set for Calibration – Defining “normal”

```
$ chactl query calibration -cluster -timeranges 'start=2016-10-28 07:00:00,end=2016-10-28 13:00:00'
```

Cluster name : mycluster

Start time : 2016-10-28 07:00:00

End time : 2016-10-28 13:00:00

Total Samples : 11524

Percentage of filtered data : 100%

### 1) Disk read (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
0.11	0.00	2.62	0.00	114.66

<25	<50	<75	<100	>=100
99.87%	0.08%	0.00%	0.02%	0.03%

### 2) Disk write (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
0.01	0.00	0.15	0.00	6.77

<50	<100	<150	<200	>=200
100.00%	0.00%	0.00%	0.00%	0.00%

### 3) Disk throughput (ASM) (IO/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
2.20	0.00	31.17	0.00	1100.00

<5000	<10000	<15000	<20000	>=20000
100.00%	0.00%	0.00%	0.00%	0.00%

### 4) CPU utilization (total) (%)

MEAN	MEDIAN	STDDEV	MIN	MAX
9.62	9.30	7.95	1.80	77.90

<20	<40	<60	<80	>=80
92.67%	6.17%	1.11%	0.05%	0.00%

# Calibrating CHA to your RAC deployment

## Creating a new CHA Model with CHACTL

- Create and store the new model

```
$ chactl query calibrate cluster -model daytime -timeranges 'start=2016-10-28 07:00:00,  
end=2016-10-28 13:00:00'
```

- Begin using the new model

```
$ chactl monitor cluster -model daytime
```

- Confirm the new model is being used

```
$ chactl status -verbose
```

```
monitoring nodes svr01, svr02 using model daytime  
monitoring database qoltpacdb, instances oltpacdb_1, oltpacdb_2 using model DEFAULT_DB
```

# Agenda

- 1 Introduction
- 2 CHA Architecture and Operation Details
- 3 Using CHA from the Command Line
- 4 Using CHA from EMCC for Alerts and Corrective Actions
- 5 Using the CHA GUI to Perform Root-Cause Analysis
- 6 Calibrating CHA to your RAC deployment
- 7 Q & A – Further Information



## For Further Information

- [Oracle 12c Autonomous Health Framework User's Guide](#)
- [Oracle 12c Clusterware Administration and Deployment Guide](#)
- [Oracle Autonomous Health Framework on OTN](#)

