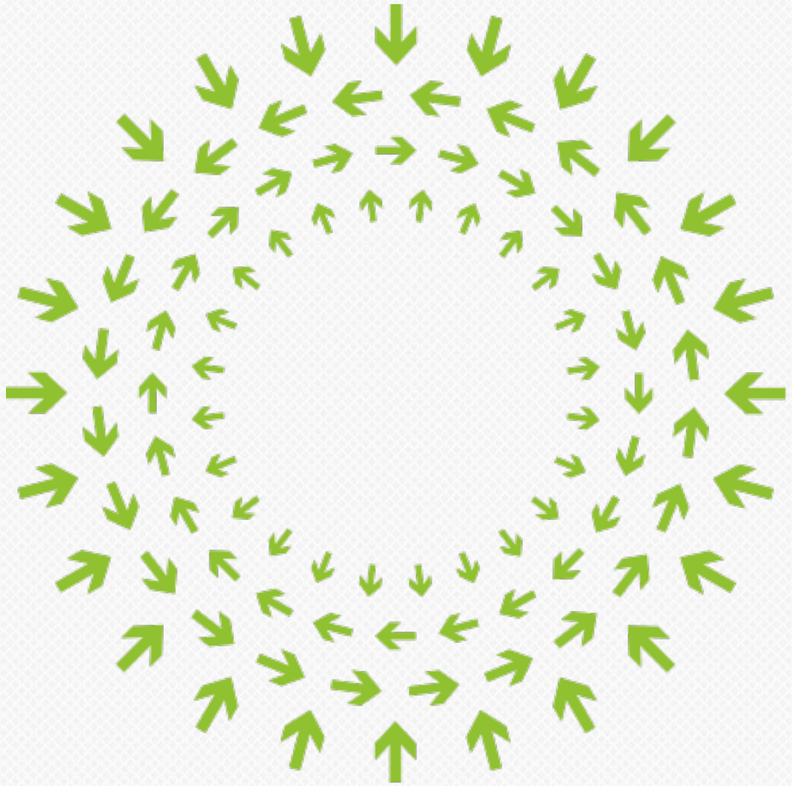

OPEN Compute Summit

March 10–11, 2015

San Jose

OPEN Compute Summit

Health monitoring & predictive analytics

To lower the TCO in a datacenter

Christian B. Madsen & Andrei Khurshudov
Engineering Manager & Sr. Director
Seagate Technology
christian.b.madsen@seagate.com

Outline

1. The opportunity
2. Our vision and implementation
3. Use cases
4. Summary

The opportunity

What if…

Seagate offered you a technology that could help you

Improve datacenter efficiency

Optimize system management

Reduce potential cost of operation

The problem

Failures in storage lead to costly outages

- **1 billion hard drives** will be used in cloud datacenters by 2020, highlighting the need to manage drive health at scale
- One total outage per datacenter is statistically expected every year
- 80% of those outages are not completely explained (or linked to root causes)
- $700,000 is the average cost per incident
  - $8,000 is the average cost per minute of an unplanned outage
- **Up to 10% of datacenter accidents are related to storage**

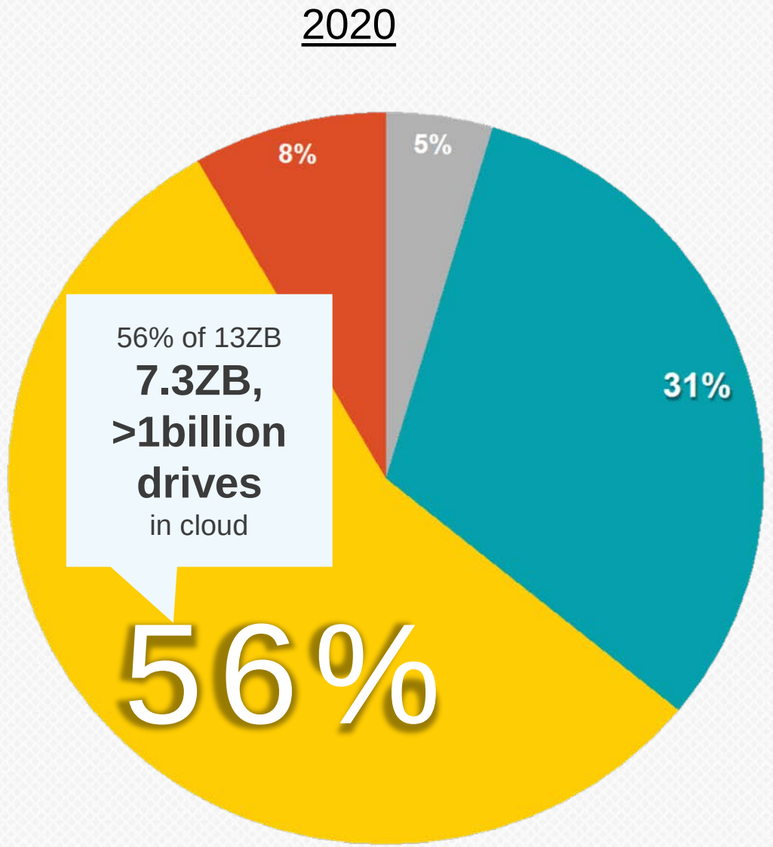

Source: Seagate Strategic Marketing and Research 2013

Better drive management will lower the TCO

Top 4 challenges in drive management

1. Drive health monitoring
   - Need reliable key performance indicators to track drive health status

2. Drive failure prediction
   - “Ultimately, we want to know when our drives will fail so we can take actions before that happens”

3. Drive failure diagnostics and management automation
   - Need to correctly identify and quickly resolve issues
   - Need to prevent false alerts to reduce cost of failure handling

4. Drive lifespan extension
   - Need to know how to optimize operating environment for better reliability
   - Need to reuse partially good drives (should be possible with in-drive diagnostic, IDD)

Our vision and implementation

Our vision

Monitoring, analytics, prediction and control – “The internet of things”[2]

Data Aggregation

MONITORING

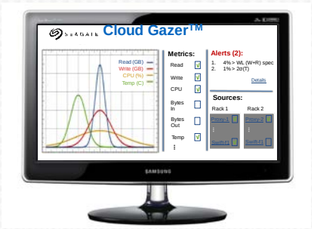

Global Access

ANALYTICS

Quick Issue Resolution

Data Center

CONTROL

Actionable Decisions

- Report storage health
- Run drive self-test
- Shut-down systems
- Repair drives
- Run auto-FA
- Point at an issue
- Highlight inefficiency
- Predict reliability
- Detect anomalies

PREDICTIONS

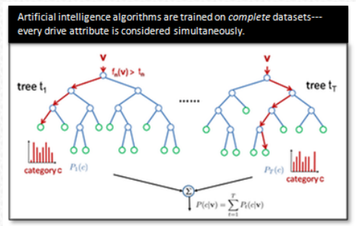

Early Warning System

- Drive-centric health monitoring
- Analytics and predictive models
- Closed-loop automation

Functional diagram

Monitoring, intelligent decisions and automation

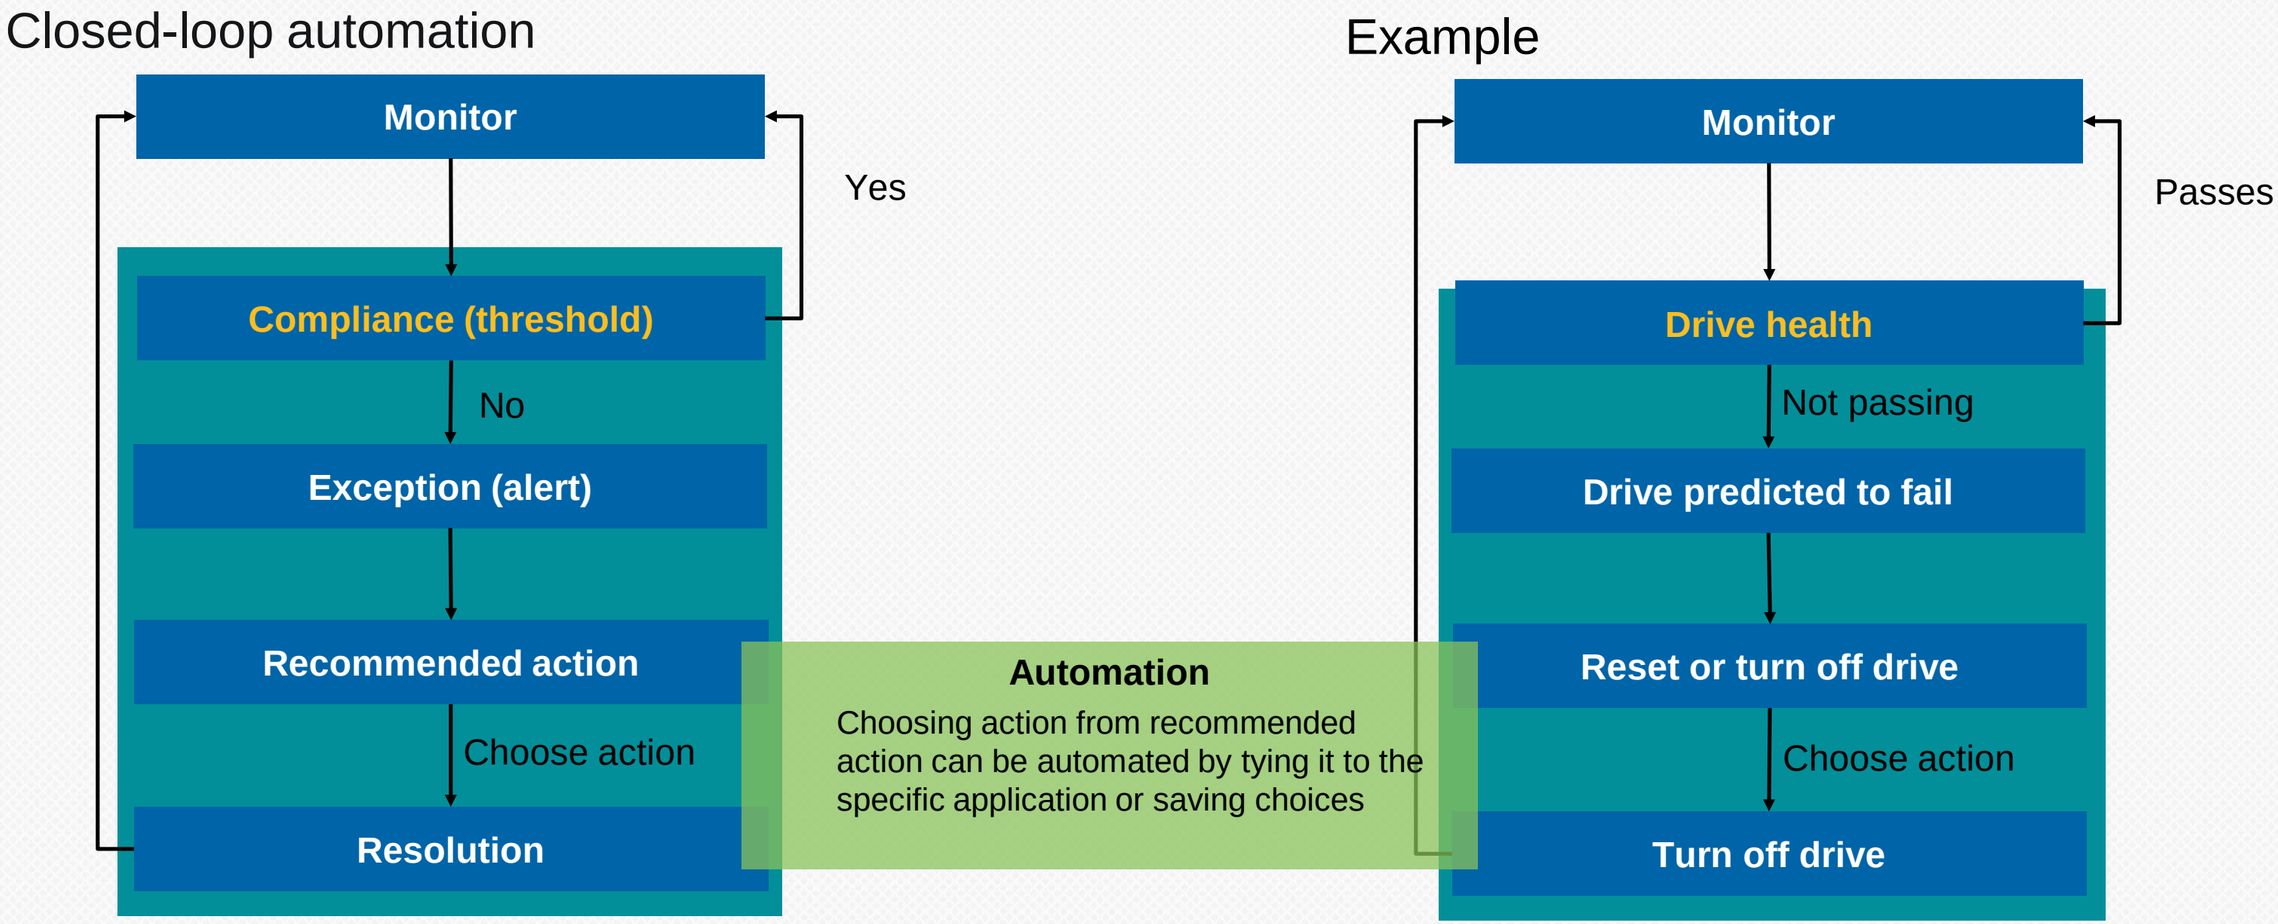

Implementation

Architecture overview

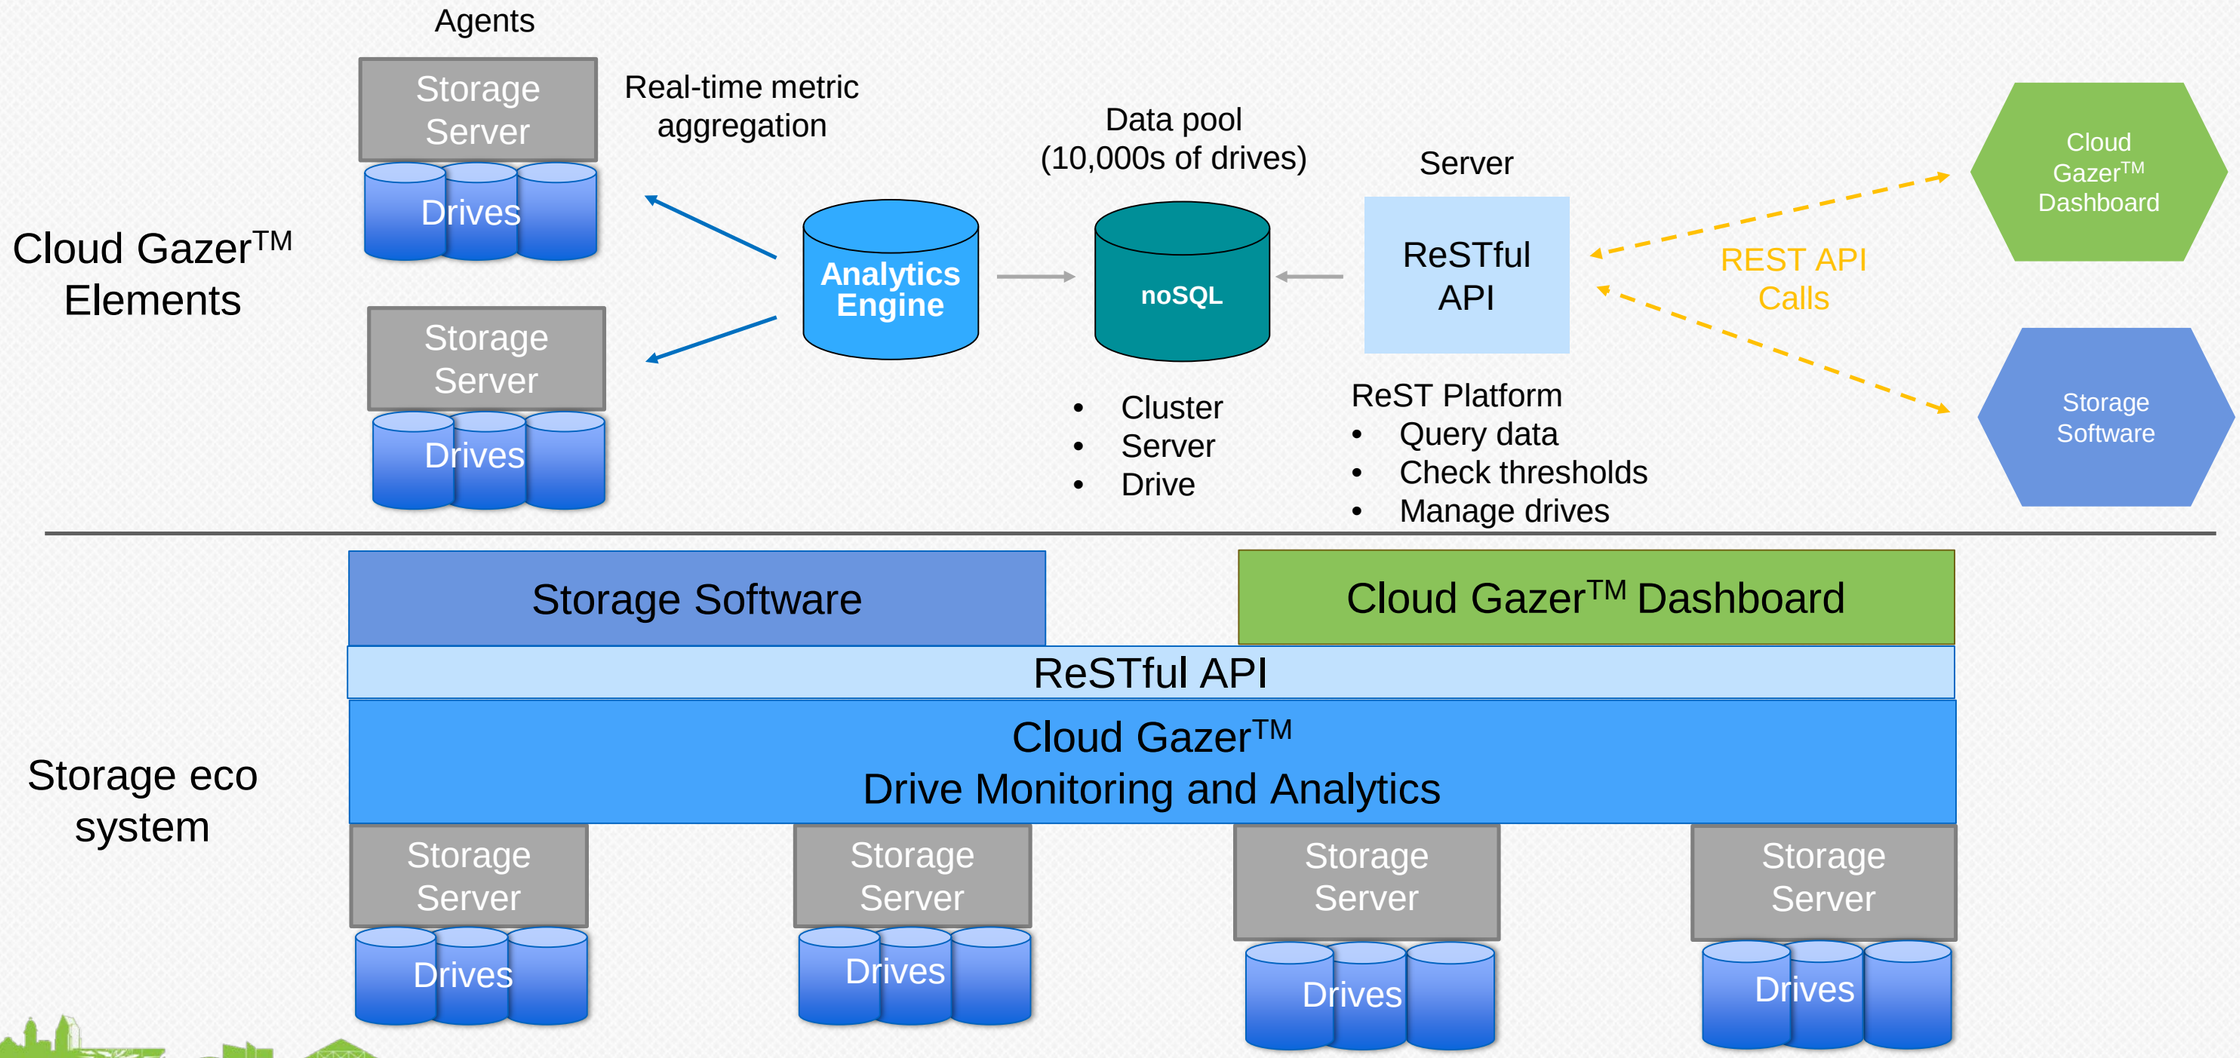

Use cases

Compliance (thresholds)

Degradation and performance warnings

19 overloaded drives: Inspect the workload statistics page, heat map or list view to see if the trend is persistent and affects reliability Hide

Health Tiles (2 minutes ago) Watching: 1 clusters, 1 servers, 1 drives

Failing Drives	Drive Capacity Utilization	Overheated Drives	Overloaded Drives
0%	0%	72%	5%
0 Drives / **402 Drives**	0 Drives / **402 Drives**	290 Drives / **402 Drives**	19 Drives / **402 Drives**

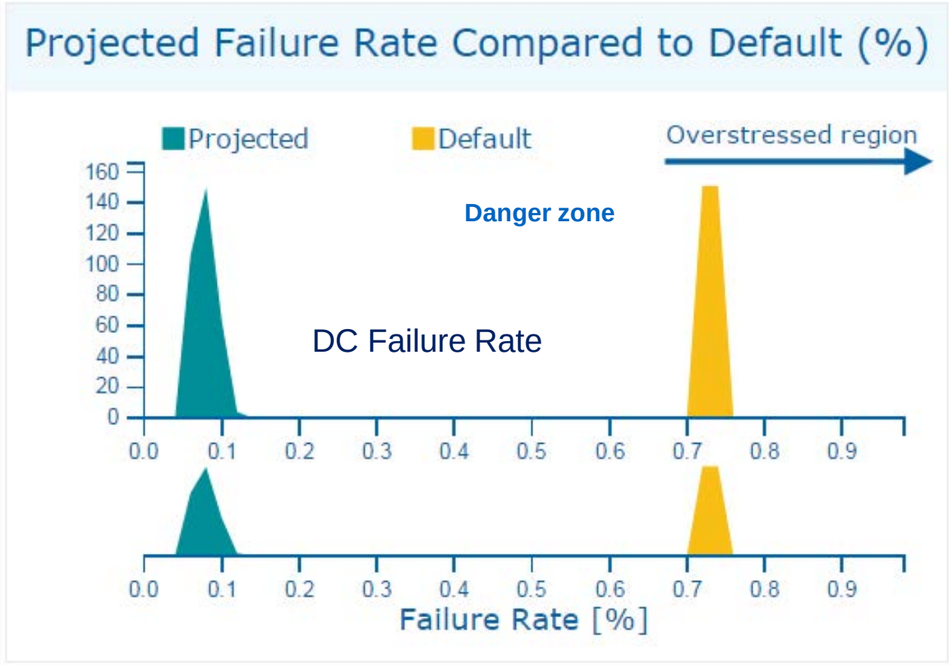

Compliance (thresholds)

Degradation and performance warnings

Overload detection
Detecting and reporting when either drives or network load exceeds design limits

19 overloaded drives: Inspect the workload statistics page, heat map or list view to see if the trend is persistent and affects reliability Hide

Health Tiles (2 minutes ago) Watching: 1 clusters, 1 servers, 1 drives

Failing Drives	Drive Capacity Utilization	Overheated Drives	Overloaded Drives
0%	0%	72%	5%
0 Drives / **402 Drives**	0 Drives / **402 Drives**	290 Drives / **402 Drives**	19 Drives / **402 Drives**

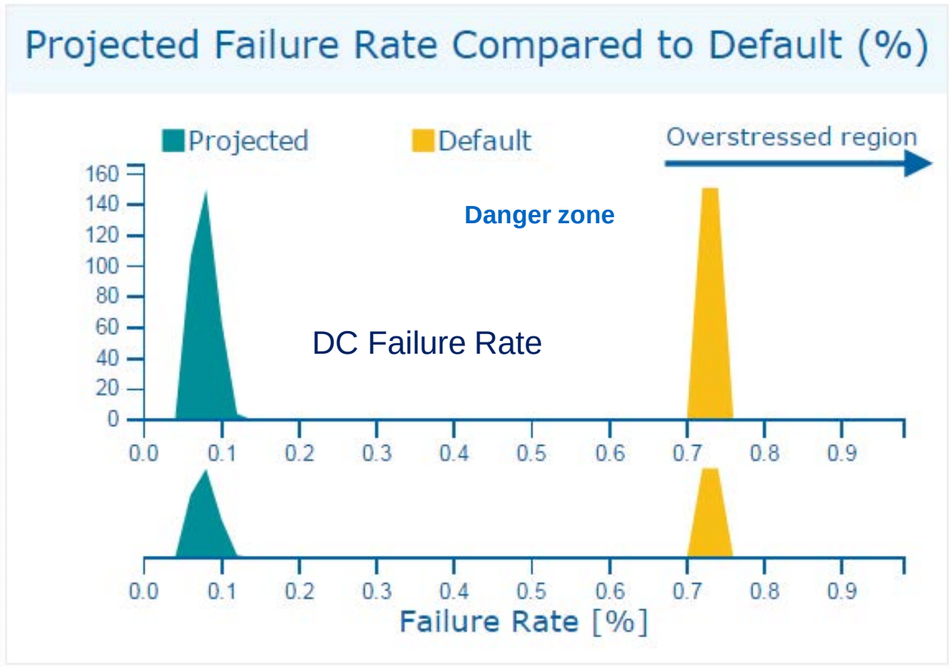

Compliance (thresholds)

Degradation and performance warnings

Recommended action
How to increase drive reliability

Overload detection
Detecting and reporting when either drives or network load exceeds design limits

19 overloaded drives: Inspect the workload statistics page, heat map or list view to see if the trend is persistent and affects reliability

Hide

Health Tiles (2 minutes ago)

Watching: 1 clusters, 1 servers, 1 drives

Failing Drives	Drive Capacity Utilization	Overheated Drives	Overloaded Drives
0%	0%	72%	5%
0 Drives / **402 Drives**	0 Drives / **402 Drives**	290 Drives / **402 Drives**	19 Drives / **402 Drives**

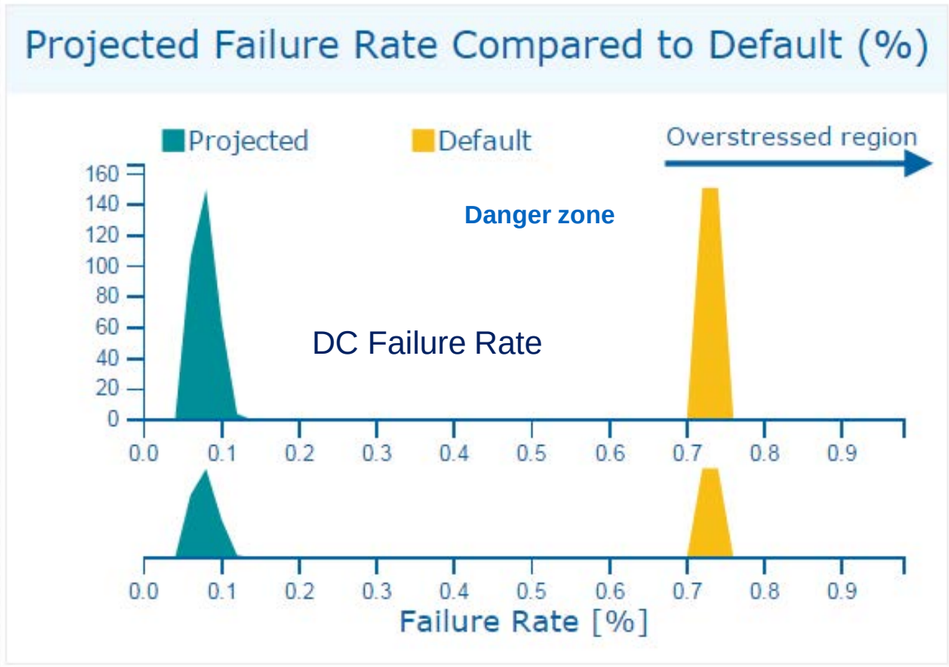

Compliance (thresholds)

Degradation and performance warnings

Recommended action
How to increase drive reliability

Overload detection
Detecting and reporting when either drives or network load exceeds design limits

HDD population failure rate
Measuring stress and estimating failure acceleration of the disk drive population in real time. Relies on the proprietary failure prediction algorithms

19 overloaded drives: Inspect the workload statistics page, heat map or list view to see if the trend is persistent and affects reliability Hide

Health Tiles (2 minutes ago) Watching: 1 clusters, 1 servers, 1 drives

Failing Drives	Drive Capacity Utilization	Overheated Drives	Overloaded Drives
0%	0%	72%	5%
0 Drives / **402 Drives**	0 Drives / **402 Drives**	290 Drives / **402 Drives**	19 Drives / **402 Drives**

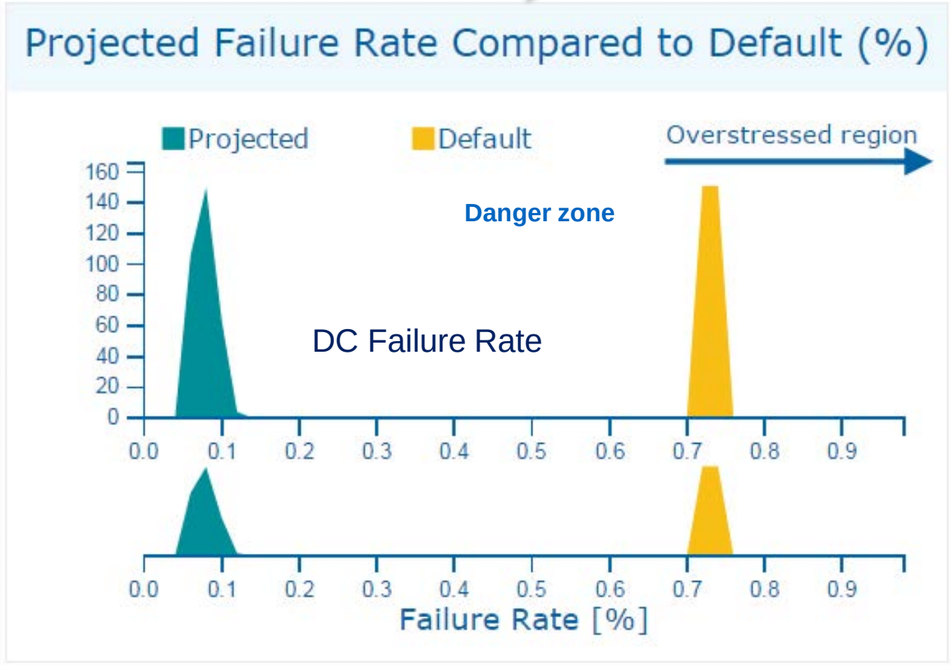

Compliance (thresholds)

Degradation and performance warnings

Failure detection
Warning about expected drive failures. Relies on the proprietary failure prediction algorithms that use unsupervised machine learning techniques. Expected average failure prediction time window is from 9 days to 12 days.

Recommended action
How to increase drive reliability

Overload detection
Detecting and reporting when either drives or network load exceeds design limits

HDD population failure rate
Measuring stress and estimating failure acceleration of the disk drive population in real time. Relies on the proprietary failure prediction algorithms

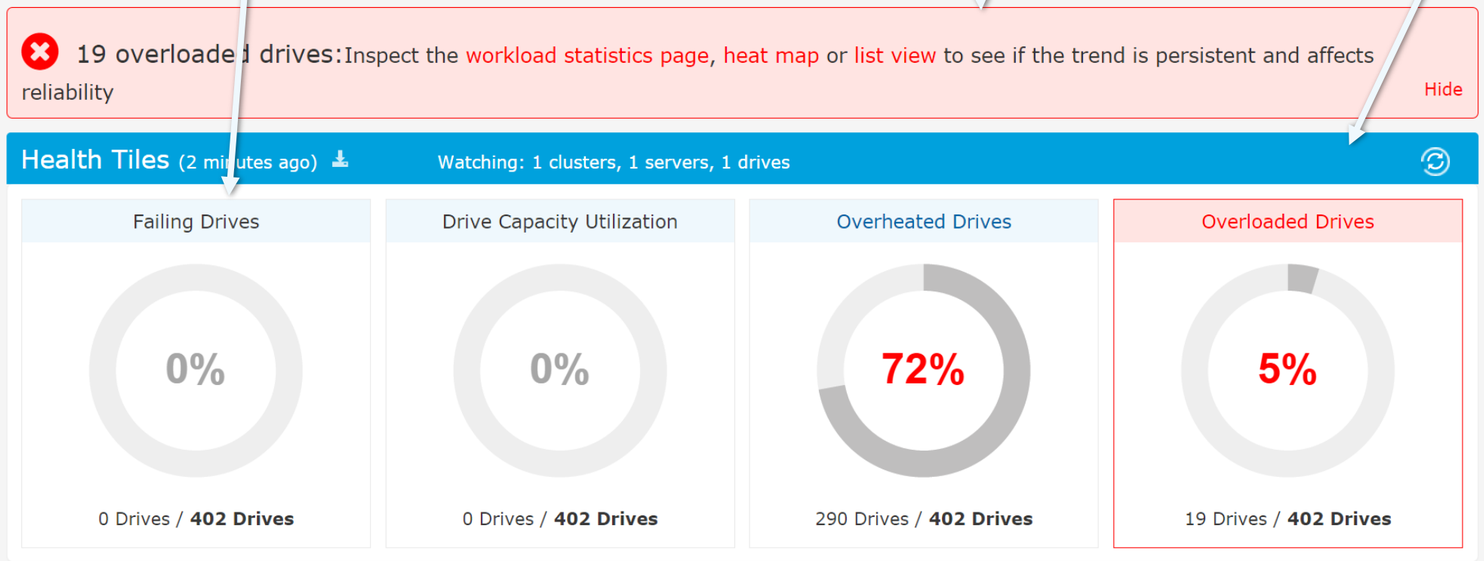

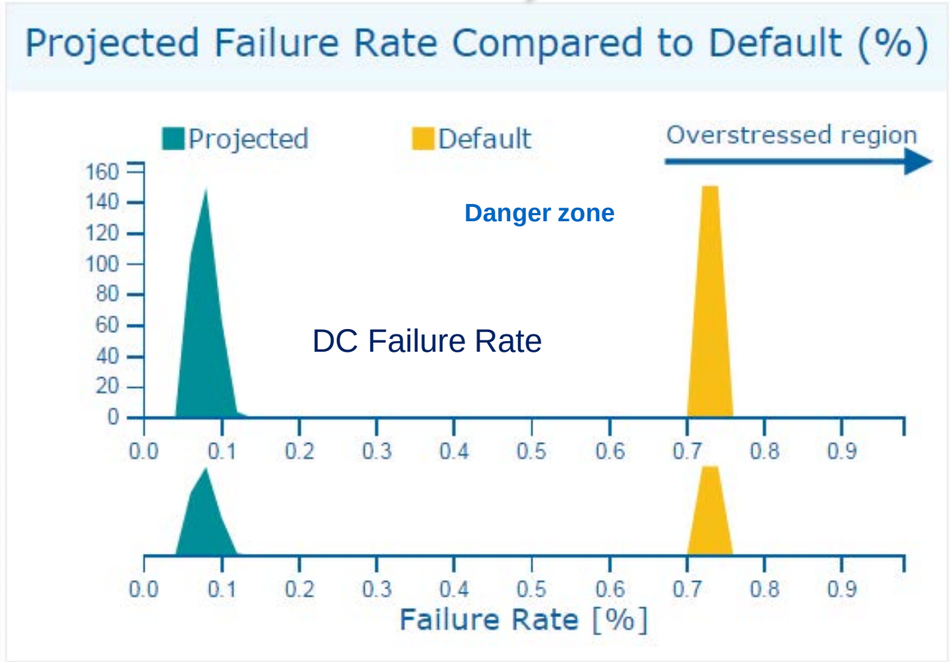

Workload optimization

Drive visibility tools to improve workload balancing

Before Load balancing issues

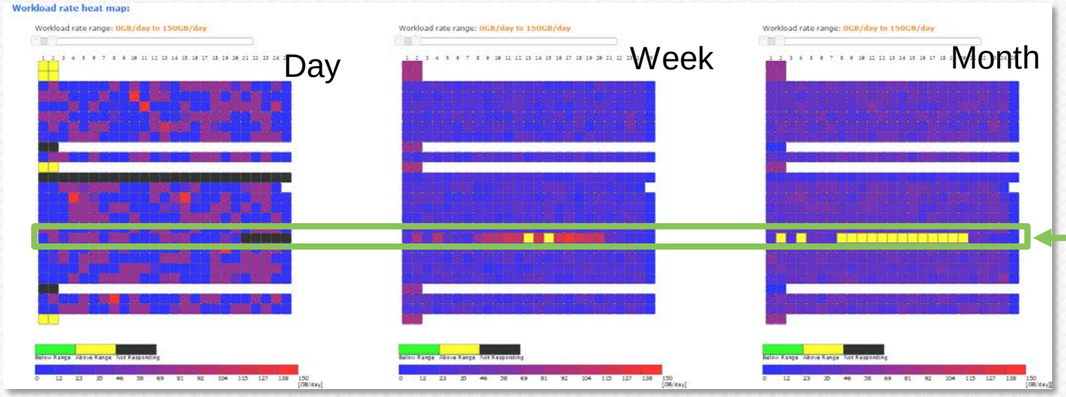

Workload predominantly hitting one server

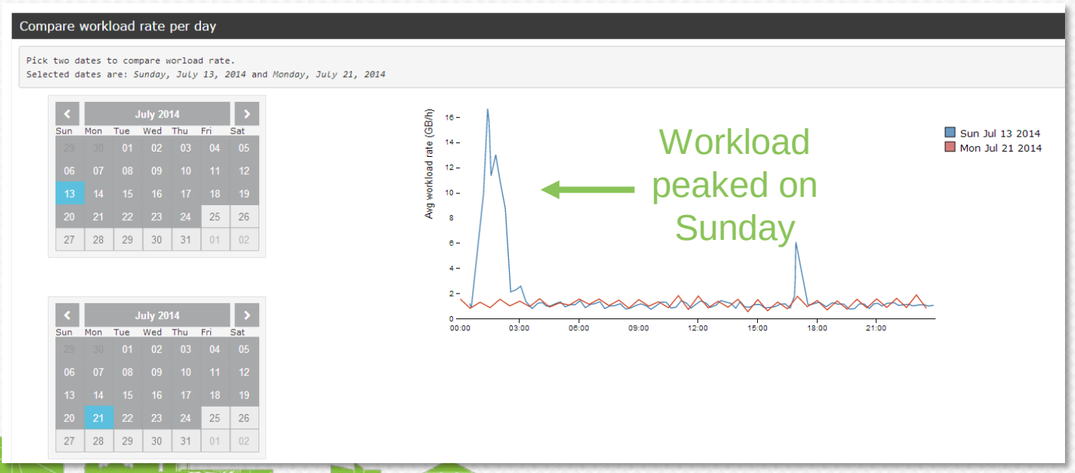

After Workload distributed over servers and time

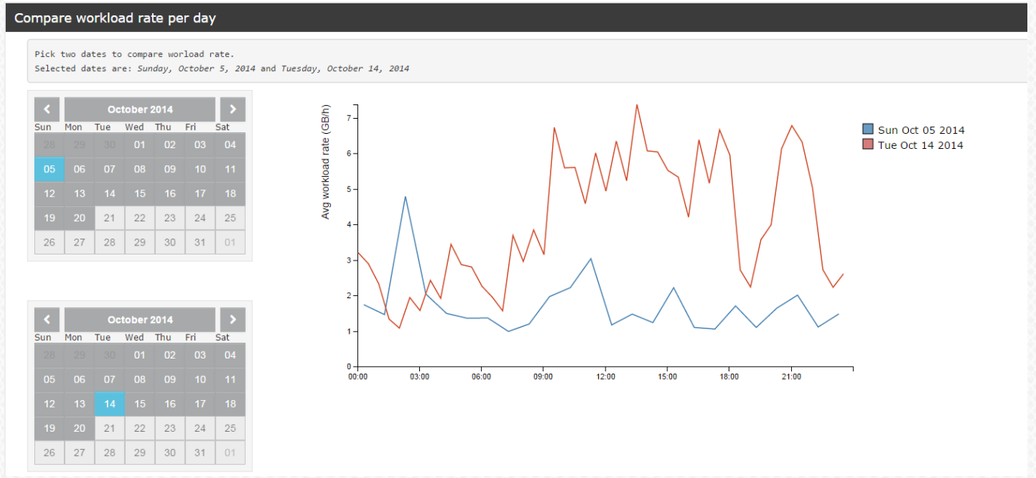

Unsupervised machine learning and failure prediction

No interaction between drive set, no prior knowledge

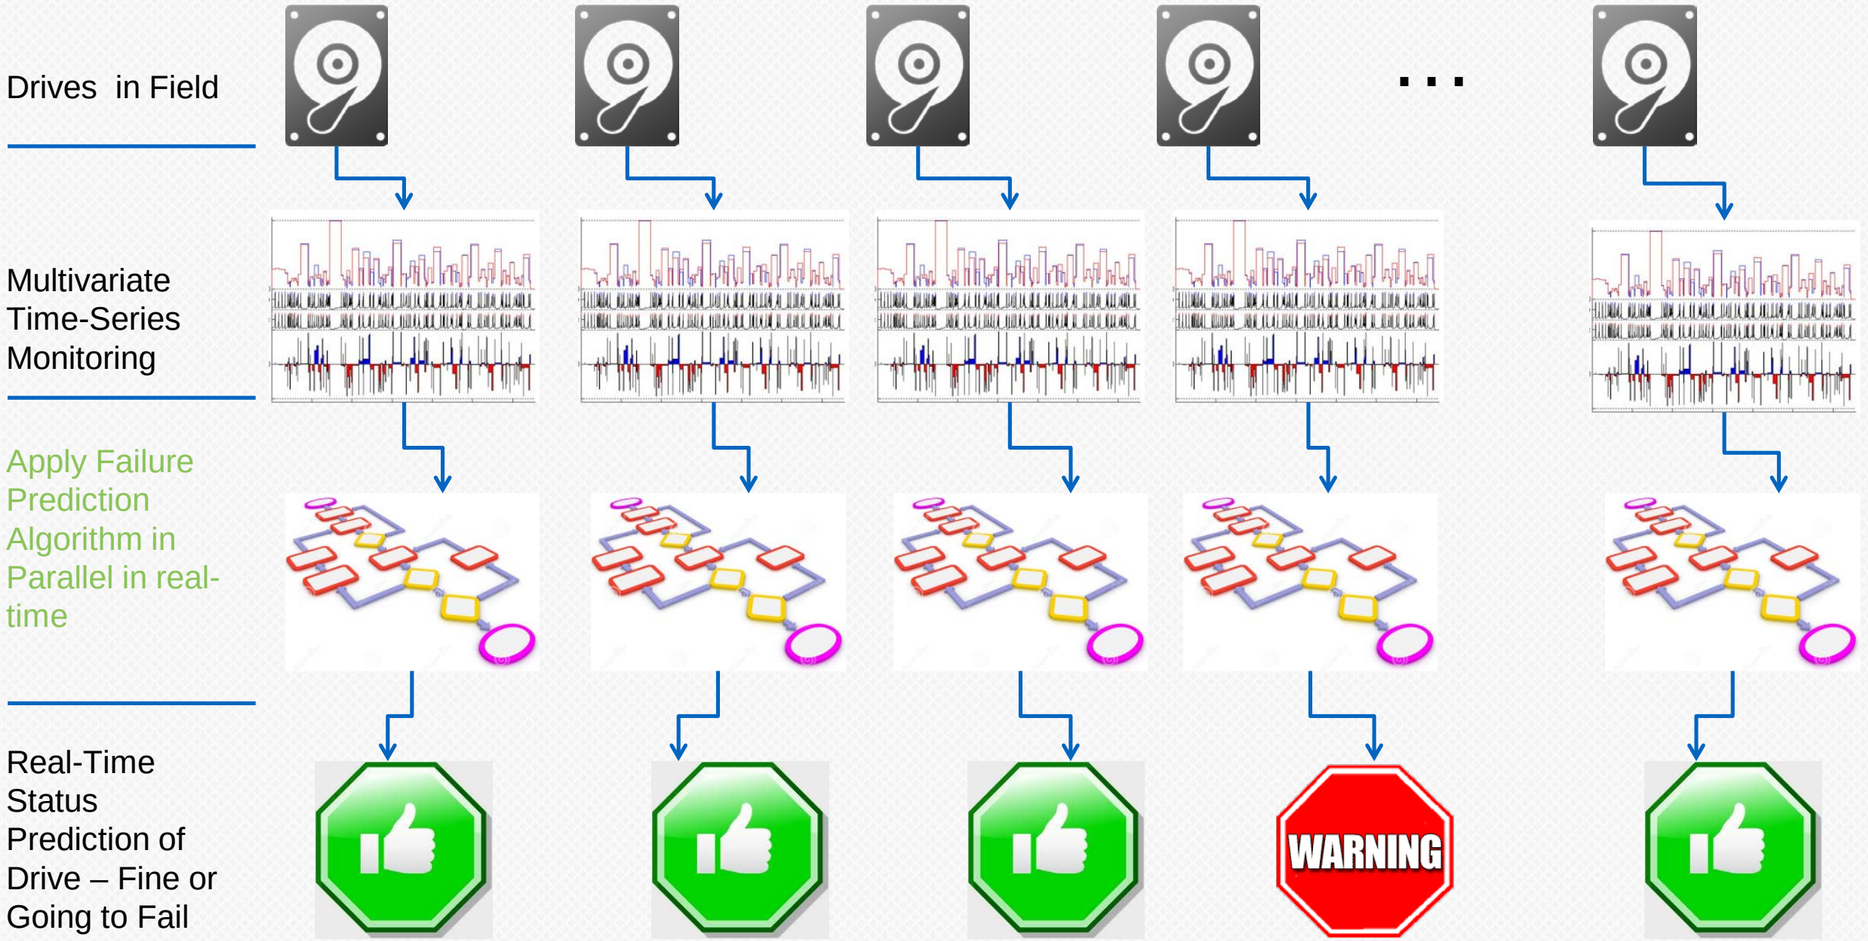

For now, an average failure prediction window is on the order of 9 to 12 days
Failure prediction accuracy ranges from 55% to 90%

Prediction and follow up actions

Heat map indicates drives at risk and you can issue drive tests (DST, IDD,…) to resolve or corroborate

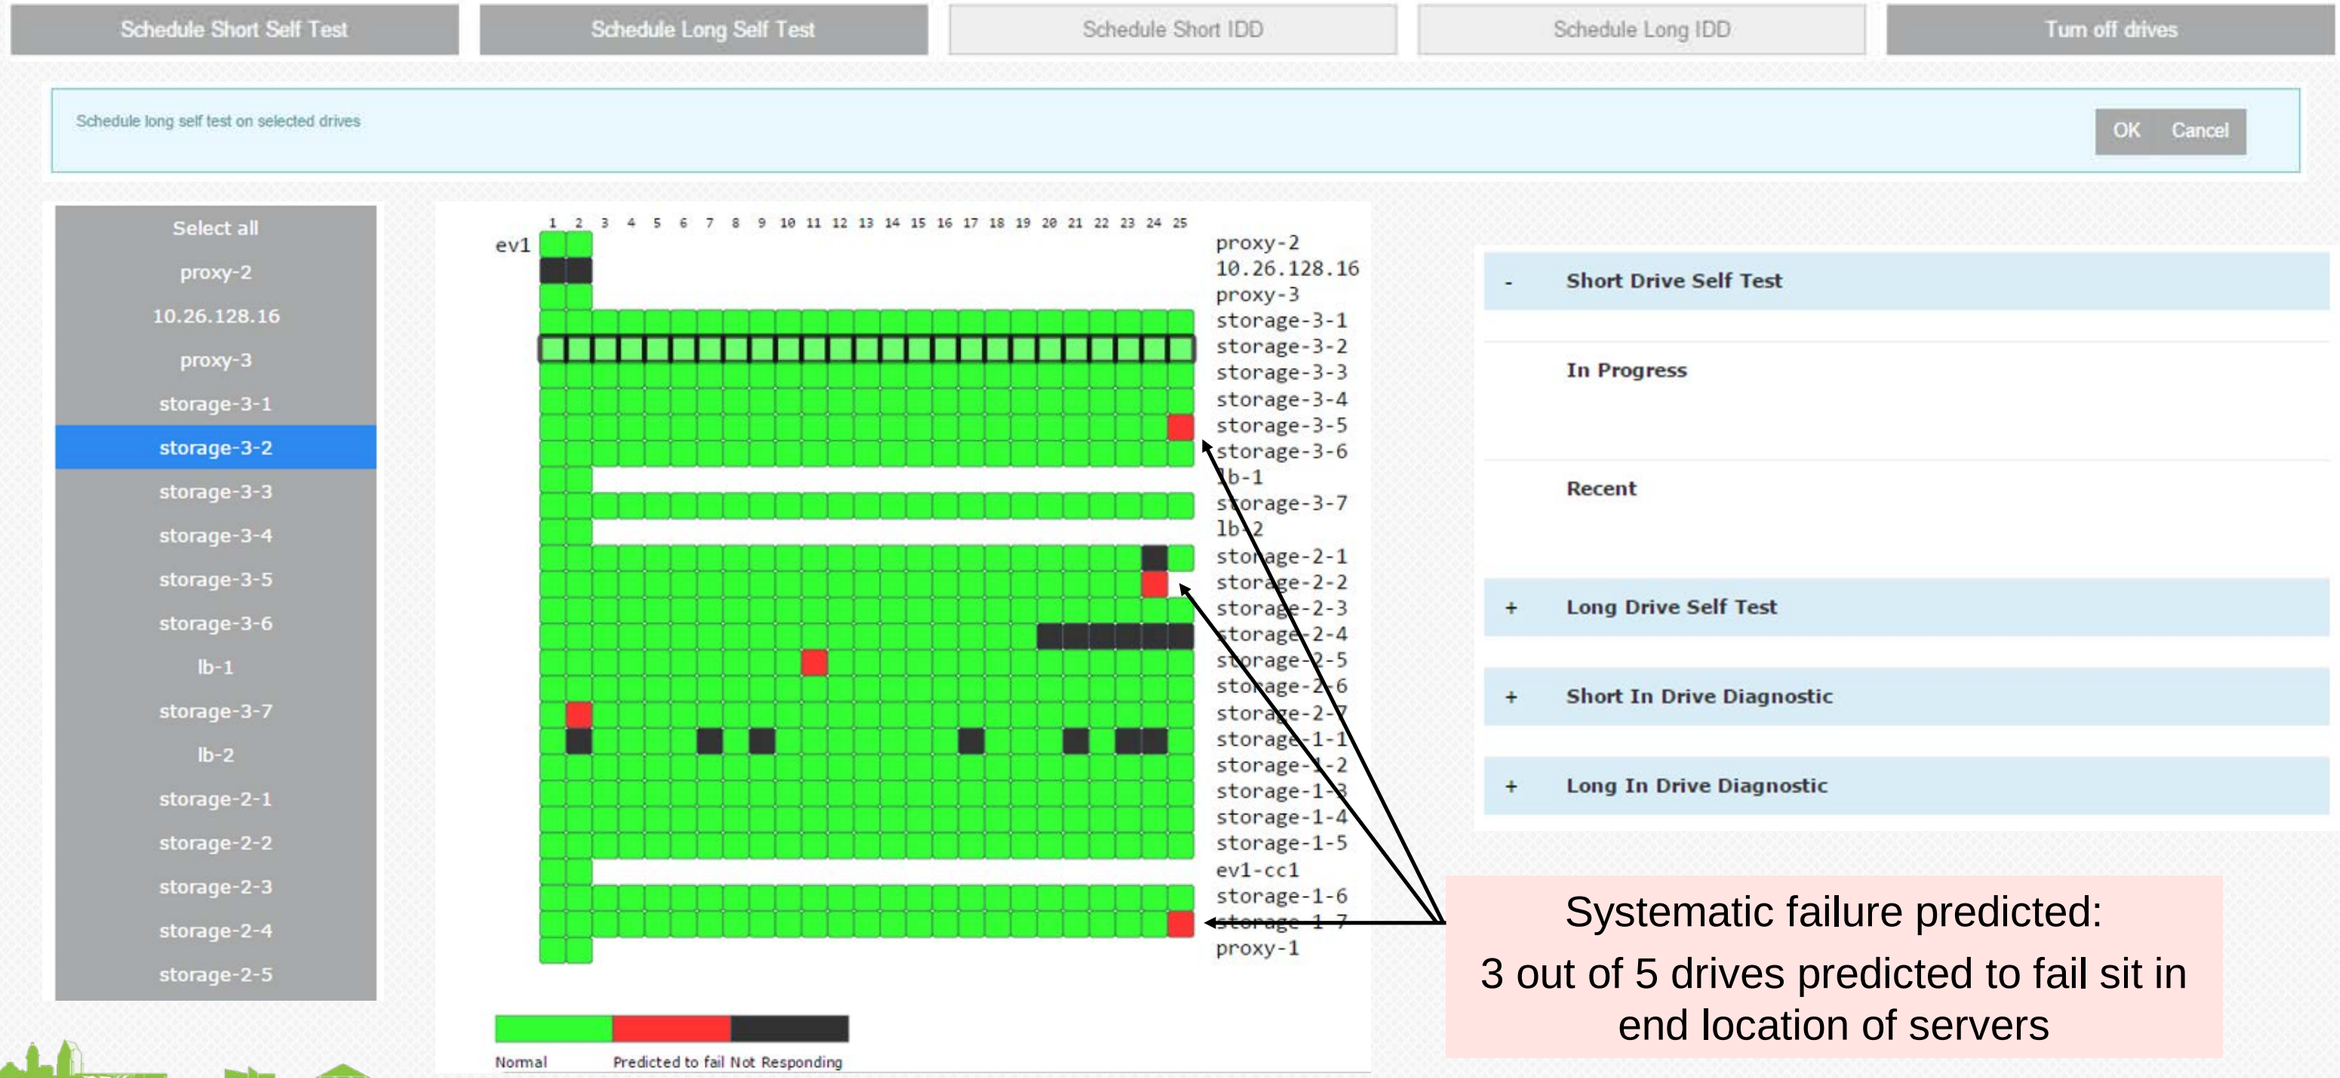

Systematic failure predicted:

3 out of 5 drives predicted to fail sit in end location of servers

Find failure triggers

Root cause tools including a temperature heat map can help you triage the cause of your drive issues

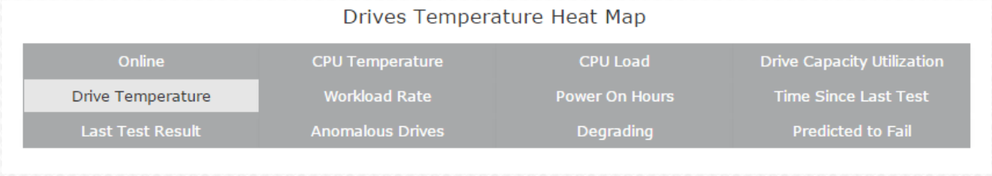

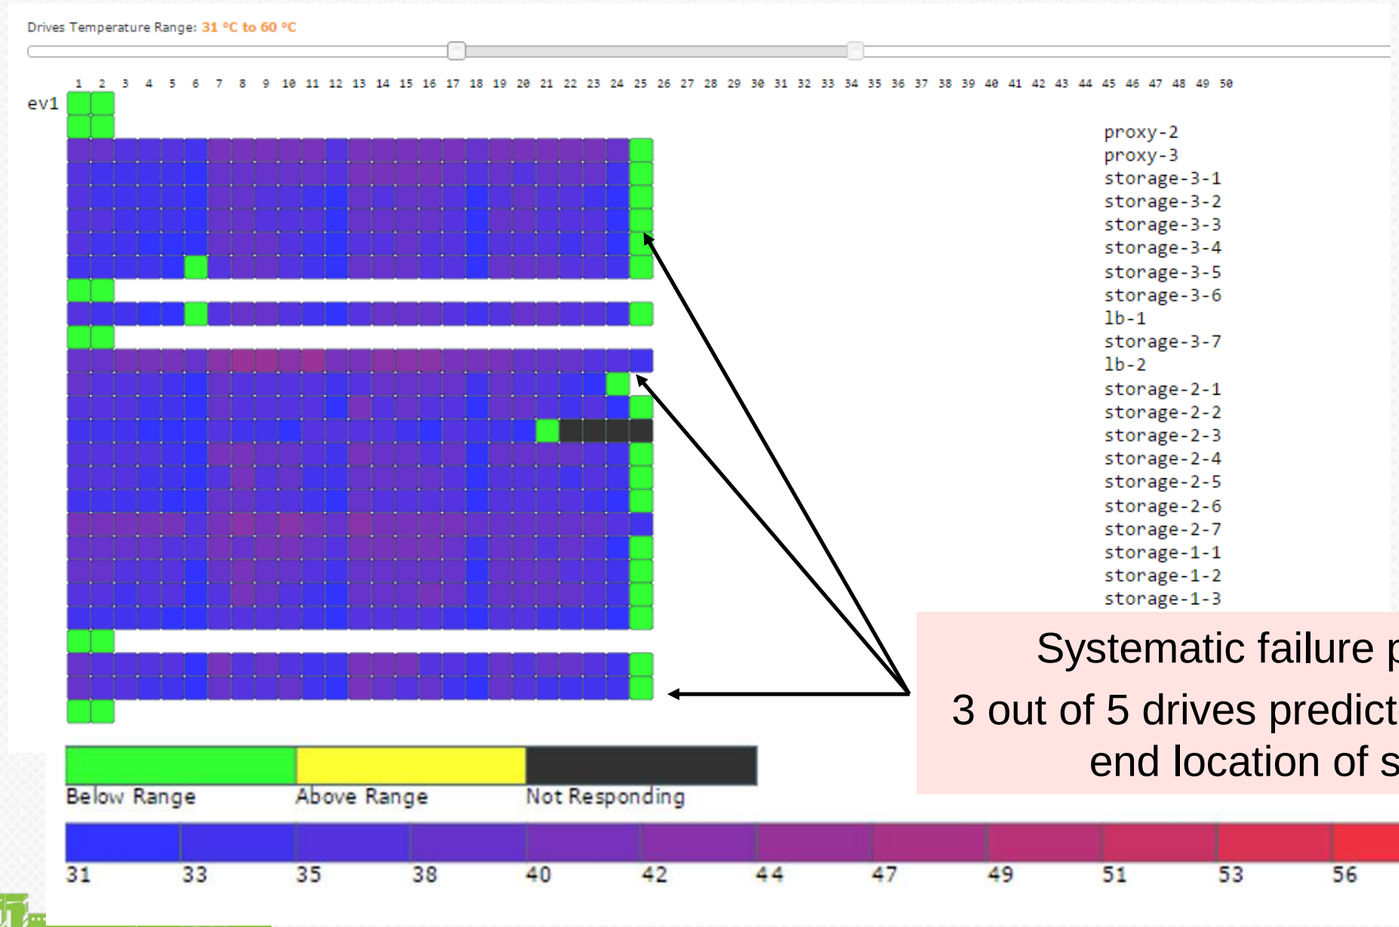

Common factors for drives in the end position is a cooler temperature. Therefore increasing the server temperature may reduce the (dominant) failure mechanism and increase drive reliability

Systematic failure predicted:

3 out of 5 drives predicted to fail sit in end location of servers

Failure prediction lead time

We can predict drives will fail on average 9-10 days before the failure

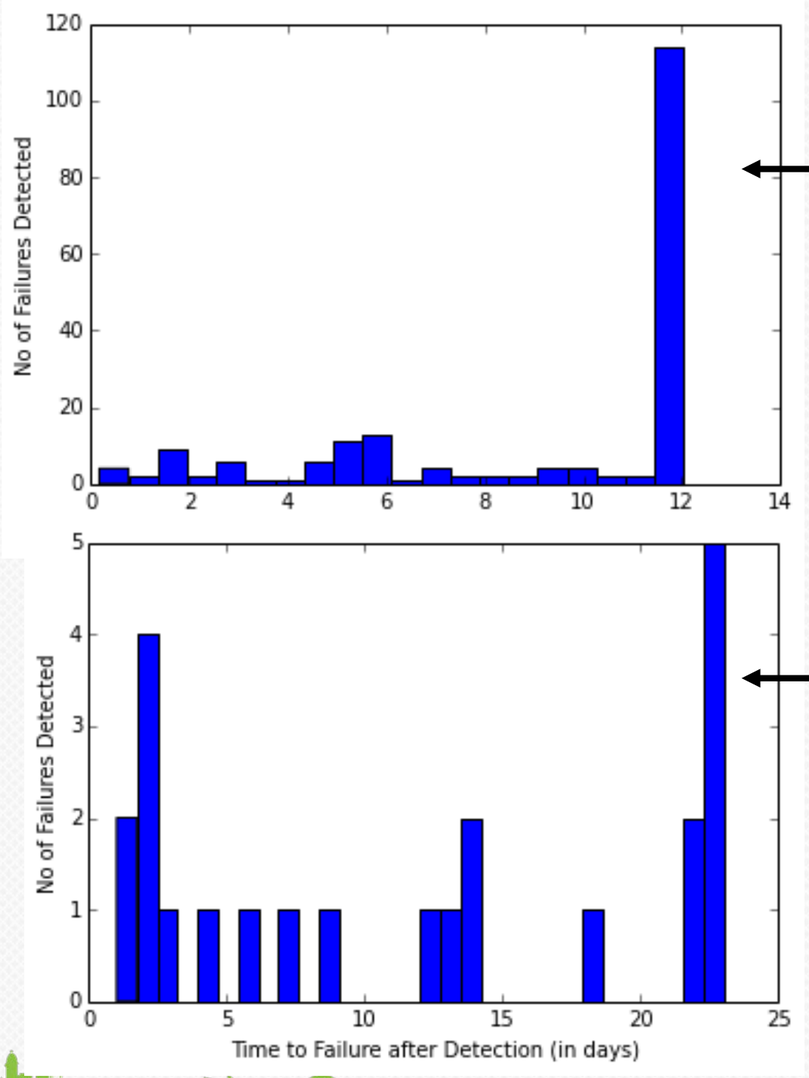

Case study 1, we predicted most drives (118 drives) to fail 12 days prior to failure

Case study 2, we predicted 5 drives to failed 23 days prior to failure, 2 drives prior to failure,… 2 drives just one day in advance

Currently catch 55-90% of failures ahead of time

Summary

Why Cloud Gazer?

- Truly drive-centric management tool for the cloud
- Most efficient tool for extracting drive health information using Seagate IP
  - Nobody knows drives better than us
  - Freeware utilities are frequently wrong
- Runs on any Linux system with little overhead (<1%)

  Windows is next
- Data can be collected, monitored and analyzed locally or in the Cloud
- ReSTful API to interact with other software
- New Analytics, Prediction, AI, and Control capabilities are added continually
- Drive repair will be possible with in-drive diagnostic
- Enclosure control will be possible by summer 2015

Simply SMARTer

	Seagate's CloudGazer*		Competition	
Attribute	SATA	SAS	SATA	SAS
Raw Error Rate	Yes	Yes	Partial	No
Spin-Up Time	Yes	Yes	No	No
Start/Stop Count	Yes	Yes	Yes	No
Retired Sectors Count	Yes	Yes	No	No
Seek Error Rate	Yes	Yes	Partial	No
Power On Hours	Yes	Yes	Partial	No
Power Cycle Count	Yes	Yes	Yes	No
Reported Uncorrectable	Yes	Yes	No	No
Command Timeout	Yes	Yes	Partial	No
High Fly Writes	Yes	Yes	Yes	No
Emergency Retract Count	Yes	Yes	Yes	No
Load/Unload Count	Yes	Yes	Yes	No
Temperature	Yes	Yes	Yes	No
ECC Count	Yes	Yes	No	No
Load Cycle Count	Yes	Yes	Yes	No
Pending Sparing Count	Yes	No	No	No
Head Flight Hours	Yes	No	Partial	No
Lifetime Writes	Yes	Yes	Partial	No
Lifetime Reads	Yes	Yes	Partial	No
Flags1	Yes	Yes	no	No
RV Abs Mean	Yes	Yes	no	No
Motor Power	Yes	Yes	No	No
Critical Event Errors	Yes	Yes	No	No

*Seagate drives

Questions?