

Gluster X

Gluster X


Gluster X (10 in Latin...) is a proposal to plan some features that will be implemented across Gluster versions 8, 9 and 10.

There's nothing special about the number 10, it's just 3 versions away and is a cool number.


Planning ahead allows better risk management, design, implementation flexibility and the ability to break stuff!

**Backward
Compatibility**

Innovation



The 4 key areas we would like to focus on



Performance

Performance enhancements, in both the data path and the control path.



Debugging

Improve maintainability of the project, with structured logging, tracing and instrumentation



Quality

Continue the efforts to improve code quality, using static analysis and code coverage, CI efficiency and accuracy



Scalability

Support more nodes, more volumes, more clients, ...

Network layer optimization

Replace the existing XDR network stack with a modern one, to reduce the network layer overhead. Look at zero-copy and improve other bottlenecks in the area.

Targeted version: 10

Complexity: High

Risk: High

Priority: High

Confidence level: Low

Performance Layer optimization

- Is -l issues
- Other perf xlators bottleneck
- DB specific workloads

Targeted version: 9

Complexity: ?

Risk: ?

Priority: High

Confidence level: ?

Path-based Geo-Replication

See

https://docs.google.com/document/d/1gW5ETQxNiy9tt4uV1ohRH1g5AMmWLtbQYD3QPs_v8Ec/edit?usp=sharing

Targeted version: 8/9

Complexity: ?

Risk: ?

Priority: High

Confidence level: Medium

Global thread pool

Fine-tune the implementation to avoid regression in some workloads. Potentially related is work to reduce contention on io-threads.

Targeted version: 9

Complexity: ?

Risk: ?

Priority: High

Confidence level: Medium

3k-4k volumes support

Release some bottlenecks that prevent Gluster from fully supporting thousands of volumes

Targeted version: 8
Complexity: Medium
Risk: Low
Priority: Medium
Confidence level: High

RIO

Bring RIO to production quality to increase scalability in terms of nodes.

Targeted version: 10

Complexity: Medium

Risk: Medium

Priority: Low

Confidence level: Low

Code Instrumentation / Tracing

Use OpenTracing or SystemTap/dtrace and friends to better find bottlenecks in the system


Targeted version: 10
Complexity: Medium
Risk: Medium
Priority: Low
Confidence level: High

Structured logging

Move to use the implemented structured logging facility across the project.

Targeted version: 10
Complexity: Medium
Risk: Medium
Priority: Low
Confidence level: High

GFID split brain resolution



<TBD>

Targeted version: 9
Complexity: Medium
Risk: Medium
Priority: High
Confidence level: High

Breaking changes are those that in the best case will require a disruptive upgrade and in the worst case scenario, data migration.

While we should strive to avoid them, it's critical to notify the community and perhaps offer LTS version of Gluster 9

