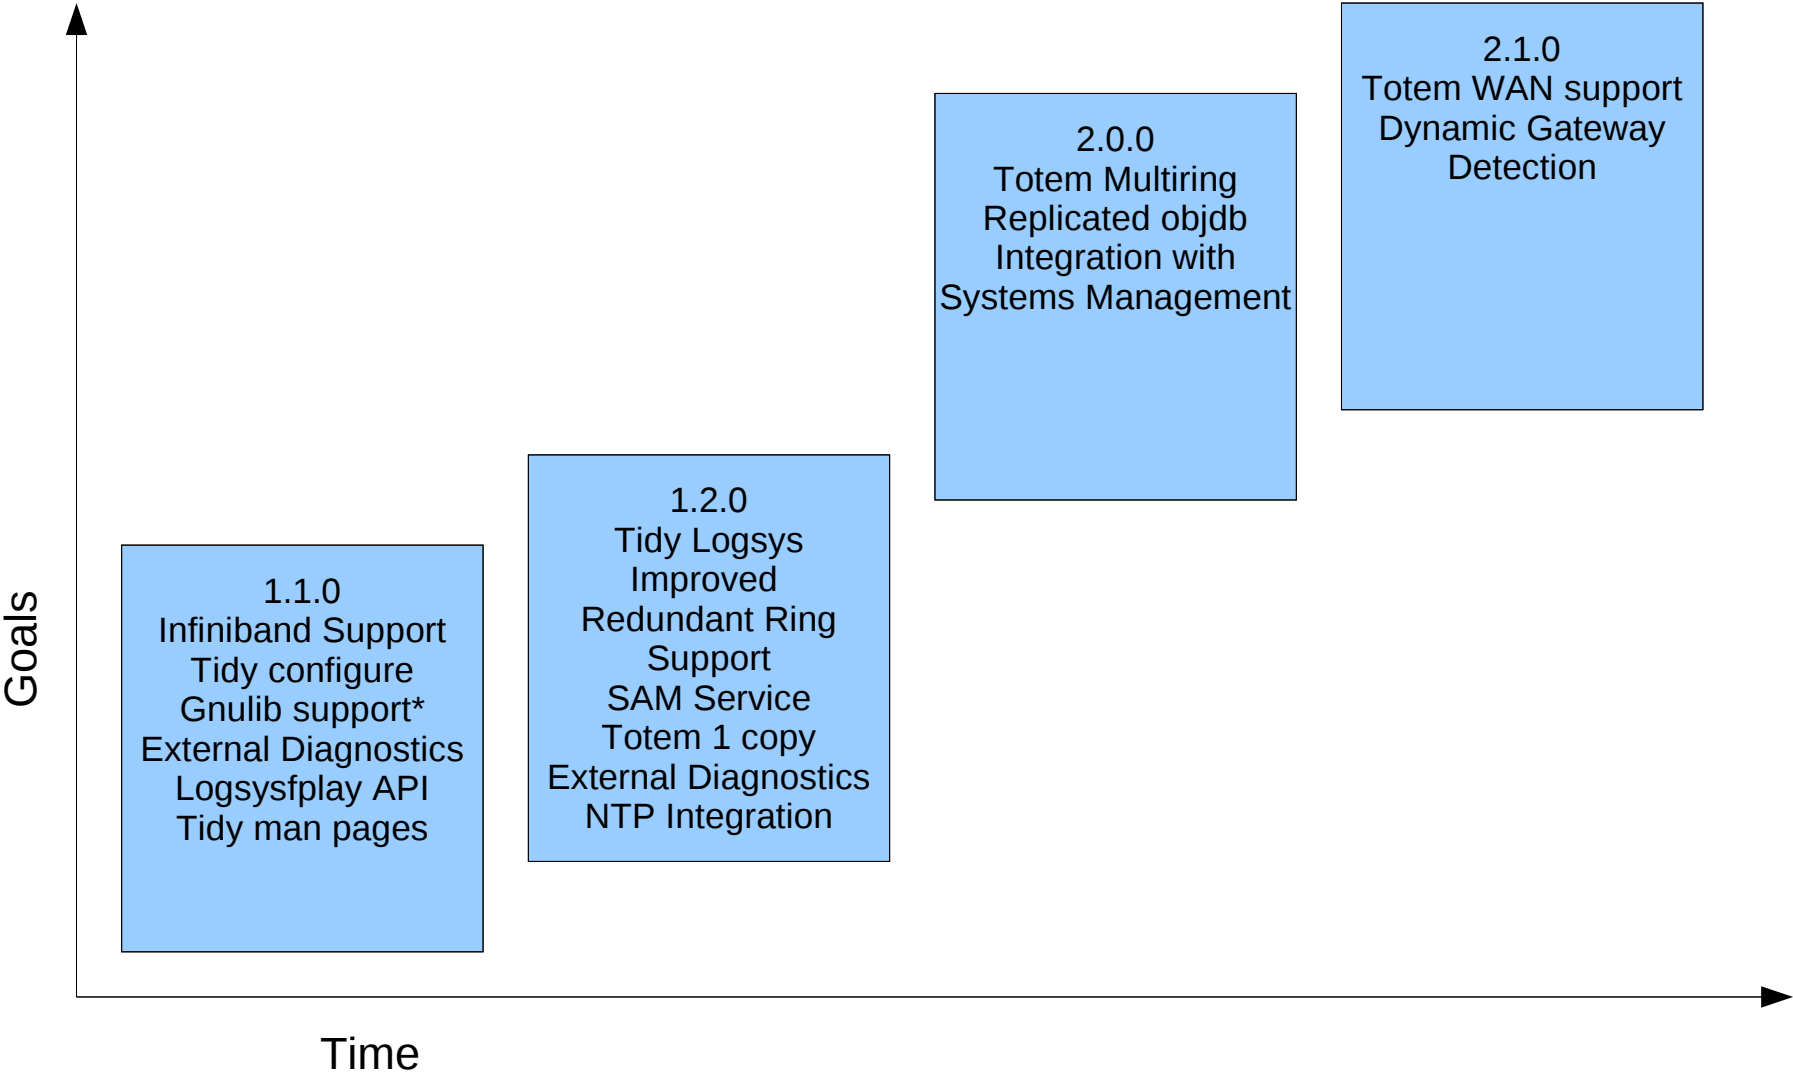


# The Corosync 2009-2010 Roadmap

The mailing list was polled about what features they would like to see in The Corosync Cluster Engine in the future. This roadmap describes a prioritization of those features covered on the mailing list. Like all things open source, our plans are subject to change and directed by individuals who want to solve their most pressing issues. If you share our issues, feel free to contribute in any way you choose.

# The Corosync Cluster Engine Community 2009-2010 Roadmap



\* depending on legal review of potential license issue

# Goal Descriptions

- Infiniband Support
  - Support the Infiniband transport by transmitting over RDMA.
- Tidy Configure
  - The configuration system depends on many platform specific features. Test for those in configure.ac instead of checking the platform on which Corosync run. Move linking to libtool.

# Goal Descriptions

- Integrate gnulib support.
  - The gnulib project provides many portability and correctness checking features. Depends on legal review that including gnulib doesn't compromise the Corosync Modified BSD license.
- External Diagnostics
  - Some basic diagnostic features available as an external CLI application.

# Goal Descriptions

- Logsysfplay C API
  - The logsys system provides a mechanism to create records of information. Unfortunately there is no portable way currently to read this information. The logsysfplay API abstracts the reading of fplay records for other applications.
- Tidy man pages
  - Man pages need some love.

# Goal Descriptions

- Tidy logsys implementation
  - The implementation of the flight recorder backend is difficult to understand. We should also consider how to remove the spinlocks from the front end of the IO path of logsys for 100x+ performance boost.
- Improved Redundant Ring Support
  - The Redundant Ring support in corosync needs more testing, especially around boundary areas such as 0x7FFFFFFFF seqids. Redundant ring should have an automatic way to recover from failures by periodically checking the link and instituting a recovery of the ring.

# Goal Descriptions

- Simple Availability Manager Service
  - This service provides a nonclustered mechanism to monitor CPG processes for failure and restart them if they fault. The model of CPG is that if a process fails, it can be restarted and will execute recovery. This is the missing component to high availability via CPG replication.
- Totem 1 copy
  - Remove as many copies as possible from the Totem stack since totem memcpy() consume 50% of Corosync CPU utilization.

# Goal Descriptions

- External Diagnostics
  - Provide more tools and features to support externally diagnosing internal state of corosync.
- NTP Integration
  - Integration with NTP clock changes.
- Totem Multiring
  - Implement the multiring protocol for better scalability up to hundreds of nodes.



# Goal Descriptions

- Replicated Configuration Object Database
  - Allow our object database to replicate certain shared information, or keep data private if it is node local.
- Integration with Systems Management
  - Integrate with future system management solutions.
- Totem WAN support
  - Allow Totem to operate over high latency links (depends on Multiring).

# Goal Descriptions

- Dynamic gateway detection
  - Allow Totem to dynamically detect its gateways so they don't have to be configured.