WLCG Worldwide LHC Computing Grid

Availability, Reliability and Alternative Forms of Magic

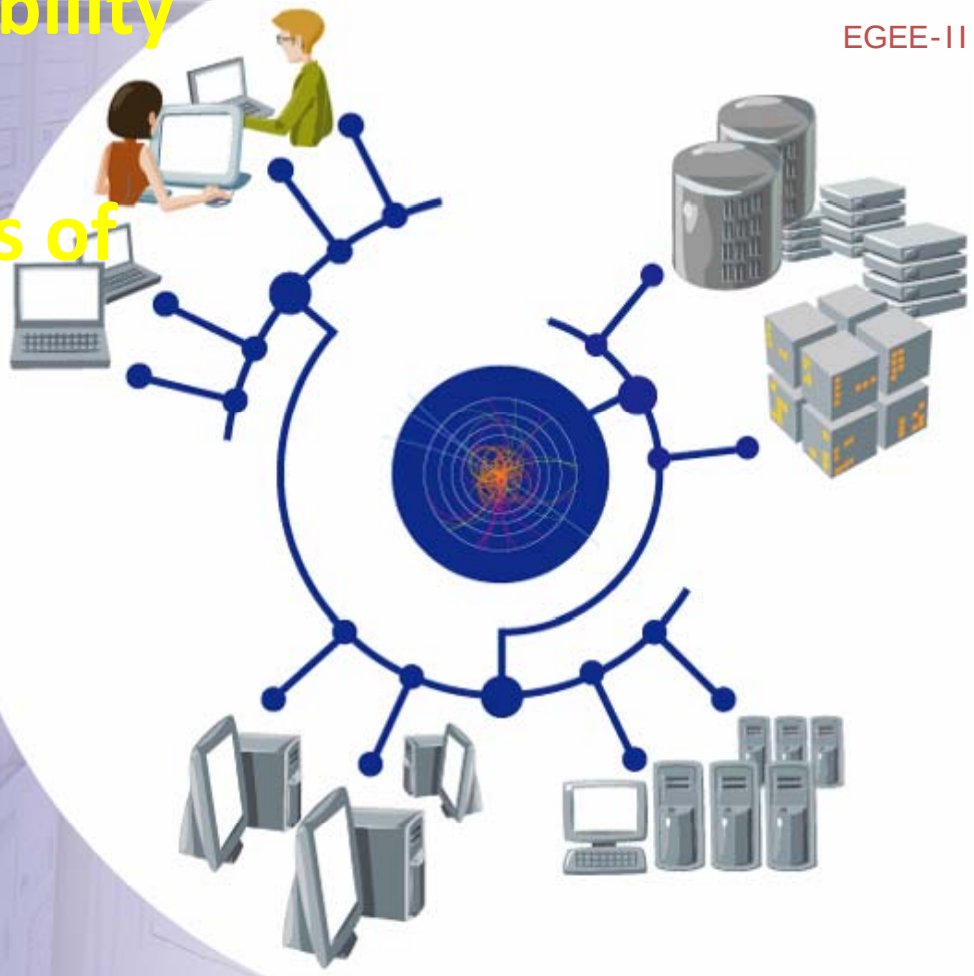

EGEE-II INFSO-RI-031688

Information Society and Media

Markus Schulz
CERN IT-GD-ITR & EGEE SA3
Markus.schulz@cern.ch

Definitions

- Reference: TheDocument

- States: UP, DOWN, SCHEDULED DOWN, UNKNOWN
  - UP means the tests passed
- Availability:
  - Fraction of UP time during a given period
    - 15 minutes UP in a given hour → 0.25
- Reliability: 13th Feb 2007 LCG MB
  - Reliability = Availability / ScheduledAvailability
  - ScheduledAvailability = 1 – ScheduledDownTime – UnknownInterval
  - UnknownInerval = TimeWithUnknownResults/Time
  - UnscheduledDownTime = 1 – Availability - ScheduledDownTime - UnknownInterval

- Reliability = Availability / (Availability + UnscheduledDowntime)

Example

- UP 12 hours, ScheduledDown 6 hours, Unknown 6 hours
  - Availability = 0.5
  - Reliability = 1
- Multiple readings within an hour are ignored
  - Last result before hour ends “wins”
  - 10:04 UP, 10:40 UP, 10:55 DOWN -> 10 – 11 counted as DOWN
- Problem: SAM launched at different intervals → continuous timescale
-
- Current algorithm differs with respect to:
  - a) Service status computation on a continuous time scale
  - b) Consideration of Scheduled Downtime
  - c) Handling of UNKNOWN status
  - d) Validity of Test Results

New versus Old

- Service status computation on continuous time scale
  - Avoids the block effect
- Consideration of Scheduled Downtime
  - Old: A service can pass tests while being SheduledDown in GOCDB
    - This is part of the Availability
    - Leads to confusing Reliability figures
  - New: During ScheduledDown test results are ignored
    - One critical service takes the full site down ←
- Handling the UNKNOWN status
  - New: UP only if all critical tests are known and O.K.
    - GOCDB and BDII are used → Vo <-> service mapping
    - No ghosts!
    - Site state is aggregated from services (more later)
  - Old: Service is UP if at least 1 critical test is O.K.
    - all other can be UNKNOWN
    - Services can drop out of sight

New versus Old

- Validity of Test Results
  - Old : 24h extrapolation for all tests
  - New: Defined by Vos, scheduled interruptions lead to UNKNOWN

- Aggregated state in the current algorithm:

- Service: If at least one instance is available, the service is available

- Site : All critical site services (CE, SE,…) are considered
  - The LOWEST state wins (UP UP UP DOWN == DOWN)
- Global services (WMS, LFC…)
  - UP if at least one instance is UP somewhere

Metrics

- Individual Service Instance Status, Availability and Reliability
- Individual Service Status, Availability and Reliability
- Individual Site Status, Availability and Reliability
- Aggregate T1/0 Availability
- Central Services Status, Availability and Reliability

- Availability is computed for each VO every hour
  - based on continuous state
- All other quantities are derived from this

- States:
  - Service Instance:
    - ScheduledDown (GOCDB, tests are ignored)
    - DontCare (no critical tests),
    - UP, DOWN, (All critical tests O.K) (at least one failed)
    - UNKNOWN, (All critical tests that have run are O.K. >=1 not available)

Metrics

- States:
  - Service :
    - ScheduledDown
      - No instance UP or DontCare, one or more ScheduledDown
    - DontCare (no critical tests defined),
    - UP (at least one instance is UP)
    - DOWN
      - >=1 DOWN and none UP or DontCare or ScheduledDown
    - UNKNOWN
      - All instances are in state UNKNOWN
  - Site:
    - All site level services (GOCDB or BDII)
    - DOWN (>=1 one service is down)
    - ScheduledDown (No service is DOWN, >=1 is ScheduledDown)
    - UNKNOWN
      - No service is DOWN or ScheduledDown, >=1 in UNKNOWN
    - UP (All services are UP or DONTCARE)

Availability & Reliability

- Hourly Availability
  - Service Instance, Service, Site :
    - DontCare as good ad UP
- Hourly Reliability
  - Service Instance, Service, Site
    - From continuous states

- Daily, Weekly, Monthly Availability, Scheduled Downtime and Unknown Interval
  - Based on hourly values

- T1/0 based on the average of all sites

Critical Tests

- All can be browsed here:
  - Critical tests
- Different tests for NDGF , EGEE/WLCG, OSG
- Not clear how much this reflects the current usage patterns and middleware
- Many are “timeless”
  - Brokerinfo, CA certs, environment tests, basic job submission….
- Datamanagement tests are using: InfoSystem, WN config, LFC, SRM-V1
  - “Orthogonalization” needed
- SiteBDII tests done in gstat
  - Laurence, Felix and Minh are reworking (improving them)
  - Tests are already quite deep
  - We need to add the same tests for the top level BDIIs
    - They are a critical service
    - We have 70 of them……