Summary

Currently it is impossible to compare data for exact periods of time, for example, compare average of data in June vs July and alert user in case of significant difference. It makes monitoring based on time shifts rather limited and not user friendly. Besides there are performance related issues since Zabbix always use historical data for calculation of trigger expressions.

It would be great if Zabbix supported this functionality. It would enable generation of valuable insights and also significantly improve business value of Zabbix.

Zabbix acceptance

1. A new set of statistical functions must be introduced for analysis of data for absolute periods of time like "previous month" (say, June), "this year", "this day one month ago", etc
    a. New functions: avg_period, count_period, delta_period, max_period, min_period, sum_period
        i. These functions will accept same parameters as existing functions without the "_period" prefix except:
            1. Parameters "sec|#num" and "time_shift" of the new functions will be renamed to "period" and "period_shift" respectively
                a. period: time period with time suffix starting from 1h. It must support suffixes d (hour), d (day), w (week), M (month) and y (year).
                b. period_shift: must support syntax similar to one used for time navigator in Zabbix UI. Some examples:
                    i. Average for current hour: avg_period(1h, now/h). If current time is 12:35 then now/h references to 13:00:00. Zabbix will calculate avg for 1h, i.e. 12:00-13:00 (including data for 12:00 and excluding data for 13:00)
                    ii. Average for prev hour: avg_period(1h, now/h-1h)
                    iii. Average for hour before prev hour: avg(1h, now/h-2h)
                    iv. Average for the same month year ago: avg(1M, now/M - 1y)
2. Zabbix must use only trend data for new functions
    a. Trend cache will be introduced
        i. Zabbix must stay up and running in case if trend cache is fully utilized
        ii. New metrics for internal monitoring must be introduced
        iii. Trend cache must keep one record if data is selected for a period of time for efficiency. For example: when calculating avg_period(1M, now/M-M) only one record with avg,min,max and count of values will be stored for previous month
    b. Trigger expressions having only new functions must be calculated only once per minimum used period
        i. Daily functions must be recalculated daily, monthly functions - monthly, etc
        ii. Zabbix must not skip recalculation of such trigger expressions in case of restart of Zabbix Server
        iii. Zabbix must spread calculation of such trigger expression across some period of time
            1. For example: if some function (like avg(1M, now/M - 1y)) needs to be re-calculated monthly then it must be re-calculated once we have all trend data in place for the previous month within one day (or N hours?)
3. Triggers must support new optional text attribute "Event name"
    a. If event name is given then its value will be used as event name, otherwise trigger name will be used (as it is now)
    b. Event name must support trigger expression formulas using syntax {?formula} and all other macros already supported in trigger name
        i. New function fmtnum(digits) will be supported to optionally round and format numeric output of the formula
            1. digits: number of digits after decimal point. No trailing zeros will be produced.
        ii. For example:
            1. Load of Exchange server increased on {?100*{{HOST.HOST}:load.avg(1M,now/M-1M)}/{{HOST.HOST}:system.cpu.load[percpu].avg(1M,now/M-2M)}}%
                a. Output: Load of Exchange server increased on 24.3413523%
            2. Load of Exchange server increased on {{?100*{{HOST.HOST}:load.avg(1M,now/M-1M)}/{{HOST.HOST}:system.cpu.load[percpu].avg(1M,now/M-2M)}}}.fmtnum(2)}%
                a. Output: Load of Exchange server increased on 24.34%

    c. Event name will also support all macros already supported in trigger name, also macro {TIME}
4. Existing macro {TIME} will support formatting using fmttime(format, time_shift) function
    a. format: optional format string, "%H%M%S". It should match strftime() basic formatting, no need to support advanced formatting options if there are any
    b. time_shift: optional time shift. No time shift by default.
        i. For example: {{TIME}.fmttime(%B,-1M)}. Result: July
    c. The macro and formatting function will use locale of Zabbix Server

Example

1. Trigger definition
    a. **Name:** Load of Exchange server increased on more than 10% last month
    b. **Event name:** Load of Exchange server increased on {{?100*{{HOST.HOST}:load.avg(1M,now/M-1M)}/{{HOST.HOST}:system.cpu.load[percpu].avg(1M,now/M-2M)}}.fmtnum(0)}% in {{TIME}.fmttime(%B,-1M)} ({{?{{HOST.HOST}:system.cpu.load[percpu].avg(1M,now/M-1M)}}.fmtnum(2)}) comparing to {{TIME}.fmttime(%B,-2M)} ({{?{{HOST.HOST}:system.cpu.load[percpu].avg(1M,now/M-2M)}}.fmtnum(2)})
    c. **Expression:** {exchange:system.cpu.load[percpu].avg_period(1M,now/M-1M)}>1.1*{exchange:system.cpu.load[percpu].avg_period(1M,now/M-2M)}
    d. **Allow manual close:** [X]
2. As a result we will start getting monthly alerts like: **Load of Exchange server increased on 24% in July (0.69) comparing to June (0.56)**

Nonfunctional requirements

1. N/A

Use cases

1. I want Zabbix to send me insights like "Load of Exchange server increased on 24% in July (0.69) comparing to June (0.56)"

Decisions made

1. Zabbix will automatically schedule execution of trend based functions
2. Do not use calculated items
    a. Hard to manage, we need to create extra item which is not connected to parent
    b. Unnecessary history
    c. It will be hard to auto-create such triggers globally
3. Event name must fully describe what happened with real data included. it must be super user friendly and readable

Open questions

1. How to deal with formatting for {?formula}. I may want to get percentages without decimal part.